Changing Inflation Dynamics, Evolving Monetary Policy

Gonzalo Castex, Jordi Galí,
and Diego Saravia
editors

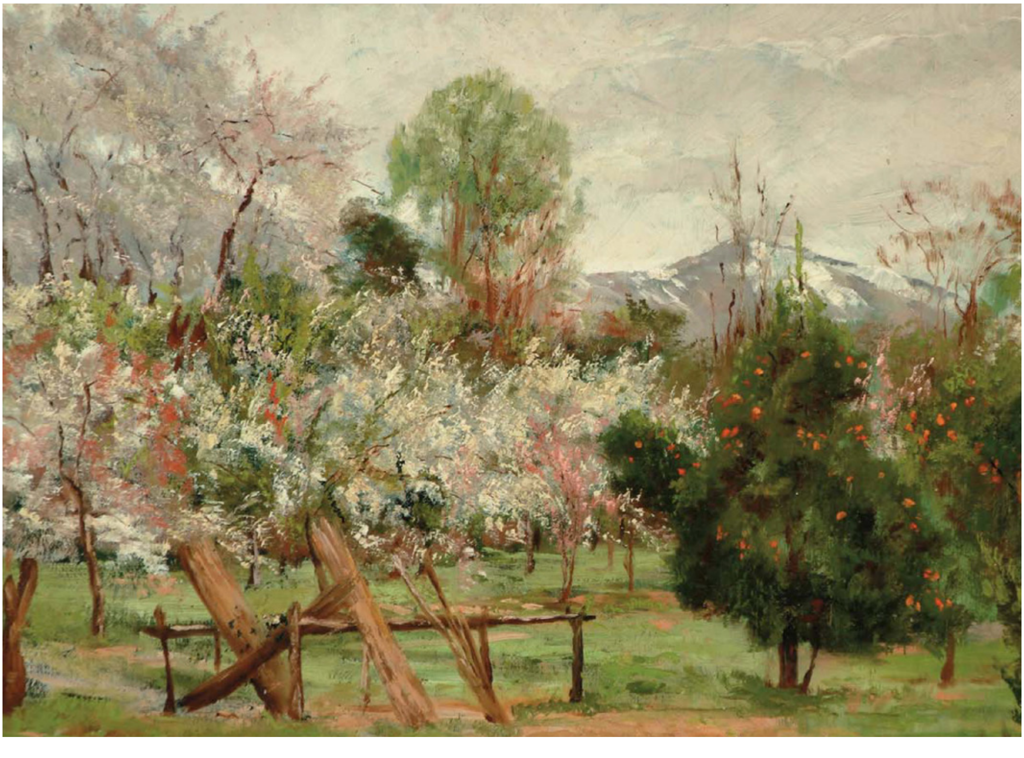

Banco Central de Chile / Central Bank of Chile

Changing Inflation Dynamics, Evolving Monetary Policy

Empirical models have failed to explain inflation behavior over the last 20 years in most developed economies. The unusual inflation dynamics—the 'missing deflation' during recessions and the 'missing inflation' during recoveries—points to a failure of Phillips curve predictions. Several hypotheses have been proposed to explain the 'twin puzzle' phenomenon while at the same time have imposed challenging implications to conduct monetary policy. It is of utmost importance to understand the challenges for monetary policy conduct in an environment where inflation dynamics is hard to unravel. This volume contributes to the study of the 'twin puzzle' phenomenon and the challenges facing monetary policy. It gathers a selective group of distinguished scholars and policy makers to discuss the latest academic findings on inflation dynamics.

Primavera
Eugenio Guzmán Ovalle
Oil on canvas, 63.5 x 40 cm
Collection of the Central Bank of Chile

Changing Inflation Dynamics, Evolving Monetary Policy

Gonzalo Castex
Jordi Galí
Diego Saravia
Editors

Central Bank of Chile / Banco Central de Chile

Series on Central Banking, Analysis, and Economic Policies

The Book Series on "Central Banking, Analysis, and Economic Policies" of the Central Bank of Chile publishes new research on central banking and economics in general, with special emphasis on issues and fields that are relevant to economic policies in developing economies. The volumes are published in Spanish or English. Policy usefulness, high-quality research, and relevance to Chile and other economies are the main criteria for publishing books. Most research in this Series has been conducted in or sponsored by the Central Bank of Chile.

Book manuscripts are submitted to the Series editors for a review process with active participation by outside referees. The Series editors submit manuscripts for final approval to the Editorial Board of the Series and to the Board of the Central Bank of Chile. Publication both in paper and electronic format.

The views and conclusions presented in the book are exclusively those of the authors and do not necessarily reflect the position of the Central Bank of Chile or its Board Members.

Changing Inflation Dynamics, Evolving Monetary Policy

Gonzalo Castex
Jordi Galí
Diego Saravia
Editors

Central Bank of Chile / Banco Central de Chile

Agustinas 1180
Santiago, Chile

Published in Santiago, Chile by the Central Bank of Chile
Manufactured in Chile

ISBN (print) 978-956-7421-67-1
ISBN (digital) 978-956-7421-68-8
Intellectual Property Registration 2020-A-2881

ISSN 0717-6686 (Series on Central Banking, Analysis, and Economic Policies)

Production Team

Editors:
Gonzalo Castex
Jordi Galí
Diego Saravia

Copy Editor:
María Marta Semberoiz

Designer:
Maru Mazzini

Proof Reader:
Dionisio Vio

Printer:
Andros Impresores

Contributors

The articles in this volume are revised versions of the papers presented at the Twenty-second Annual Conference of the Central Bank of Chile on Changing Inflation Dynamics, Evolving Monetary Policy, held in Santiago on 25-26 October 2018. The list of contributing authors and conference discussants follows.

Contributing Authors

Fernando Álvarez
University of Chicago
NBER

Laurence Ball
Johns Hopkins University

Elena Bobeica
European Central Bank

Gonzalo Castex
University of New South Wales

Matteo Ciccarelli
European Central Bank

Jordi Galí
CREI – Universitat Pompeu Fabra – Barcelona GSE

Luca Gambetti
Collegio Carlo Alberto, Università di Torino, UAB Barcelona GSE

Simon Gilchrist
New York University
NBER

Bart Hobijn
Arizona State University

Òscar Jordà
Federal Reserve Bank of San Francisco
University of California, Davis

Sandeep Mazumder
Wake Forest University

Fernanda Nechio
Federal Reserve Bank of San Francisco

Pablo Andrés Neumeyer
Universidad Torcuato Di Tella

Diego Saravia
Central Bank of Chile
Proficio Investment, Argentina

James H. Stock
Harvard University

Isabel Vansteenkiste
European Central Bank

Mark W. Watson
Princeton University

Egon Zakrajšek
Federal Reserve Board

Conference Discussants

Gonzalo Castex
University of New South Wales

Luca Gambetti
Collegio Carlo Alberto, Università di Torino, UAB Barcelona GSE

David López-Salido
Board of Governors of the Federal Reserve System

Fernanda Nechio
Federal Reserve Bank of San Francisco

Ernesto Pastén
Central Bank of Chile

Argia Sbordone
Federal Reserve Bank of New York

Carlos Viana de Carvalho
Banco Central do Brasil

Eduardo Zilberman
Central Bank of Chile

Table of Contents

Changing Inflation Dynamics, Evolving Monetary Policy: An Overview
Gonzalo Castex, Jordi Galí, and Diego Saravia 1

The Passthrough of Large-cost Shocks in an Inflationary Economy
Fernando Álvarez and Pablo Andrés Neumeyer 7

The Nonpuzzling Behavior of Median Inflation
Laurence Ball and Sandeep Mazumder 49

The Link between Labor Cost Inflation and Price Inflation in the Euro Area
Elena Bobeica, Matteo Ciccarelli, and Isabel Vansteenkiste 71

Has the U.S. Wage Phillips Curve Flattened?
Jordi Galí and Luca Gambetti 149

Trade Exposure and the Evolution of Inflation Dynamics
Simon Gilchrist and Egon Zakrajšek 173

The Supply-Side Origins of U.S. Inflation
Bart Hobijn 227

Inflation Globally
Òscar Jordà and Fernanda Nechio 269

Trend, Seasonal, and Sectorial Inflation in the Euro Area
James H. Stock and Mark W. Watson 317

Changing Inflation Dynamics, Evolving Monetary Policy: An Overview

Gonzalo Castex
University of New South Wales

Jordi Galí
CREI - Universitat Pompeu Fabra- Barcelona GSE

Diego Saravia*
Proficio Investment, Argentina

Understanding the dynamics of inflation has become an important challenge for both policymakers and researchers over the past decade. Empirical models linking inflation and economic activity—versions of the so-called Phillips curve—have failed to account for the behavior of inflation in many advanced economies. In particular, inflation in the U.S. and other countries was higher during the 2008-2009 Great Recession than the conventional empirical Phillips curve would imply. As noted by some economists, this "missing deflation" phenomenon may have already started in the mid-2000s. Just as puzzling, during the subsequent recovery, inflation has remained subdued relative to the predictions generated by existing models, despite the aggressive expansionary monetary policies implemented in many advanced economies.

Economists have labeled the previous developments the "twin puzzle". A number of hypotheses to explain these unusual inflation dynamics have been put forward, with significant implications for the conduct of monetary policy. The XXII Annual Conference of the Central Bank of Chile gathered several researchers and policymakers to discuss and analyze the causes and consequences of these changing inflation dynamics, their potential policy implications, and the challenges they represent for central banks.

* At the time of producing this book, Mr. Saravia was Economic Research Manager at the Central Bank of Chile.

The eight papers presented and discussed at the XXII Annual Conference can be found in the present volume. Their content spans a wide range of issues regarding inflation and its changing relationship with economic activity. Several articles document the weakened relationship between real activity and prices. Some authors propose alternative measures of inflation for which the link with real activity seems restored. Other articles explore the strength of the Phillips relationship in specific sectors and countries, and conditional on the nature of the shocks driving economic fluctuations.

Next we provide a brief summary of the papers included in this volume.

In "The Passthrough of Large Cost Shocks in an Inflationary Economy," Fernando Álvarez and Andy Neumeyer analyze several episodes involving large changes in the nominal price of inputs in Argentina over 2012–2018 by using microprice data for the city of Buenos Aires. They focus on input-price changes resulting from large changes in regulated prices or exchange rates. They find a high short-term pass-through to prices. They compare the observed price dynamics to the predictions of a menu-cost model of price setting, where firms face both idiosyncratic and aggregate cost shocks. They show that the evidence and theory can be reconciled if both large shocks and a high underlying inflation are assumed. By contrast, the authors argue, neither flexible-price models nor models with time-dependent price setting can be easily reconciled with the evidence.

In "The Nonpuzzling Behavior of Median Inflation," Laurence Ball and Sandeep Mazumder analyze the performance of the U.S. Phillips curve since the Great Recession of 2008–2009, with a special focus on the 2017–2018 period. The authors propose an alternative measure of inflation for which there is no sign of breakdown in the Phillips curve relationship. Their proposed measure is the weighted median of industry inflation rates, after excluding food and energy sectors. This measure is argued to filter out large relative-price changes unrelated to aggregate forces. It is also less volatile than traditional core inflation. And most importantly, it displays a stronger, largely unbroken relationship with the unemployment rate.

In "The Link between Labor Cost Inflation and Price Inflation in the Euro Area," Elena Bobeica, Matteo Ciccarelli, and Isabel Vanteenkiste document the strong relationship between price inflation and labor costs in Europe. Their analysis focuses on different economic sectors (construction, manufacture, services) in four main economies (Germany, France, Italy, and Spain) by using quarterly data from 1981Q1 to 2018Q1. The authors estimate a structural vector autoregression (VAR) to understand the high frequency relationship between labor costs and price inflation. They document that this relationship depends on the state the economy and the type of shock that the economy is subject to. The interpretation of their finding is related to the cost-push/price-markup view of the inflationary process. They find that the passthrough is highest in the construction sector in France, services in Germany and Italy, and manufacturing in Spain. Their findings shed light on the circumstances under which labor costs are the main driver of inflation.

In "Has the U.S. Wage Phillips Curve Flattened? A Semi-Structural Exploration," Jordi Galí and Luca Gambetti start by documenting the decline in recent years in the estimated slope coefficient of a reduced-form wage Phillips curve for the U.S. economy, as well as the shrinking role of lagged price inflation in the determination of wage inflation. They provide estimates of a conditional wage Phillips curve, based on a structural decomposition of wage, price, and unemployment data generated by a VAR with time-varying coefficients, identified by a combination of long-run and sign restrictions. Their estimates show that the key qualitative findings from the unconditional reduced-form regressions also emerge in the conditional evidence, thus suggesting that they are not entirely driven by endogeneity problems or possible changes over time in the importance of wage-markup shocks. The conditional evidence, however, suggests that actual changes in the slope of the wage Phillips curve may not have been as large as implied by the unconditional estimates.

In "Trade Exposure and the Evolution of Inflation Dynamics," Simon Gilchrist and Egon Zakrajšek analyze the potential role played by globalization as a factor behind the weakening of the link between price inflation and economic activity. They use a panel of industry-level data for the U.S. economy, with information on prices, wages, output, and employment. Their data allows them to exploit cross-sectional heterogeneity and to control for aggregate dimensions for inflation and economic activity. They focus on comovements between inflation and measures of resource utilization driven by disturbances to the financial intermediation process, a specific form of aggregate demand shocks. Their analysis points to a significant effect of international trade exposure on the responsiveness of inflation to economic activity at the industry level, with the Phillips curve slope coefficient being about three times larger for low trade intensity industries as compared with their high trade intensity counterparts.

In "The Supply-Side Origins of U.S. Inflation," Bart Hobijn argues that the weak Phillips relationship observed in recent years can be explained by the coexistence of demand and supply shocks. He argues that the monetary policy transmission remains valid once we allow monetary policy to affect also the short-run aggregate supply. The author uses growth-accounting techniques to decompose the sources of U.S. inflation at different horizons. His analysis suggests that nearly half of the variance of inflation is driven by changes in the price of imports, with oil being one of the most important factors. An important message of the paper is that policymakers have to think beyond the need to stabilize aggregate demand in order to avoid fluctuations in inflation.

In "Inflation Globally," Òscar Jordà and Fernanda Nechio address two important issues. First, they seek to understand the global trends of inflation after the financial crisis. Secondly, they assess whether tighter credit conditions affect inflation. The authors use a long panel database (20 years, quarterly data) including 45 countries (advanced and emerging economies). The identification of the Phillips curve is not trivial—if output gap is correlated to supply shocks, the estimation may be affected by simultaneity bias. To deal with this issue, the authors adopt an instrumental-variable approach by using the Germany or the U.S. interest rate as an instrument for the corresponding interest rate in countries with an exchange-rate peg. As the authors note, observed changes in the Phillips curve may be spuriously attributed to the crisis even if they have a different origin. To assess this possibility, Jordà and Nechio pursue a diff-in-diff approach by using, as a treated group, a list of countries that were affected by the crisis and, as a control group, a list of countries that were not affected by the crisis. They find that inflation has declined globally, while at the same time is now more forward-looking. They provide mixed evidence on the hypothesis that the missing deflation was caused by firms facing credit constrains being forced to raise prices when demand was low. They document a gradual change in the Phillips curve, with an increasing weight on expected inflation and a declining weight on backward-looking terms. That development seems to affect all economies, regardless of their exposure to the financial crisis.

The last paper in the volume is "Trend, Seasonal, and Sectorial Inflation in the Euro Area", by James H. Stock and Mark W. Watson. They estimate an unobserved components model with stochastic volatility for euro-area inflation. Their goal is to come up with a measure of underlying or trend inflation that is cleansed from seasonal and irregular components. The authors first use a univariate model to decompose inflation into its trend, seasonal, and irregular components. A drawback of the univariate approach is that the resulting estimates of trend inflation are highly imprecise. This motivates the analysis of a multivariate unobserved components model that exploits the heterogeneity in the time-series properties of 13 sectoral inflation measures, while allowing for stochastic volatility in the seasonal components. By estimating a multivariate model, the authors can obtain much precise estimates for the trend component of inflation. Trend inflation is shown to display a substantial correlation with measures of cyclical activity.

The Passthrough of Large-cost Shocks in an Inflationary Economy

Fernando Álvarez
University of Chicago and NBER

Pablo Andrés Neumeyer
Universidad Torcuato Di Tella

This paper surveys and modestly extends the theory of menu-cost models of the behavior of the aggregate price level after large-cost shocks. It does so in the context of an economy with a high underlying rate of inflation. It concentrates on the effect of large permanent and unexpected increases in the nominal price of inputs on the price level at different horizons. We use a simple theoretical model where increases in nominal cost will increase aggregate prices one for one in the long run. We study how the nominal rigidities implied by a menu cost distribute the increases in the price level between the impact effect immediately after the cost shock and the subsequent price adjustment until the price catches up with its long-run increase. In other words, we study the passthrough of large-cost shocks at different horizons. We pay particular attention to the role of the underlying inflation rate as well as to the role of the size of the cost shock, since both elements are important to determine the dynamics of aggregate prices.

Our interest in that question comes both from interest in testing aspects of price-setting theories and from a practical monetary policy point of view. On the theoretical side, the differential effect of large versus small shocks is the hallmark difference between menu-cost models and time-dependent models of price adjustment. Hence, the characterization of the model's behavior is important to be able to

Prepared for the XXII Conference of the Central Bank of Chile, Santiago, October 25–26. We thank participants at the XXII Conference in Santiago de Chile; Seminars at the Federal Reserve Banks of Minneapolis, St Louis, and Chicago; and at the Mantel lecture at the AAEP meeting in Argentina. We especially thank the comments by David López-Salido.

Changing Inflation Dynamics, Evolving Monetary Policy edited by Gonzalo Castex, Jordi Galí, and Diego Saravia, Santiago, Chile.

discern between theories with observed experiences. On the monetary policy side, we are interested in this particular question because of the recent experience in Argentina, an economy where inflation has been quite high for international standards in the last decade, and where due to changes in macroeconomic policies in the last four years there have been large-cost shocks. In particular, there have been large changes in exchange rates as well as extremely large changes in the price of regulated prices, mainly inputs related to energy. In this context we ask the question of whether the response to these large-cost shocks is close to one of an economy with fully flexible prices or to one with time-dependent price setting, such as the Calvo model.

We give a short narrative of instances where, in a single month, there have been very large-cost changes in the inputs for goods that make up the core CPI for Argentina between 2012 and 2018. We also use a comprehensive, never used in the price adjustment literature, micro-data set underlying the construction of the core CPI for the city of Buenos Aires to compute the objects analogous to the ones we describe in the theory. From the comparison between statistics in the model and in the data, we conclude that either full price flexibility or time-dependent rules are quite counterfactual. Nevertheless, in the cases of very large positive cost shocks in an economy with an underlying high inflation, the passthrough of the shock to prices is quite fast. In the most extreme case, when the exchange rate jumped almost 25% in one day (from 31.94 pesos per dollar on August 29th to 39.60 on August 30th), the fraction of firms changing prices, which is 24% in normal times, rose to almost 60% between August and September 2018.

We believe that through this narrative approach we can make a plausible case that regulated prices can be thought of as exogenous changes in cost, and that the same will be true for some of the large changes in exchange rates. The change in regulated prices followed a period of many years during which, despite the fact that inflation was very high for international standards, regulated prices were frozen in nominal terms. The situation became untenable from the fiscal point of view, costing at least 3% of GDP, and also due to the distortions that this policy generated. The normalization of prices occurred in steps due to the very large increases that it implied, as well as due to challenges in the courts. To get an idea of the magnitude of the change in relative prices, during the sample period, the price of natural gas relative to the core CPI was multiplied by a factor of five and the relative price of electricity by a factor of three. On the exchange-rate changes there were different reasons for the observed large changes. In 2014 there was a devaluation within a regime with severe capital controls and dual exchange rates. At the end of 2015 and beginning of 2016, there was a large devaluation, when the multiple exchange-rate regime was abandoned essentially overnight. The sharp depreciation of the exchange rate throughout April–August 2018 is probably the result of a mix of a reaction to a change in economic policies (perhaps larger than what policymakers had anticipated) and exogenous shocks within the framework of a dirty floating exchange-rate regime. In our model, cost shocks are once-and-for-all unexpected changes, which may not be completely accurate for some of the episodes.[1]

We use the microdata underlying the construction of the consumer price index (CPI) in the city of Buenos Aires to compare the predictions of the model with actual observations after several of the large-cost increases that occurred between 2014 and 2018. Some readers may be aware that the National Statistical Agency (*Instituto Nacional de Estadística y Censos, Indec*) produced price indices that are widely regarded as underestimating the inflation rate between 2008 and 2015.[2] The price index of the city of Buenos Aires has the advantage that it was independent from the National Statistical Agency and produced figures very similar to the privately produced price indices.

We propose a theoretical model where firms have both idiosyncratic as well as aggregate changes in costs. We assume that firms are monopolistic competitors and that they face a fixed menu cost for changing prices. The solution to the firms' price-setting problem gives rise to a classical sS rule for price changes, which postulates price increases as well as price decreases. The optimal decision rule, as well as the response of the aggregate price level, depends crucially on the ratio of the inflation rate to the variance of the idiosyncratic shocks. While idiosyncratic shocks make the analysis more complicated, we think that they are essential to the answer of the problem for two reasons. First, they are required to reproduce the large fraction of price decreases that are observed even when inflation rate is above

1. For instance, the normalization of the exchange-rate system and removal of capital controls is a policy that was to a large extent announced by the two main parties before the elections at the end of 2015. Indeed there is an unsettled debate in the Argentine economic circles on whether the effect of the likely increase in the exchange rate that the unification of the exchange-rate markets will entail was anticipated and included in price changes months before it happened. We discuss these episodes in more detail in section 1.

2. See, for instance, Cavallo (2013).

25% per year. Second, the behavior of the passthrough depends on the magnitude of these shocks relative to the level of inflation. In the theoretical section, we compare the effect of cost increases in three cases: a very low inflation rate, which is the case of most economies; high inflation rates, of the order of Argentina during this period (say, 25% per year); and very high inflation. The degree of passthrough depends on both the size of the steady-state inflation rate and the size of the shocks. It turns out that, even for inflation rates as large as 25% per year, still one can see clear effects of price stickiness. Yet, for large-cost increases of the magnitude that occurred in Argentina, the passthrough occurs in an extremely short time. In particular, there is a very large impact effect, with most of the adjustment occurring at the time of the shock, and a very short half-life, smaller than two months, for the remaining adjustment. Indeed this matches what we see during the months of large-cost increases in Argentina: a very sharp increase in the fraction and in the size of price increases, almost no change in the fraction and in the size of price decreases, and a jump on the inflation rate close to the size of the cost increase.

It is well known that cost shocks explain a large fraction of the variance of inflation in standard estimates of medium size New Keynesian models. Nevertheless, these costs are typically a residual in the standard specification. On the other hand, there is a literature that tries to use identified cost shocks to evaluate different price-setting mechanism. Our paper adds to the literature that studies the effect of large-cost shocks in price-setting models with menu costs of price adjustment. Early examples of this literature are Gagnon (2009) and Gagnon and others (2013). These papers consider the experience of Mexico in the mid-90s, when there was both a large step devaluation (about 40%) and changes in the VAT (from 10% to 15%). Another similar recent exercise is the one in Karadi and Reiff (2019), which uses changes in the VAT in Hungary. In both cases, a version of a menu-cost model is used to interpret the microdata underlying the construction of CPI and to compare the predictions of this class of models with the data. A related, yet different evidence, is the study of the exchange-rate passthrough to export and import prices in customs data in Bonadio and others (2016), comparing small changes in the Swiss franc's value with the large change that occurred when the Swiss National Bank abandoned its peg to the euro in January 2015. Finally, Álvarez, Lippi, and Passadore (2016) estimate panel regressions of the short-term passthrough of exchange-rate changes to consumer prices that include non-linear terms in the size of the exchange-rate changes.

Differently from the previous studies, these panel regressions do not use microprice statistics. Relative to Gagnon (2009), Gagnon and others (2013), and Karadi and Reiff (2019), and motivated by the levels of inflation in Argentina during the period of interest, this paper has a more systematic treatment of the role of the running (or steady-state) inflation and of the size of the cost shocks on the level of short-term passthrough, and on the overall speed of adjustment. Finally, relative to Álvarez and others (2019) we study a different period of time for Argentina, but more importantly, in this paper we concentrate on the effect that large unexpected-cost shocks have on price dynamics, as opposed to the effect of different steady-state inflation levels, which is the main point of Álvarez and others (2019).

The remainder of the paper is organized as follows. In section 1 we provide a brief narrative of macroeconomic events in Argentina in the period 2012–2018 as they pertain to nominal cost shocks and inflation. The theoretical analysis is in section 2. The firm's problem is described in subsection 2.1; subsection 2.2 describes the steady-state distribution of price markups over nominal marginal costs; subsection 2.3 explains how inflation affects optimal decision rules; subsection 2.4 derives analytically expressions for the impulse responses of consumer prices to cost shocks; and subsection 2.5 contains numerical simulations of how cost shocks affect consumer price dynamics for small and for large-cost shocks, and for three different inflation rates (low, high, and extremely high). Finally, section 3 compares the predictions of the model with the evidence emerging from city of Buenos Aires consumer price microdata.

1. Nominal Cost Shocks and Inflation, Argentina 2012-2018

In this section we provide some background on the evolution of inflation and the nominal value of some key inputs during our sample period, July 2012 to December 2018. We first comment on the monetary policy framework and exchange-rate developments, and later on regulated price policies for energy inputs.

We divide the analysis of the monetary policy framework in three periods: the dual exchange-rate regime prior to December 2015, the inflation-targeting regime between March 2016 and December 2017, and the abandonment of the inflation-targeting regime throughout 2018.

During the first period, the monetary policy framework was a dual exchange-rate regime. The Central Bank fixed the Argentine peso price of the U.S. dollar for transactions related to international trade as well as for some limited financial ones. Capital controls were binding and a shadow exchange-rate market that carried a premium over the official one developed. We believe that this was caused by the high rate of money growth due to the monetary financing of deficits. Part of the monetary financing of deficits was sterilized with Central Bank debt that reached close to 6% of GDP in March 2016. The average rate of core inflation between July 2012 and December 2015 in the city of Buenos Aires was 31% (with a median of 28%).[3] The official exchange rate crawled at a median rate of 17% (annualized median monthly rate of devaluation). Between November 2013 and March 2014, a succession of jumps in the exchange rate resulted in a cumulative increase of 32% in the exchange rate. Core CPI prices in those months rose by 15%.

On December 15th, 2015, exchange-rate controls were removed and the exchange rate started to float with limited central bank intervention. The unification of the foreign-exchange market entailed a depreciation of the peso of over 50% between December 2015 and February 2016.[4] In March 2016 the Central Bank adopted an interest rate peg as its policy instrument in the context of an incipient inflation-targeting regime, which was formally adopted in September. The national inflation targets were 12–17% for 2017, 10% ± 2% for 2018, and 5% ± 1.5% for 2019. The actual core inflation rates in the city of Buenos Aires were 35% in 2016, 24% in 2017, and 43% in 2018. Wages increased by 34% in 2016, 24% in 2017, and 21% in the first nine months of 2018 (prices increased by 30% in this period).[5]

Starting in December of 2017, there was a gradual abandonment of the inflation-targeting framework. On December 28th, 2017, the inflation target for 2018 was raised from 10% to 15%, and in January the Central Bank lowered interest rates by 150 basis points from 28.75% to 27.25%. The market perception was that this interest rate move was motivated by political pressure. A speculative run on the Central Bank's debt unraveled between April and September 2018. In the period between April and September (end of period), the Central Bank paid (did not rollover) 529 billion pesos of short-term debt (40% of its debt and 53% of the monetary base in April), and the monetary

3. Average and median of annualized monthly inflation.
4. See table 1 and figure 1.
5. Core inflation for the city of Buenos Aires. Wages are from *Ministerio de Trabajo, Empleo y Seguridad Social* (2018). [Ministry of Labor, Employment and Social Security]

base increased by 250 billion pesos in the same period.[6,7] This led to a sharp depreciation of the Argentine peso. The peso-dollar exchange rate quivered around 17 pesos per dollar in the period July–December 2017, then around 20 pesos between late January and the end of April, and reached 41 pesos per dollar in late September. The evolution of the (log) exchange rate is depicted in figure 1.

Public utilities prices were also an important source of large nominal cost shocks during our sample period. The nominal price of regulated goods such as electrical power, natural gas, water, and public transportation was practically frozen for a decade up to 2014. With the price level rising, the relative price of regulated goods, especially energy items, steadily eroded. The gap between the nominal marginal cost of these regulated goods and their sale prices was covered with government subsidies that in 2015 are conservatively estimated around 3% of GDP. A small attempt at normalizing these prices resulted in large nominal increases in the price of natural gas during 2014. A more ambitious normalization started in 2016.

Figure 1. Cost Shocks and Inflation

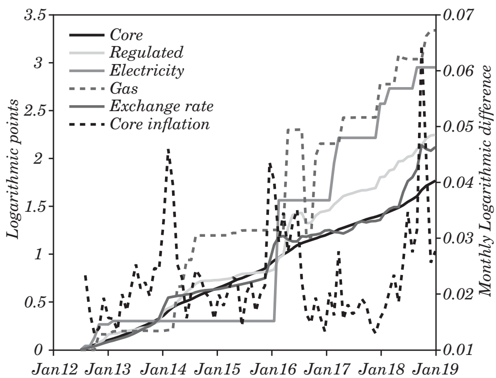

Source: Authors' calculations.
Note: Prices on the left axis are the natural logarithm of the price index published by the city of Buenos Aires normalized to one at the initial date. The exchange rate is the natural logarithm of the monthly average published by the Central Bank. Core inflation is the monthly log difference of the core price level (rest), which excludes seasonal and regulated goods and services.

6. See Bassetto and Phelan (2015) for a related theoretical model.

7. Establishing the reasons behind this run is difficult and goes beyond the purpose of this paper. Several explanations have been proposed in local economic circles, including the importance of current account deficits, the effect of negative productivity shocks on agriculture (a very large drought), the size and maturity of the Central Bank debt, changes in local taxation of capital flows, the fear of fiscal dominance in the near future, and the adverse international financial circumstances, just to cite a few.

Figure 1 provides an overview of the evolution of the nominal variables described in the preceding paragraphs. All prices are from the city of Buenos Aires CPI.[8] The solid lines depict the natural logarithm of the core CPI, regulated prices in the CPI, the price of electricity, the price of natural gas, and the exchange rate. Core inflation excludes seasonal products and regulated prices. All five are normalized to their July 2012 value.

The black solid line is the price level for core goods and services while the dotted black line represents its rate of change (in log differences). The dark grey line represents the exchange rate expressed as pesos per dollar. There are three major devaluations. During the dual exchange-rate period, the exchange rate exhibits a rate of growth below that of core prices except for the period of the devaluation around January of 2014. There is a second sharp depreciation of the peso between December 2015 and February 2016 after the removal of capital controls. The last episode of depreciation of the peso, about 80%, occurred between April and September of 2018.

Regulated consumer prices in the city of Buenos Aires are represented by the light grey line in figure 1. As it was the case for the exchange rate, prior to December 2015 these prices grow at a slower pace than core prices and there are two important relative price corrections, one in 2014 and another in 2016. The main drivers of regulated prices are the prices of energy, especially electricity, and gas. The middle grey line represents the price of electricity. It is a step function with jumps of 253% in February 2016, 92% in January–March 2017, 68% in December–February 2018 and 25% in August 2018. The price of natural gas follows a similar pattern.[9,10]

We summarize the behavior of nominal cost shocks in a proxy variable, which we refer to as cost proxy of cost shock. We assume that consumer goods are produced with labor, tradable inputs, and regulated goods. In our model we assume that there is a consolidated sector that produces and retails consumption goods. It purchases an

8. We work with data for the city of Buenos Aires as the reliability of national statistics between 2007 and 2016 has been severely questioned. For example, on February 2013, the International Monetary Fund issued a declaration of censure against Argentina in connection with the inaccuracy of CPI data from the National Statistical Agency, the Indec. Also see Cavallo (2013) for a comparison of national statistics and online prices across several Latin American countries, including Argentina.

9. See table 1.

10. The price of natural gas fell 68% in August 2016 because the Supreme Court stayed the April price increase on procedural grounds. The price increases were resumed after the Executive remedied the judicial objection to the previous price increase.

aggregate input in a flexible price competitive market and sells its output to consumers in a monopolistically competitive market. To construct our intermediate input measure, we assume following Jones (2011) that the share of intermediate inputs in the production of the final consumption goods is 50%, and that the remaining is a labor share of 50%. The weights of energy and tradable intermediate goods are 10% and 40%, respectively. Hence, our measure of cost shocks is obtained by computing first a geometric price index for this aggregate input as $p_R^{0.1} E^{0.4} W^{0.5}$, where p_R denotes regulated prices, E is the peso/U.S. dollar exchange rate, and W are nominal wages.

Figure 2 depicts the cost shock proxy variable, wholesale prices and core inflation.[11] The proxy for nominal cost shocks tracks wholesale prices closely with a correlation of 0.92. Observe that spikes in nominal cost shock inflation are associated with spikes in core inflation.

Table 1. Large Nominal Cost Shocks

	Elect.	*Natural gas*	*Exchange rate*		*Elec.*	*Natural gas*	*Exchange rate*
Oct-12	9			Feb-17	48		
Nov-12	15	18		Mar-17	30		
Jan-14			12	Apr-17		31	
Feb-14			11	Jul-17			7
Apr-14		57		Dec-17	43	42	
Jun-14		32		Jan-18			8
Aug-14		29		Feb-18	17		
Dec-15			19	Apr-18		36	
Jan-16			19	May-18			17
Feb-16	253		8	Jun-18			12
Apr-16		184		Aug-18	25		9
Jul-16			5	Sep-18			28
Aug-16		–68		Oct-18		27	
Oct-16		134		Nov-18		6	
Nov-16		14					

Source: Authors' calculations.
Note: We report the percentage change between prices in month t and t–1. Price data for electricity and natural gas is from the Consumer Price Index for the city of Buenos Aires. Exchange rate is the change in the monthly average exchange rate.

11. There is no official reliable data on wholesale prices prior to December 2015.

Figure 2. Cost Shocks and Inflation

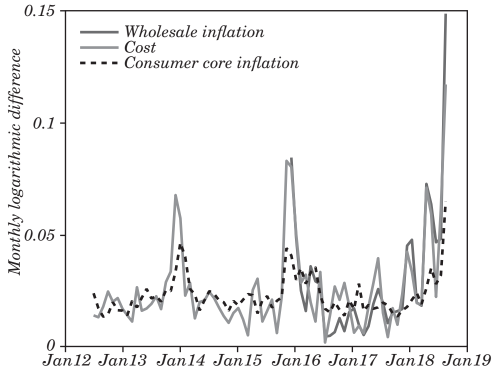

Source: Authors' calculations.
Note: Prices on the left axis are the natural logarithm of the price index published by the city of Buenos Aires normalized to one at the initial date. The exchange rate is the natural logarithm of the monthly average published by the Central Bank. Core inflation is the monthly log difference of the core price level (rest), which excludes seasonal and regulated goods and services.

In the next section we present a theoretical model of the pricing decision of a retail firm in order to study the speed and the magnitude of the passthrough of these cost shocks to consumer prices.

2. Passthrough of Cost Shocks to Prices: Theory

In this section we study the nominal passthrough to consumer prices of the nominal cost shocks in a menu-cost model of price adjustment. We consider the problem of a monopolistically competitive firm that faces idiosyncratic demand and productivity shocks in an environment in which the cost of an aggregate input grows at a constant inflation rate. We study how the aggregate consumer price level reacts to an unexpected jump in nominal marginal costs in this environment.

We first describe the firm's price-setting problem in subsection 2.1 and then look at the cross sectional distribution of prices in subsection 2.2. We then proceed to describe how steady-state inflation affects decision rules in subsection 2.3, which refers to Álvarez and others (2019). Finally, in subsection 2.4, we report the impulse response of price distributions to cost shocks and, in subsection 2.5, we compare the reaction of prices to small and large-cost shocks and find interesting nonlinear effects.

2.1 Price-Setting Problem for the Firm

We consider an economy where firms' marginal nominal costs have a common and an idiosyncratic component. The common component is given by a nominal cost, which we will assume that, after an initial value is realized, will grow at a deterministic constant rate. The (log of the) idiosyncratic component follows driftless Brownian motion, with innovation variance σ^2. Firms face a downward sloping demand, act as monopolistic competitors, and must pay a fixed menu cost to adjust their price. The firm's marginal (and average) cost will then be $W(t)x(t)$, where $W(t)$ is the nominal cost of the aggregate input and $x(t)$ is an idiosyncratic shock. We assume that $x(t) = exp(\sigma B(t))$, where $B(t)$ is a standard Brownian motion, independent across firms. We will start the firms at a steady state, where W_t has been growing at a constant inflation rate π, so that $W(t_0+T) = W(t_0)e^{T\pi}$.

We will use g to describe the logarithmic deviation of the firm's current markup relative to the one that maximizes instantaneous profits. We refer to this variable as the "price gap". The price gap is positive, g > 0, if the price of the product is relatively high or if the firm's cost is relatively low. We will assume that the optimal markup is independent of the level of the costs and of the demand, as it will be the case with an iso-elastic demand function and constant marginal cost. Note that during a period where the firm does not change prices, its markup changes only because its cost changes, i.e., $dg(t) = -\,d\log W(t) - d\log x(t)$. For instance, at steady state, we have that $dg = -\pi dt - \sigma dB$ during the period at which the price of the good does not change. At the times when the firm decides to pay the fixed cost and change prices, $g(t)$ changes discretely and in equal proportion (equal in log points) to the change in price.

The firm problem can be summarized by the Bellman equation:

$$V(g) = \max_{\tau} \mathbb{E}\left[\int_0^{\tau} e^{-\rho t} F\left(g\left(t\right)\right)dt + e^{-\rho\tau}\left[\psi + \max_{\tilde{g}}\ V\left(\tilde{g}\right)\right] \middle| g(0) = g\right] \qquad (1)$$

subject to

$$dg(t) = -\,\pi dt - \sigma dB(t) \text{ for all } 0 \le t \le \tau \text{ and } g(0) = g. \qquad (2)$$

In this problem ψ is the fixed cost, ρ is the (real) discount rate, π is the inflation rate (of the aggregate input), σ is the idiosyncratic volatility of the cost shocks, and $F(\cdot)$ is the instantaneous profit function, written as a function of the price gap. The expectation in the objective function is with respect to the cost shocks, the values of the path for $B(t)$. Mathematically speaking, the object of choice τ are *stopping times*. A *stopping time* indicates the time and circumstances under which prices will be adjusted. In this kind of problem, the optimal rule is that τ occurs the first time that $g(t)$ is outside a range of inaction. After paying the cost and deciding the optimal price we find the optimal return point as $g^{*} = arg\max_{\tilde{g}} V\left(\tilde{g}\right)$, or $V(g^{*}) = \max_{\tilde{g}} V\left(\tilde{g}\right)$. In appendix B we write the ordinary differential equation (o.d.e.) and boundary conditions which simultaneously determine the function $V(\cdot)$ and the values of $\underline{g}, g^{*}$ and $\overline{g}$. In appendix C we formulate and compute a discrete time and discrete state-space problem that approximates the continuous one.

The optimal policy for the firm is an sS rule described by three numbers $\underline{g}, g^{*}$ and $\overline{g}$. We refer to $\underline{g} < \overline{g}$ as the boundaries of the range of inaction, and to $g^{*} \in (\underline{g}, \overline{g})$ as the optimal return point. The inaction region is thus described by the interval $[\underline{g}, \overline{g}]$. If the price gap g is in the inaction region, the price of the firm stays constant and the price gap changes with the cost changes—with opposite sign. If the price gap is lower than $\underline{g}$, so that the markup is very low, then the firm will pay the fixed cost ψ and increase prices so that, right after the price change, the price gap will be g^{*}. Likewise, if the price gap is higher or equal than $\overline{g}$, so that the markup is very high, the firm will pay the fixed cost ψ and decrease its price, so that, right after the price change, the price gap becomes g^{*}.

Note that we are writing the profit function F as a function of the price gaps exclusively. In general, even at steady state, it should be a function of both the price and the cost. Below, we explain the conditions under which this simplification can be obtained. We choose units so that F is measured in either terms of units of the aggregate input, or as deviations of maximized steady-state profits—our preferred choice. Of course, the fixed cost ψ has to be measured in the same units as F. The optimal decision rule depends only on the ratio ψ/B, since the value function is homogeneous of degree one in ψ and B. Intuitively, the fixed cost matters only through the relative advantages of changing prices, captured by the curvature of F, rather than on each of them separately. Furthermore, the discount factor ρ is a real interest rate measuring the intertemporal price of the aggregate input. The inflation rate π is also the nominal change in the price of the aggregate input. Since ρ, π and σ^2 are rates per time unit, the optimal decision rule depends, apart from on ψ/B, only on the ratios $\{\rho/\sigma^2, \pi/\sigma^2\}$.

Derivation of the profit function F. We have written the profit function F having the price gap g as its only argument. This is a simplification which provides lots of tractability. In Álvarez and others (2019), we work without this simplifying assumption and obtain essentially the same results.[12] We describe here the assumptions so that we can use the simplified version. Let $Q(p/W)z$ be the quantity demanded as a function of p/W, the ratio of the nominal price of the good p and the nominal price of the generic input W. In steady state we can write this relative price or the relative price with respect to some aggregate good without loss of generality. It also turns out that in the set-up described by Golosov and Lucas (2007), which is in no way pathological, this is a consequence of the general equilibrium structure. The variable z is a multiplicative shifter of the demand, which we use for two different illustrations. The nominal marginal cost is xW. We will proceed by steps. First we will derive the profit $\tilde{F}$ with (g, z, x) as arguments. Then we will add assumptions to eliminate (x, z) from it.

Let's use m for the log of the optimal markup, i.e., let the nominal price that maximize instantaneous profits be: $P^{*} = e^{m}xW$. Thus g is defined as:

$$g \equiv \log\left(\frac{\frac{P}{Wx}}{\frac{P^{*}}{Wx}}\right) = \log\left(\frac{\frac{P}{Wx}}{M}\right) = \log\left(\frac{P}{Wx}\right) - m \text{ or } P = Wxe^{g+m}. \tag{3}$$

Now we are ready to write the profit function. The units of profits will be first in terms of the real value of the aggregate input. Profits in nominal terms are $[P - xW]Q(P/W, z)$, i.e., nominal markup times quantity. By dividing this expression by W, we obtain profits in time t units of the aggregate input.

$$\tilde{F}(g,x,z) = \left[\frac{P}{xW} - 1\right] Q\left(\frac{P}{Wx}x\right) zx = \left[e^{g+m} - 1\right] Q\left(e^{g+m}x\right) zx \tag{4}$$

Furthermore, *assuming that Q is iso-elastic*, with elasticity equal to η, so that $Q(P/W) = A(P/W)^{\eta}z$ for some constant A, then the optimal markup M, or its log m, will be constant and equal to $m = \log(\eta/(\eta - 1))$. In this case profits will be:

$$\tilde{F}(g,x,z) = \left[e^{g+m} - 1\right] e^{-\eta(g+m)} A x^{1-\eta} z.$$

12. See the section on the model with random walk shocks and CES demands, which we refer to Kehoe and Midrigan (2015) version of the Golosov and Lucas (2007)'s model.

It should be clear that, by definition, F is maximized when $g = 0$.

Adding the *extra assumption* that the demand shifter satisfies $z = x^{\eta-1}$, then we obtain that $\tilde{F}$ does not depend on (x, z). To be honest, this is a strange assumption; it requires the shock that increases cost to simultaneously push the demand up, so that the maximized profit remains the same for any value of x. On the other hand, it simplifies the algebra a lot! As mentioned above, in Álvarez and others (2019) we work out both versions of the models and find very small differences.

An alternative is to use a *second order approximation* of the function $\tilde{F}$ (g, x, z) and to retain only the leading terms on g. In particular, we use a second order approximation of $\tilde{F}$ (g, x, z) around $g=0$, $x=\bar{x}$ and $z=\bar{z}$. Using that $g = 0$ maximizes profits, and the multiplicative separable nature of the profit function into three terms, ignoring the terms that are higher than second order and the terms not involving g, we have

$$F(g) \equiv -\frac{(\eta-1)\eta}{2} g^2 = -Bg^2 \quad \text{so that} \quad B \equiv \eta(\eta-1)/2 \tag{5}$$

where we are measuring profits relative to the maximized profit at $x=\bar{x}$ and $z=\bar{z}$, which equals $\tilde{F}(0,\bar{x}) = \left(\frac{\eta}{\eta-1}\right)^{-\eta}\left[\frac{A\bar{x}^{1-\eta}\bar{z}}{\eta-1}\right]$. Of course, we can combine these assumptions too.

Lack of First-Order Strategic Complementarity. Finally, the result in equation (5) is useful because it states that in the model we focus on, there is no *first-order strategic complementarity*. This is because we can summarize the behavior of the rest of the firms in z. This will be the case in the standard New Keynesian model, where there will be at least two effects on the profit function coming from aggregate consumption: one is that with CES demand for each firm, higher output of each of the other firms shifts the demand up, and the other, in opposite sign, that higher output decreases the Arrow-Debreu price of the good in the current period. But as we have seen, these effects—captured by z—are of third order in the profit function in the current set-up.[13] Interestingly, this means that we can use as an accurate approximation the firm's decision rules characterized by $\{\underline{g}, g^*, \bar{g}\}$ even if the rest of the economy is not at a steady state, as long as the firm expects constant growth rate of the aggregate input

13. The general equilibrium version in Golosov and Lucas (2007) also makes the value of the nominal aggregate input, labor in their case, depend only on the path of nominal money.

prices. We will rely heavily on this property to characterize the impulse response of prices to a one-time common shock to their cost.

2.2 Steady-State Distribution of Price Gaps and Price Changes

We let $f(g)$ be the steady-state density of the distribution of price gaps. This density has support $[\underline{g}, \overline{g}]$. It solves a simple ordinary differential equation, balancing the flows in and out when a price change occurs. Its shape depends on the parameters $\{\underline{g}, g^*, \overline{g}, \pi/\sigma^2\}$. Of course, $\underline{g}, g^*, \overline{g}$ depend on all the parameters that define the steady-state problem of the firm described above, namely $\{\psi/B, \rho/\sigma^2, \pi/\sigma^2\}$ or, using the simplification described in appendix B, it is described by just $\{\psi/B, \pi/\sigma^2\}$. The equations that determine the steady-state density given the decision rules and parameters are:

$$0 = \pi f'(g) + \frac{\sigma^2}{2} f''(g) \text{ for all } g \in \left[\underline{g}, g^*) \cup (g^*, \overline{g}\right]$$

with boundary conditions:

$$0 = f(\underline{g}) = f(\overline{g}) \text{ and } \int_{\underline{g}}^{\overline{g}} f(g)\,dg = 1\,. \tag{6}$$

The first line is the Kolmogorov forward equation, and the second line has the three relevant boundary conditions. The solution is given by the sum of two exponential functions. The boundary conditions on the extreme of the inaction regions indicate that there is zero density in those points. Intuitively, this is because they are exit points, so it is "hard" to accumulate density near them. This result will be important for some of the results below.

The shape of the density f depends on the inflation rate relative to the idiosyncratic variance π/σ^2, both through their direct effect as seen in equation (6), and indirectly through their effect on $\underline{g}$, $g^*$and $\overline{g}$. For finite π/σ^2, the distribution has zero density in its two boundaries, it is strictly increasing in $(\underline{g}, g^*)$, non-differentiable at g^*, and strictly decreasing in$(g^*, \overline{g})$. In particular for $\pi/\sigma^2 = 0$, the distribution is *symmetric* around $g^* = 0$, and has a tent-shape, with $f''(g) = 0$. As π/σ^2 increases, the value of g^* becomes positive, and the shape of f becomes concave from $[\underline{g}, g^*)$ and convex between $(g^*, \overline{g}]$. As $\pi/\sigma^2 \to \infty$, the distribution converges to a uniform distribution between $[\underline{g}, g^*]$ and to a zero density everywhere else, as in the model of Sheshinski and Weiss (1979), which has $\sigma^2 = 0$ and finite π.

The change in the shape of the invariant distribution f reflects, in a very intuitive way, the different strength of the idiosyncratic shocks (measured by σ^2), which are symmetric, and the effect of inflation (measured by π), which is asymmetric. As inflation increases, the price gaps g naturally tend to pile up in the left side, since the cost increase in expected value as the nominal price remains fixed. Indeed as $\pi/\sigma^2 \to \infty$, this effect is so strong that price gaps essentially march deterministically from g^* to $\underline{g}$, and hence the distribution is uniform, as stated above. Lastly, the drastic change in shape of f around $g^*$is also intuitive, since after a price is changed, the new price of the product is set so that $g = g^*$, which explains why the mass is highest at this point and why it is not differentiable, since the behavior of f around this point is governed by mass coming from the boundaries of the range of inaction.

In steady state the size of price changes is given by a very simple formula of the thresholds of the sS rule. Price increases are given by $g^* - \underline{g}$, since the firm increases its price when the markup is very low, i.e., the first time g reaches $\underline{g}$. Likewise, price decreases occur when the markup has reached the value $\overline{g}$ and are of size $\overline{g} - g^*$. Denoting the size price increases by Δ_p^+ and the size price decreases by Δ_p^-, we have that: $\Delta_p^+ = g^* - \underline{g}$ and $\Delta_p^- = \overline{g} - g^*$.

We denote the average number of price changes per unit of time by λ_a. This can be easily computed as the reciprocal of the time between price changes, by the fundamental theorem of renewal theory. The expected time until a price change $T(g)$ for a firm with current price gap g solves the following Kolmogorov backward equation:

$0 = -\pi T'(g) + \frac{\sigma^2}{2} T''(g)$ for all $g \in [\underline{g}, \overline{g}]$with boundary conditions: $0 = T(\underline{g}) = T(\overline{g})$.

The boundary conditions are quite natural, at either $\underline{g}$ or $\overline{g}$ there will be a price change, and hence the expected time to reach them is zero! Since right after a price change $g = g^*$, then the expected time until the next price change is $T(g^*)$, and hence the frequency of price changes $\lambda_a = 1/T(g^*)$. Using a similar procedure we can find the frequency of price increases λ_a^+ and the frequency of price decreases λ_a^-. For instance, letting $T^+(g)$ be the expected time until a price increase, it is easy to see that it satisfies the same Kolmogorov backward equation than T. The difference is in the boundary conditions, which are $T^{+}(\overline{g}) = T^{+}(g^{*})$, since a price decrease will occur at $\overline{g}$ but we need to keep counting; we also have $T^{+}(\underline{g}) = 0$, since a price increase occurs at $\underline{g}$. The frequency of price increases is thus: $\lambda_a^{+} = 1/T^{+}(g^{*})$. A similar argument holds for the expected time until a price decrease T^{-}, i.e., it solves the same o.d.e. than T with boundary conditions $T^{-}(\underline{g}) = T^{-}(g^{*})$ and $T^{-}(\overline{g})=0$. The frequency of price decreases is thus: $\lambda_a^{-} = 1/T^{-}(g^{*})$.

We note that the length of the range of inaction equals the sum of the average size of price increases plus the average size of price decreases: $\overline{g} - \underline{g} = \Delta_p^{+} + \Delta_p^{-}$. This observation will become useful because these average sizes can be measured, and the length of the range of inaction is important to understand when a cost shock is large. Finally, the second moment of price changes is given by $E\left[\Delta_p^2\right] = \frac{\lambda_a^{+}}{\lambda_a}(\Delta_p^{+})^2 + \frac{\lambda_a^{-}}{\lambda_a}(\Delta_p^{-})^2$.

In appendix C we give a simple alternative numerical procedure by using a discrete time and discrete state version of the model to compute the steady-state distribution and the frequency of price changes per unit of time.

2.3 Optimal Decision Rules and Inflation

This subsection analyzes how optimal pricing decisions vary with the rate of normalized inflation, π/σ, keeping the normalized fixed cost, ψ/B, constant.[14]

For $\pi/\sigma \approx 0$, then the decision rules are approximately symmetric, with $g^{*}= 0$ and $\overline{g} = -\underline{g}$.[15] This implies that around zero inflation, the price increases and price decreases have approximately the same size $\Delta_p^{+} = \Delta_p^{-}$. Moreover, the frequency of price increases and price decreases is the same around zero inflation, so $\lambda_a^{+} = \lambda_p^{-}$.

For higher inflation rates relative to idiosyncratic volatility π/σ^2, the optimal return markup g^{*} becomes larger. This is because due to inflation, during inaction markups decrease in expected value, and thus this expected effect, is compensated by starting with a higher markup. In Álvarez and others (2019) we show that for small inflation the main adjustment is *not* in the frequency of price changes λ_a, but

14. Appendix B shows that, for a given fixed cost relative to the curvature of profits ψ/B, the optimal decision rules depend on the normalized level of inflation π/σ^2.

15. We say *approximately* because under the quadratic approximation at exactly $\pi/\sigma^2 = 0$ the decision rules are symmetric.

instead in the difference between the frequency of price increases and decreases, i.e., the derivative of $\lambda_a(\pi)$ with respect to inflation is zero at $\pi/\sigma^2 = 0$, but the derivative of $\lambda_a^+(\pi) - \lambda_a^-(\pi)$ is strictly positive. We also show analytically that it accounts for 90% of the change in inflation at low inflation; the other 10% is explained by changes in $\Delta_p^+ - \Delta_p^-$. Summarizing, as we move from zero steady-state inflation (where frequency and size of price increases and decreases are symmetric) to positive steady-state inflation, the model predicts that the frequency of price increases $\lambda_a^+(\pi)$ will be higher than the one of decreases $\lambda_a^-(\pi)$, and that the size of price increases and decreases will be approximately the same, i.e., $\Delta_p^+ \approx \Delta_p^-$. Nevertheless, while we show that the average size is similar, we also show that as we move from zero steady-state inflation, $\Delta_p^+ > \Delta_p^-$.[16]

For very large inflation, $\lambda_a^+(\pi) \to \lambda_a(\pi)$ and $\lambda_a^-(\pi) \to 0$, as $\pi \to \infty$, so most price changes are increases, and the model converges to Sheshinski and Weiss (1979), in the sense made precise in Álvarez and others (2019).

2.4 Impulse Response of the Price Level to Unexpected Cost Shocks

In this subsection we characterize the impulse response of the aggregate (log of the) price level to a once-and-for-all increase in the nominal price of the aggregate input of size δ, measured in logs. We will consider different inflation levels π and different sizes of the shock δ.

We start with an economy that is in the steady-state distribution of price gaps. This economy is characterized by parameters $\{\psi/B, \rho, \sigma^2\}$ and π. Firms will face a once-and-for-all jump on the nominal price of the aggregate input and believe that, after this jump, the price of the inputs will rise at the same inflation rate π as before the jump. We are interested in computing the effect on the (log of the) aggregate price level of the goods produced by these firms. We will distinguish between the impact effect on the price level, i.e., the effect on the moment that the unexpected jump occurs, and the subsequent effects which lead the price level to adjust up to the full amount δ of the shock.

We can think of the price level just before the shock as the limit: $\bar{P} \equiv \lim_{t \uparrow 0} P(t;\pi)$. Likewise, we can let the value of the (log of the) price

16. For a precise statement, see Propositions 1 and 3 in Álvarez and others (2019). Numerically, we find the changes given by these derivatives at zero inflation to be accurate even up to inflation rates of 30% per year.

of the aggregate input just before the shock be $\overline{W} \equiv \lim_{t \uparrow 0} W(t)$. Of course, for the price of the aggregate input we have that the jump is $\delta = \lim_{t \downarrow 0} W(t) - \lim_{t \uparrow 0} W(t) \equiv \lim_{t \downarrow 0} W(t) - \overline{W}$. As mentioned above, we assume that $d\log W(t)/dt = \pi$ for $t \neq 0$. Throughout the exercise, the parameters $\{B, \psi, \rho, \sigma^2\}$ and the invariant distribution implied by the optimal decision rules, $\{\underline{g}, g^*, \overline{g}\}$ are fixed.

We will denote the price level t periods after the shock $P(t; \delta, \pi)$ for an economy that is hit by the shock of size δ when the inflation rate of the aggregate input before and after the cost shock is π. We will let the price level right before the shock to be denoted by $\overline{P}$. We distinguish between the impact effect, which we denote by $\Theta(\delta, \pi)$, and the subsequent rate of change of the (log of the) price level, denoted by $\theta(t; \delta, \pi)$. Thus

$$P(t;\delta,\pi) = \overline{P} + \Theta(\delta,\pi) + \int_0^t \theta(s;\delta,\pi)ds \tag{7}$$

Whenever it is clear, we omit δ and π from the expression for P, Θ and θ. Also, while P is the log of the aggregate price level, whenever it is clear, we will refer to it as just the price level.

Since we are measuring P in logs, then $\theta(s; \delta, \pi)$ is the inflation rate of the CPI s periods after the shock δ has occurred in an economy with steady-state inflation rate π. In particular, after the impact effect at time $t = 0$, the term $\theta(s; \delta, \pi)ds$ yields the contribution to the average (log) price of the firms that are adjusting prices at times between s and $s + ds$. This contribution is equal to $\theta(s; \delta, \pi) = [\Delta_p^+ \lambda_a^+(s) - \Delta_p^- \lambda_a^-(s)]$ where Δ_p^+ and Δ_p^- are the same as in steady state, but the frequency of price increases and decreases, $\lambda_a^+(s)$ and $\lambda_a^-(s)$, are time-varying. The reason these frequencies are time varying is that the density of the distribution of firms indexed by their price gap g, denoted also by $f(g,t)$, is time varying. This density changes through time because the cost shock δ and the price changes that occur right after the cost shock have displaced it from its steady state. Thus, even following the same time invariant decision rules as in steady state, it takes time for the density to return to its steady-state level described by equation (6). See the appendix C for the discrete time and discrete state space analog computation of the path of the distribution and of λ's, or see Álvarez and Lippi (2019) for the continuous time characterization of the impulse response by using an eigenvalue-eigenfunction decomposition.

We will compare the path of the log of the aggregate price level $P(t;\, \delta,\, \pi)$ against the path for the price level of the aggregate input. Recall that the aggregate input grows at rate π before and after the shock, and jumps by δ log points at time $t = 0$. In the long-run prices will increase by as much as the shock to the aggregate input, so that

$$\delta = \lim_{t\to\infty}\left[P\left(t;\delta,\pi\right) - \pi t\right] - \bar{P} = \lim_{t\to\infty}\left[W(t) - \pi t\right] - \bar{W} \tag{8}$$

Impact effect. We define the impact effect Θ as the jump in the price level at $t = 0$, i.e.,

$$\Theta\left(\delta,\pi\right) = \lim_{t\downarrow 0} P\left(t;\delta,\pi\right) - \bar{P} \tag{9}$$

To be clear, if prices are fully flexible, we will have $P(t) = W(t) + \mu$ for some constant μ at all times. With menu costs, after the jump in the aggregate input prices, we expect that $\Theta \leq \delta$ and that over time the price level $P(t)$ catches up with the increases in the path of $W(t)$. Later we show that this is true for values of δ and π/σ^2 that are not too large.

The impact effect is simple to compute following this two-step procedure. First we shift the distribution of price gaps from the steady state to the one right after the shock but before the prices adjust, so the new density is $f(g + \delta)$ with support $[\underline{g} - \delta, \bar{g} - \delta]$. This is so because with the common increase in cost, the price gap of each firm decreases by δ. Second, using the lack of first-order strategic complementarity, all the firms that end up with price gaps g below $\underline{g}$ will increase their prices from their new value for g to $\bar{g}^*$. Thus:

$$\Theta\left(\delta,\pi\right) = \int_{\underline{g}-\delta}^{\underline{g}} \left(g^* - g\right) f\left(g + \delta\right) dg \tag{10}$$

This yields the following derivatives:

$$\frac{\partial}{\partial\delta}\Theta\left(\delta,\pi\right) = \left(g^* - \underline{g} + \delta\right) f(\underline{g}) + \int_{\underline{g}-\delta}^{\underline{g}} \left(g^* - g\right) f'\left(g + \delta\right) dg$$

$$\frac{\partial^2}{\partial\delta^2}\Theta\left(\delta,\pi\right) = f\left(\underline{g}\right) + \left(g^* - \underline{g} + \delta\right) f'(\underline{g}) + \int_{\underline{g}-\delta}^{\underline{g}} \left(g^* - g\right) f''\left(g + \delta\right) dg\,.$$

Evaluating these expressions at $\delta = 0$, and using that, by definition $\Theta\,(0, \pi) = 0$ and that at the exit points of the invariant density we have $f(\underline{g})$, we obtain the following expansion of Θ on δ:

$$\Theta\left(\delta,\pi\right) = \frac{1}{2}\Delta_p^+ f'\left(\underline{g}\right)\delta^2 + o(\delta^2) \tag{11}$$

As argued elsewhere,[17] if the shock δ is small, then the impact effect Θ is very small. Mathematically speaking, Θ is of second order in δ. We can also see that the leading coefficient of δ increases with inflation since, as explained above, as π/σ^2 increases, the density becomes more concave in the lower segment, until in the limit $f'(\underline{g}) \to \infty$ as $\pi/\sigma^2 \to \infty$. Thus the impact effect is of smaller order than δ, but the coefficient of δ^2 increases with inflation. Whether this is an important effect for the level of inflation rates for the period of Argentina under consideration is an important issue that we will discuss below.

Two extreme examples help to organize ideas: First, consider the case where $\pi/\sigma^2 = 0$, then $f'(g) = 1/(\Delta_p^+)^2$ is constant, and thus for $\delta \le \Delta_p^+$, equation (10) becomes

$$\Theta(\delta,\pi) = \frac{1}{(\Delta_p^+)^2}\int_{\underline{g}-\delta}^{\underline{g}} \left(g^* - g\right)\left(g - \underline{g} + \delta\right)dg = \frac{\delta^2}{2}\frac{1}{\Delta_p^+}\left(1 + \frac{1}{3}\frac{\delta}{\Delta_p^+}\right).$$

As in the general case, for small δ, the value of Θ is very small. Yet when δ is large, say in the order of magnitude of Δ_p^+, the impact effect can be large. For instance, if $\delta = \Delta_p^+$ we have $\Theta(\delta, \pi) = 2/3\delta$, which is smaller than δ, but of the same order of magnitude than δ. Second, consider the other extreme case, where $\pi/\sigma^2 \to \infty$. In this case f converges to a uniform distribution between $[\underline{g}, g^*]$ and thus equation (10) becomes

$$\Theta(\delta,\pi) = \frac{1}{\Delta_p^+}\int_{\underline{g}-\delta}^{\underline{g}} \left(g^* - g\right)dg = \delta\left(1 + \frac{\delta}{2}\frac{1}{\Delta_p^+}\right),$$

which is of order δ. Note that for small δ, this gives the same answer as the case of full price flexibility, i.e., $\Theta \approx \delta$, and indeed it converges to a version of Caplin and Spulber (1987)'s neutrality case. Interestingly, when δ is not infinitesimal, then $\Theta > \delta$. In this case, since $P(t) - \overline{P} - \pi t \to \delta$ as $t \to \infty$, there must be an overshooting in the short run, and thus prices should have an eco and oscillate as they converge to their path. Indeed if $\delta = \Delta_p^+$ we have $\Theta(\delta, \pi) = 3/2\delta > \delta$.

17. See Álvarez and Lippi (2014), Álvarez, Le Bihan, and Lippi (2016), and Álvarez, Lippi, and Passadore (2016).

The stark difference between the cases with $\pi/\sigma^2 \approx 0$ and $\pi/\sigma^2 \to \infty$ calls for an evaluation in the case of Argentina during the period of interest where inflation rate is quite high, but far away from hyper-inflationary levels, say on the order of 25% per year when we exclude the peaks. Is this closer to the $\pi/\sigma^2 \to \infty$ limit, or is it closer to the $\pi/\sigma^2 \to 0$ limit? Additionally, is the size of the cost changes for δ for this period in Argentina large enough that we have to go beyond the approximation in equation (11)? Note that a relevant theoretical comparison is how large δ is relative to Δ_p^+. Motivated by these considerations, we will evaluate the relevant expressions for calibrated parameters values and compare them with the "observed" impact effects.

Initial slope of the impulse response. We now characterize the initial inflation rate, just after the impact effect. For this we consider a very short-time interval right after the shock, which we denote by Δ. On the one hand, we note that immediately after the shock, there is no density near the upper bound $\bar{g}$. On the other hand, there is a strictly positive density at the lower-bound $\underline{g}$. Recall that price increases will occur for the firms in the lower bound of the inaction region, which gets an idiosyncratic increase in cost, and hence a decrease in g. Using the assumption of the Brownian motion for the idiosyncratic shocks, it can be shown that about half of the firms at the lower bound will increase prices in the very small interval of time following the aggregate cost shock. This means that, for a very short interval immediately after the impact effect, there is an extremely large number of price increases and almost no price decreases.

In particular, after the shift due to the common cost shock, the density is zero in the upper interval $g \in [\bar{g} - \delta, \bar{g}]$. Thus, $\lambda_a^-(\Delta) \to 0$ as $\Delta \to 0$. On the other hand, after the impact effect, and differently from what happens at steady state, there is a positive density at $g = \underline{g}$. This density is equal to $0 < f(\underline{g} + \delta) < f'(\underline{g})\delta$, where the inequality holds due to the concavity of the steady-state density f in $[\underline{g}, g^*]$. Consider the discrete-time discrete-state approximation developed in appendix C, where each step of the process for g and of the discretized steady-state distribution are of size $\sqrt{\Delta}\sigma$. The number of firms changing prices per unit of time $\lambda_a^+(\Delta)$ equals the density at the boundary times the step size $\sqrt{\Delta}\sigma$ times the probability that those firms have an increase in cost, denoted by p_d, and divided by the length of the time period Δ. For a diffusion, as $\Delta \to 0$ then $p_d \to 1/2$. Thus the fraction of firms changing prices per unit of time is $f(\underline{g} + \delta)\sigma/(2\sqrt{\Delta})$. This implies that $\lambda_a^+(\Delta) \to \infty$ as $\Delta \to 0$. We have then $\theta(\Delta) \to \lambda_a^+$

$(\Delta)\Delta_p^+ \to \infty$ as $\Delta \to 0$. Yet, $\theta(\Delta)\Delta \to 0$, as $\Delta \to 0$, so the integral for $P(t)$ is still well defined. An alternative more general way to show that the slope of the impulse response is infinite is in Álvarez and Lippi (2019), who use an eigenvalue-eigenfunction decomposition of the relevant linear operator.

2.5 Comparative Static of Cost Shocks

In this subsection we compare the effect of a small and a large-cost shock, say $\delta = 0.01$ and $\delta = 0.1$ for three economies: a low inflation one $\pi = 0.025$, a large inflation one, $\pi = 0.25$ and one close to the hyperinflationary range, $\pi = 2.5$. Recall that cost shocks are measured in logs, so we are trying 1% and 10% once-and-for-all shocks. Inflation rates are measured as annually continuously compounded (c.c.), so we are trying 2.5%, 25% and 250%, but in the last two cases recall that continuously compounded and annually compounded can be meaningfully different.[18]

Calibration. We use the same parameters for all cases. The parameters for the firm problem are chosen so that at $\pi = 0.25$, i.e., 25% annual continuously compounded inflation, the steady-state statistics resemble the same statistics in Argentina for the period under study. We use $\eta = 7$, which has a markup of just above 15% and a fixed cost of $\psi = 0.012$ yearly frictionless profits. We use an annual discount rate $\rho = 0.04$ and an annual volatility of idiosyncratic shock of $\sigma = 0.20$.

With these parameters the model implies $\Delta_p^+ = 0.12$ and $\Delta_p^- = 0.097$. The average number of price changes per year are $\lambda_a^+ = 2.88$ and $\lambda_a^- = 0.89$.[19] These figures are similar to the averages for Argentina, when we omit periods of abnormal cost increases. The size of cost changes in the data is around 10%–11%.[20] The annual number of price changes measured as the fraction of outlets changing prices in a "normal" month times 12 is $\lambda_a^+ \approx 0.24 \times 12 = 2.88$ and $\lambda_a^- \approx 0.073 \times 12 = 0.88$.[21]

We will discuss the effect of small ($\delta = 0.01$ or 1%) and large ($\delta = 0.10$ or 10%) cost shocks for each of the three continuously

18. Continuously compounded yearly inflation at rate π implies that the ratio of prices (or cost) at the end of the year relative to prices at the beginning of the year is e^{π}. For example, with $\pi = 0.25$, this ratio is $e^{0.25} \approx 0.28$, and for $\pi = 2.5$, this ratio is $e^{2.5} \approx 12.2$!

19. We use a time period $\Delta = 1/365/12$, so two hours.

20. See figure 8.

21. See figure 7.

compounded annualized inflation rates we consider: $\pi = 0.025$, $\pi = 0.25$ and $\pi = 2.5$. The plots corresponding to each of the inflation rates are in figures 3, 4, and 5, respectively. The small shock is of a size we consider in the upper bound of what a normal monetary shock is. The large-cost shocks are similar of the type of shocks we argue occurred in Argentina during the period under consideration, which we view as extremely large. The three inflation rates correspond to: (i) a "common" inflation rate for a developed economy ($\pi = 0.025$ or 2.5% c.c. per year), (ii) a high inflation rate which is about the average running inflation during the period of study for Argentina when we exclude the spikes we associate to the jumps in cost ($\pi = 0.25$ or 25% c.c. per year), and (iii) a very large inflation on the hyperinflation range of almost 23% monthly compounded inflation ($\pi = 2.5$ or 250% c.c. per year).[22] For each of the inflation rates, we present three panels of plots: the first panel with the path of the log of the aggregate CPI level and with the path of the log of the price of the aggregate input, the second panel with the density of the invariant distribution right before and right after the cost shock, and the third panel with the path for the (monthly moving average of the) frequency of price increases and price decreases. Each panel displays two cases, corresponding to the small-cost change (1%) on the left side of the panel, and to the large-cost change (10%) on the right side of the panel. In total there are nine subplots for each of the three inflation rates.

Some general comments on the objects of figures 3, 4, and 5 are in the same order.

First, in the top panel for each figure, we have the path of the log of the nominal price for the CPI and of the path of the log of the nominal price of the aggregate input. We normalize the price of the nominal input so that, at the time of the cost shock, both the core CPI and the price of the aggregate input are equal. The time of the cost shock is labeled $t = 1$, and time is measured in years for the first and third panels. Second, in the middle panel for each figure we have the invariant distribution for the price gaps at steady state and its version right after the cost shock, but before the price changes. In the horizontal axis we have indicated the values of $\underline{g}, g^*, \overline{g}$. The price gap is measured in log point deviations, so that $g = 0.1$ represents 10% log point difference between the static maximizing markup and

22. To be clear, for $\pi = 2.5$ annual continuously compounded rate, we have that the monthly compounded rate is $(e^{\frac{2.5}{12}} - 1) \times 100 \approx 23$.

the markup corresponding to that value of g. We have shaded in grey the distribution of price gaps right after the cost shocks. This shaded area measures the fraction of firms (or products) that change prices on impact, i.e., at the time of the cost shock. Third, the bottom panel displays the average number of price changes per year. This is defined as follows: we take the fraction of firms (or products) that change prices per model period and then we divide it by the length of the model period. In the plot we display a centered monthly moving average of this number.[23] We take a monthly moving average for comparability with the Argentine data, which uses monthly frequency, and also to smooth out the large jump. Note that at the time of the increase of the cost, these frequencies increase smoothly for about a month due to the fact that we use a centered moving average.

We end this section with a brief discussion on how the pattern of the different statistics in these figures illustrates the analytical properties derived above. First we discuss the difference between the impact effect of the jump in the cost of the aggregate input on the price level and the rate of converge of the price to the cost. Let's first concentrate on the case of low ($\pi = 0.025$) and high inflation ($\pi = 0.25$), i.e., figures 3 and 4, respectively. Even though the inflation rate that roughly corresponds to Argentina is the large case (π = 25% c.c. per year), the pattern for the passthrough is similar in figures 3 and 4. For both inflation rates, the instantaneous passthrough is larger for the large-cost shock (as can be seen by comparing the left and right subplots). Nevertheless, as expected, in the case of high inflation (π = 25%), the passthrough is higher and the convergence is faster, i.e., the half-life of the shock with high inflation is one half relative to the low inflation one. The convergence rates in the figures with large shocks are so high, with half-lives below two months, that we are very close to full price flexibility.

23. The centered moving average at time t takes an average of half of $(1/\Delta)/12$ model periods before the date t, and half of $(1/\Delta)/12$ model periods after the date t, where Δ is the length—measured in years—of the model period. See section C for details on the computations.

Figure 3. Passthrough of Nominal Cost Shocks for Low Inflation (π = 2.5% c.c.)

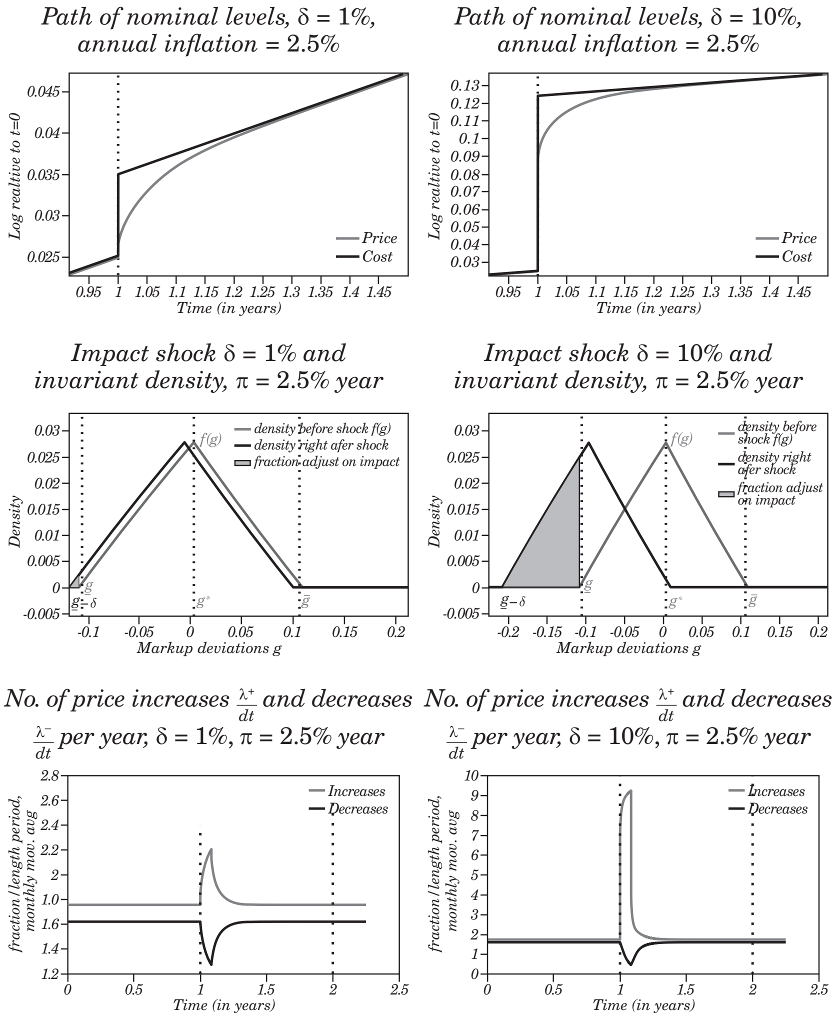

Source: Authors' calculations.

The case of very large inflation of figure 5 with a continuously compounded inflation π = 250% per year is different. As anticipated in the theoretical section, large shocks in an inflationary economy induce an overshooting of the price level on impact as shown in figure 5b.

As inflation is very high, firms that adjust prices (and pay the menu cost) find it optimal to save on future menu costs by raising prices by more than the cost shock. The case of a small shock depicted in figure 5a instead, is similar to Caplin and Spulber (1987), also as expected.

Now we turn to the middle panel of each of the three figures, which in itself is useful to understand the instantaneous passthrough of the top panel just discussed. Note that for low inflation ($\pi = 2.5\%$ c.c. per year), the invariant distribution is almost a tent-map, as it should be for exactly zero inflation. Indeed, it is theoretically known that the effect of steady-state inflation around zero inflation is very small. For high inflation ($\pi = 25\%$ c.c. per year), the invariant distribution is convex-concave, as explained in the theoretical section above. Also as explained in the theoretical section, by comparing the small ($\delta = 0.01$) and large-cost shocks ($\delta = 0.1$) corresponding to the left and right panels respectively, it is seen that the number of firms (or products) that change prices (i.e., the size of the grey-shaded area) increases more than proportionally as the shock increases from 1% to 10%. Alternatively, we can see that the approximation that the impact effect on prices is proportional to the square of the cost shock δ is accurate for this range of shocks. Moreover, the size of the grey-shaded area for the 10% shock in the low inflation rate $\pi = 2.5\%$ case is smaller than in the high inflation rate $\pi = 25\%$ case, due to the convex-concave nature of the invariant distribution for the higher inflation rates. Again, the case of very large inflation ($\pi = 250\%$ c.c. per year) is different.

The share of firms changing prices on impact is much closer to be proportional to δ than in the cases of lower inflation, as can be seen in figure 5d.

Lastly we turn to the behavior of the frequency of price increases and decreases, the bottom panel of the figures for each inflation rate. First, note that for low inflation ($\pi = 2.5\%$ c.c. per year in figure 3) while the rise in the frequency of price increases is moderately larger than the decline in the frequency of price decreases for small-cost shocks ($\delta = 0.01$ in the left panel), this difference is much larger for the case of high-cost changes ($\delta = 0.10$ in the right panel). The pattern is similar in the case of high inflation ($\pi = 25\%$ c.c. per year in figure 4), except that the differences between increases and decreases are a bit more stark. Instead, again, the situation for very large inflation ($\pi = 250\%$ c.c. per year, or figure 5) is different. Since there are almost no price decreases, there is no detectable change in them. On the expected number of price decreases, the behavior is very different between small-cost shocks ($\delta = 0.01$ in the left panel) and large-cost shocks

(δ = 0.1 in the right panel). For small-cost shocks, we have the one-time blip that is characteristic of the mechanism in Caplin and Spulber (1987). Instead, as explained in the theory above, for large-cost shocks, there is overshooting which leads to a subsequent echo effect, which is seen in the damped oscillations in the path of the frequency of price increases.

Figure 4. Passthrough of Nominal Cost Shocks for High Inflation (π = 25% c.c.)

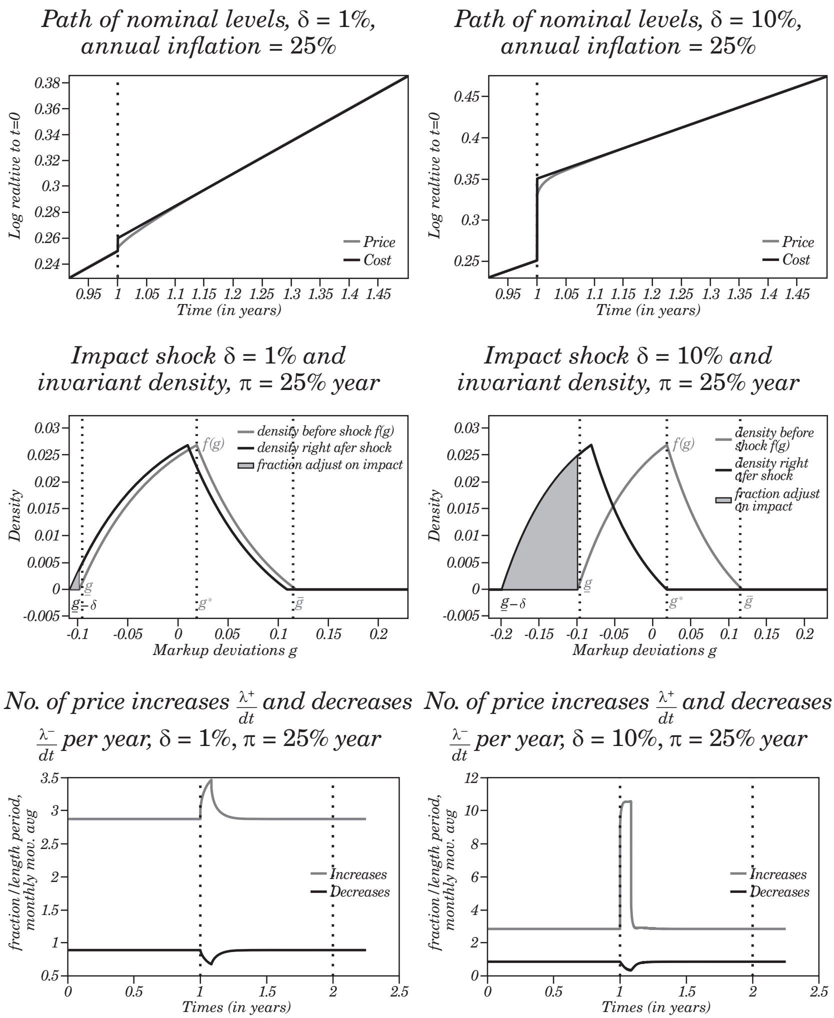

Source: Authors' calculations.

Figure 5. Passthrough of Nominal Cost Shocks for Very High Inflation (π = 250% c.c.)

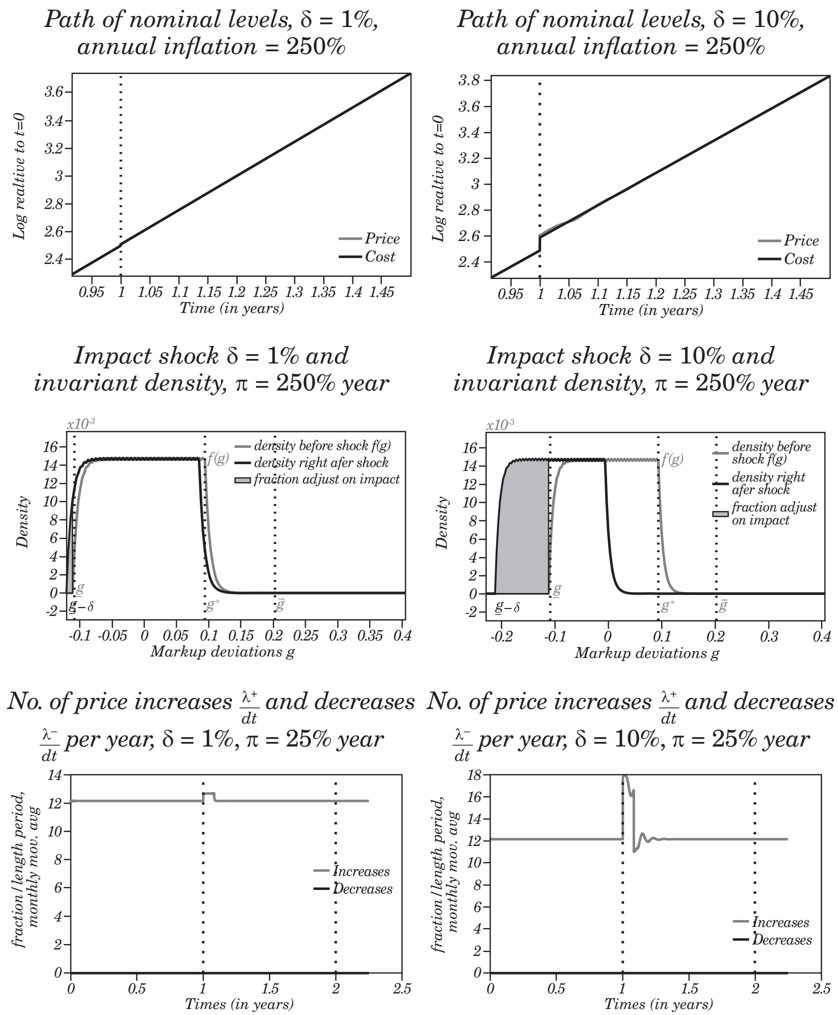

Source: Authors' calculations.

3. Argentina's Evidence on Large-cost Shocks and Price Dynamics

In the previous section we studied how firms facing menu costs of price adjustments react to unexpected shocks to nominal marginal costs. We then studied the aggregate behavior of prices, paying particular attention to the effect of these shocks on the average price level, on the size of price changes, and on the frequency of price adjustment, both on impact and over time. In this section we look at how prices in the city of Buenos Aires reacted to the large nominal shocks described in section 1 and draw conclusions on the applicability of the model.

Figure 6 provides an overview of the behavior of our proxy of the nominal cost and of core prices. We prefer to look at core inflation because, as it excludes seasonal and regulated goods, this measure of prices avoids the mechanical direct impact that changes in regulated prices have on the CPI.

Figure 6. Nominal Costs and Pricve Levels

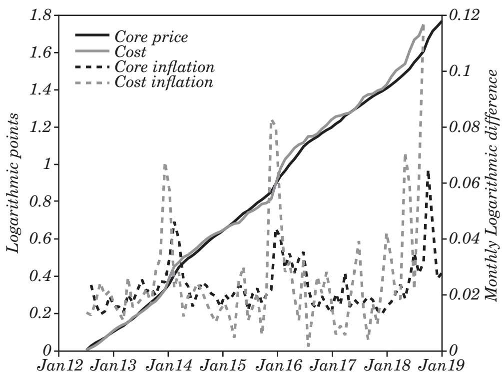

Sources: Prices are from for the city of Buenos Aires Statistical Office. Exchange rates are from the Central Bank, and wages are from *Ministerio de Trabajo, Empleo y Seguridad Social* (2018) [Ministry of Labor, Employment and Social Security].

Note: Cost is a weighted geometric average of regulated prices, the exchange rate and wages with weights 0.1, 0.4 and 0.5, respectively.

There are three large jumps in costs: one in early 2014, a second one in early 2016, and a third one in May–September 2018. The first one is mainly due to a 23% devaluation that took place in the second half of January. As our measure of costs is based on monthly averages, our cost proxy jumps in December 2013 and in January 2014. The 40% weight of tradable goods in our cost measure implies that the jump in cost is slightly above 9%, roughly in synch with the size of price changes in the data[24] and in the simulated examples in subsection 2.5. Figure 6 also shows that there is a spike in inflation associated to each spike in nominal costs. Between November and February 2013, the cost proxy increased by 12%, while the price level increased by 8%. The second shock took place in the first half of 2016. It consisted of a sequence of cost shocks stemming from the impact of the removal of capital controls on the exchange rate and from the change in the relative price of regulated energy prices as shown in table 1. The impact effect relative to the size of the shock is smaller than in the first shock. The persistence of the shocks is reflected in the persistence of the high inflation. Finally, in the first quarter of 2018, there are nominal shocks of about 4% related to regulated prices and, starting in May, there are two exchange-rate lead spikes in nominal costs with peaks of 7% in May and 12% in September.

Several issues prevent us from using the data underlying figure 6 to estimate the impulse response of core prices to cost shocks analyzed in subsection 2.4. (i) At the time a shock hits, prices might still be adjusting to previous shocks. (ii) The cost shock might have been partially anticipated.[25] (iii) An aggregate shock might change the relative price between consumer goods and wholesale goods. (iv) There might be other aggregate cost or productivity shocks not captured by our cost proxy. Nevertheless, we can check if the frequency and if the size of price changes in the data are consistent with the theoretical analysis in 3, illustrated in figures 3 to 5.

Figures 7 and 8 show the frequency and the size of price changes. The figures are based on the data underlying the city of Buenos Aires core consumer price index (IPCBA-rest). The city collects approximately 70,000 prices per month for 628 goods and services.[26]

24. See figure 8.

25. This issue is especially relevant for the devaluation of December 2015. See for example Neumeyer (2015) and Levy-Yeyati (2015).

26. The instructions to the enumerators and other methodological issues are described in *Dirección General de Estadística y Censos* (2013).

The frequency of price changes is computed as the fraction of prices that either increased or decreased between two consecutive observations within the goods and services included in the core measure of inflation. The size of price changes is the geometric equally weighted average of the absolute value of price increases/decreases. The methodology for computing these statistics is described in Álvarez and others (2019), where we discuss the property of this simple estimator and perform robustness checks.

Figure 7 shows the fraction of outlets changing price each month, our proxy for costs, and the core inflation level. Observe first that the average level of the fraction of price increases in "normal" times is 0.24 and the one for price decreases is 0.073. The magnitude of the frequency of price decreases is interesting because it indicates that, for the levels of underlying inflation during 2012–18 in Argentina, the benchmark menu-cost model of Sheshinski and Weiss (1979) is unlikely to be the appropriate one—something that we will also see as we examine other statistics. It rules out the case of $\pi/\sigma \to \infty$ in figure 5, pointing out the importance of idiosyncratic shocks that induce some firms to lower prices even when costs are rising at a cruising speed of 2% per month.

Figure 7. Cost Shocks and the Frequency of Price Changes

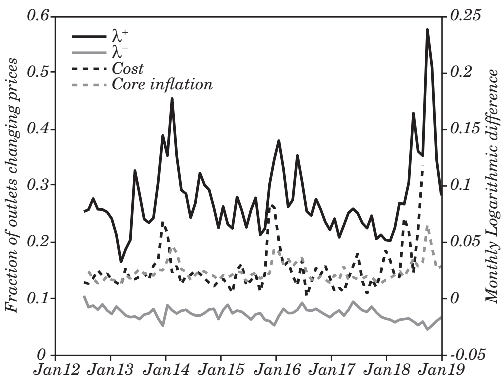

Sources: Prices are from the city of Buenos Aires Statistical Office. Exchange rates are from the Central Bank. Wages are from *Ministerio de Trabajo, Empleo y Seguridad Social* (2018) [Ministry of Labor, Employment and Social Security].
Note: $\lambda+$ is the fraction of outlets increasing prices each month. $\lambda-$ is the analogue for outlets decreasing prices. Cost is a weighted geometric average of regulated prices, the exchange rate and wages with weights 0.1, 0.4 and 0.5, respectively. The dotted line is the core inflation for the city of Buenos Aires (IPCBA - rest).

Second, the three cost shocks identified in our narrative in 2014, 2016, and 2018 are large in the sense that their size is similar to the size of the price changes in normal times. In these three cases, the reaction of the frequencies of price changes is consistent with the predictions of the model for the case of a large shock with high inflation, illustrated in figure 4f. The shock at the beginning of 2014 is short-lived and of a magnitude similar to the $\delta = 0.1$ in the simulated example. As suggested by the theory, the data shows a contemporaneous spike of the fraction of outlets raising prices to 0.45. The transitory decrease in the fraction of outlets lowering prices seems to be in the data, but it is hard to distinguish it from noise. The second episode corresponds to the sequence of cost shocks that took place in the first half of 2016.

These shocks also induced an increase in the fraction of outlets raising prices that is less pronounced and more persistent than the one in 2014. The frequency of price increases peaks at 0.38 in January 2016 and at 0.35 in May 2016. There seems to be a fall in the fraction of outlets lowering prices in the period leading to the devaluation but, again, it is hard to distinguish it from noise in the data. Finally, the fraction of outlets raising prices rises throughout 2018 reacts to the cost shocks, with peaks of 0.43 in June and 0.57 in September, coinciding with important jumps in the exchange rate.

Figure 8 describes the absolute size of price changes conditional on a price change taking place. In "normal" times, the absolute size of positive and negative price changes is similar, around 10–11%. This is consistent with the theory. However, the fact that the absolute value of price decreases is sometimes higher than the one for price increases after 2014 is inconsistent with the theory.[27]

The reaction of the size of price changes to the cost shocks also gives empirical support to the menu-cost model. In the theory, when there is a large-cost shock, many firms end up with a markup that is outside the inaction range, as shown is figure 4d. This implies that the magnitude of price increases has to grow in the presence of large unanticipated shocks. The magnitude of price decreases, on the other hand, should not change as a result of increases in nominal marginal costs, as no firm's markups are pushed above the upper band, $\overline{g}$. These two predictions of the theory are observed in the data. For the three shocks, we observe significant increases in the magnitude of price increases in figure 8, which reach 1.5 times "normal" values.

27. These differences may just be noise.

Figure 8. Cost Shocks and the Size of Price Changes

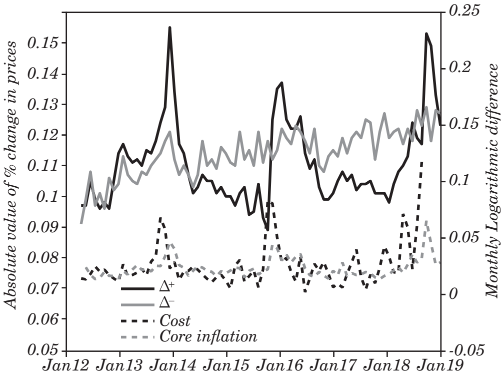

Sources: Prices are from the city of Buenos Aires Statistical Office. Exchange rates are from the Central Bank. Wages are from *Ministerio de Trabajo, Empleo y Seguridad Social* (2018) [Ministry of Labor, Employment and Social Security]. Note: Δ+ is the average percentage change in prices across outlets conditional on a price increase. Δ– is the analogue for price decreases. Cost is a weighted geometric average of regulated prices, the exchange rate and wages with weights 0.1, 0.4 and 0.5, respectively. The dotted line is the core inflation for the city of Buenos Aires (IPCBA - rest).

Also, the magnitude of price decreases does not show abnormal patterns around the time of cost shocks.

We conclude this section by saying that the behavior of the frequency and of the size of price changes supports the passthrough of costs shocks to prices predicted by the menu-cost model of price adjustment. In "normal" times with small shocks, the frequency and the size of price changes are consistent with the simulations for $\pi = 0.25$ and $\delta = 0.01$ in subsection 2.5, thus supporting the choice of parameter values for ψ/B and π/σ used in the simulations. The frequency of price decreases of 0.073 and the fact that the size of price increases and decreases imply that idiosyncratic firm shocks are important. The expected duration of prices, computed as $(\lambda^{+} + \lambda^{-})^{-1}$, is 3.2 months, thus supporting the importance of price rigidities. For the episodes that we interpret as large unanticipated cost shocks, the frequency and the size of price adjustments react as the theory predicts with large increases in the size and in the frequency of price increases, and with no discernible effects on these variables for the case of price decreases. This leads us to conjecture that the fast passthrough of costs to prices, predicted by the theory but hard to estimate in our dataset, is likely to be present in the data.

References

Álvarez, F., M. González-Rozada, A. Neumeyer, and M. Beraja. 2019. "From Hyperinflation to Stable Prices: Argentina's Evidence on Menu-Cost Models." *Quarterly Journal of Economics* 134(1): 452–505.

Álvarez, F. and F. Lippi. 2014. "Price Setting with Menu Costs for Multi-Product Firms." *Econométrica* 82(1): 89–135.

———. 2019. *The Analytic Theory of a Monetary Shock*. Tech. rep.: The University of Chicago.

Álvarez, F., F. Lippi, and J. Passadore. 2016. "Are State- and Time-Dependent Models Really Different?" *NBER Macroeconomics Annual 2016*, volume 31.

Álvarez, F., H. Le Bihan, and F. Lippi. 2016. "The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach." *American Economic Review* 106 (10): 2817–51.

Bassetto, M. and C. Phelan. 2015. "Speculative Runs on Interest Rate Pegs." *Journal of Monetary Economics* 73(C): 99–114.

Bonadio, B., A. Fischer, and P. Sauré. 2016. "The Speed of Exchange Rate Pass-through." Globalization Institute Working Papers 282, Federal Reserve Bank of Dallas.

Caplin, A.S. and D.F. Spulber. 1987. "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics* 102(4): 703–25.

Cavallo, A. 2013. "Online and Official Price Indexes: Measuring Argentina's Inflation." *Journal of Monetary Economics* 60(2): 152–65.

Dirección General de Estadística y Censos, CABA. 2013. *Índice de Precios al Consumidor de la Ciudad de Buenos Aires. Principales Aspectos Metodológicos.*

Gagnon, E. 2009. "Price Setting during Low and High Inflation: Evidence from Mexico." *Quarterly Journal of Economics* 124(3): 1221–63.

Gagnon, E., D. López-Salido, and N. Vincent. 2013. "Individual Price Adjustment along the Extensive Margin." In *NBER Macroeconomics Annual 2012, volume 27*, edited by D. Acemoglu, J. Parker, and M. Woodford. Chicago, IL: University of Chicago Press.

Golosov, M. and R.E. Jr. Lucas. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115: 171–99.

Jones, C. 2011. "Intermediate Goods and Weak Links in the Theory of Economic Development." *American Economic Journal: Macroeconomics* 3(2): 1–28.

Karadi, P. and A. Reiff. 2019. "Menu Costs, Aggregate Fluctuations, and Large Shocks." *American Economic Journal: Macroeconomics* 11(3): 111-46.

Kehoe, P. and V. Midrigan. 2015. "Prices Are Sticky After All." *Journal of Monetary Economics* 75: 35–53.

Levy-Yeyati, E. 2015. "¿De qué hablamos cuando hablamos de dólar e inflación?" *Blog Foco Económico* URL http://focoeconomico.org.

Ministerio de Trabajo, Empleo y Seguridad Social. 2018. *Boletín de Remuneraciones de los Trabajadores Registrados*.

Neumeyer, A. 2015. "El dólar y los precios: Anticipando el fin del cepo." *Blog Foco Económico* URL http://focoeconomico.org.

Sheshinski, E. and Y. Weiss. 1979. "Demand for Fixed Factors, Inflation, and Adjustment Costs." *Review of Economic Studies* 46(1): 31–45.

Appendix A

Derivation of the Profit Function

Define

$$\tilde{F}\left(g,x,z\right) = \Pi\left(g\right)Ax^{1-\eta}z \text{ where } \Pi\left(g\right) = \left[e^{g+m} - 1\right]e^{-\eta(g+m)}.$$

Thus

$$\Pi'(g) = e^{(1-\eta)(g+m)}(1-\eta) + \eta e^{-\eta(g+m)}$$

and evaluating it at $g = 0$

$$\Pi'\left(0\right) = e^{-\eta m}\left[e^{m}\left(1 - \eta\right) + \eta\right] \Rightarrow e^{m} = \frac{\eta}{\eta - 1}.$$

Likewise

$$\Pi''\left(g\right) = e^{(1-\eta)(g+m)}\left(1-\eta\right)^2 - \eta^2 e^{-\eta(g+m)}$$

and evaluating it at $g = 0$

$$\Pi''(0) = e^{(1-\eta)m}\left(1-\eta\right)^2 - \eta^2 e^{-\eta m} = \left(\frac{\eta}{\eta-1}\right)^{1-\eta}\left(1-\eta\right)^2 - \eta^2\left(\frac{\eta}{\eta-1}\right)^{-\eta}$$

$$= \left(\frac{\eta}{\eta-1}\right)^{-\eta}\left[\left(\frac{\eta}{\eta-1}\right)\left(1-\eta\right)^2 - \eta^2\right] = \left(\frac{\eta}{\eta-1}\right)^{-\eta}\left[\eta\left(\eta-1\right) - \eta^2\right]$$

$$= -\left(\frac{\eta}{\eta-1}\right)^{-\eta}\eta = -\left(\frac{\eta-1}{\eta}\right)^{\eta}\eta.$$

The level of the optimized value of Π is:

$$\Pi\left(0\right) = [e^{m} - 1]e^{-\eta m} = e^{(1-\eta)m} - e^{-\eta m} = \left(\frac{\eta}{\eta-1}\right)^{1-\eta} - \left(\frac{\eta}{\eta-1}\right)^{-\eta}$$

$$= \left(\frac{\eta}{\eta-1}\right)^{-\eta}\left[\frac{\eta}{\eta-1} - 1\right] = \left(\frac{\eta}{\eta-1}\right)^{-\eta}\left(\frac{1}{\eta-1}\right).$$

A second order expansion of $\tilde{F}$ around $(0, \overline{x}, \overline{z})$ gives:

$$\tilde{F}(g,x,z) = \Pi'(0)A\overline{x}^{1-\eta}g + \Pi(0)(1-\eta)A\overline{x}^{-\eta}\overline{z}(x-\overline{x}) + \Pi(0)A\overline{x}^{1-\eta}(z-\overline{z})$$

$$+\frac{1}{2}\Pi''(0)A\ \overline{x}^{1-\eta}\overline{z}g^2 + \frac{1}{2}\ \Pi(0)A(1-\eta)(-\eta)\overline{x}^{-\eta-1}\overline{z}(x-\overline{x})^2 + 0\frac{1}{2}(z-\overline{z})^2$$

$$+\Pi'(0)(1-\eta)A\overline{x}^{-\eta}\overline{z}(x-\overline{x})g + \Pi'(0)A\overline{x}^{1-\eta}0(z-\overline{z})g$$
$$+\Pi(0)A\overline{x}^{1-\eta}0(z-\overline{z})(x-\overline{x}) + o\left(\left\|x-\overline{x},g\right\|^2\right).$$

Since $\Pi'(0) = 0$, we have:

$$\tilde{F}(g,x,z) = \Pi(0)(1-\eta)A\overline{x}^{-\eta}(x-\overline{x}) + \frac{1}{2}\Pi''(0)A\ \overline{x}^{1-\eta}g^2$$
$$+\Pi(0)A\overline{x}^{1-\eta}(z-\overline{z})$$

$$+\frac{1}{2}\ \Pi(0)A(1-\eta)(-\eta)\overline{x}^{-\eta-1}(x-\overline{x})^2 + o\left(\left\|x-\overline{x},g\right\|^2\right).$$

Finally, ignoring the terms that are smaller than second order, or that do not involve g, and dividing by $\tilde{F}(0, \overline{x}, \overline{z})$, the maximum profit at $\overline{x}$, we have:

$$F(g) = -\frac{1}{2}\eta(\eta-1)g^2 \equiv Bg^2.$$

Appendix B

Value Function and Optimal Decision Rules

The value function $V(\cdot)$ and the optimal decision rules $\{\underline{g}, g^*, \overline{g}\}$ solve the following system of o.d.e. and boundary conditions:

$$\rho V(g) = F(g) - \pi V'(g) + \frac{\sigma^2}{2} V''(g) \text{ for } g \in [\underline{g}, \overline{g}]$$

as well as value matching and smooth pasting at the boundary of the range of inaction as well as optimality of g^*:

$$V(\underline{g}) = V(g^*) - \psi, V(\overline{g}) = V(g^*) - \psi, \text{and } V'(\underline{g}) = V'(\overline{g}) = V'(g^*) = 0$$

As explained in Álvarez and others (2019), the decision rules depend on ψ/B, π/σ^2 and ρ/σ^2. One can further simplify the problem by considering the case in which $\rho \to 0$, the so called "ergodic" control case. Indeed the solutions are very insensitive to ρ when it takes small values—one can show that the derivative of them with respect to ρ is zero when evaluated at $\rho = 0$. The appendix in Álvarez and others (2019) describes the limit that defines the ergodic control case and characterizes its solutions. One can also compare the solutions obtained numerically with the method described in appendix C, which requires $\rho > 0$; they are almost identical. Hence, to simplify, we will parametrize the solutions by two ratios: ψ/B and π/σ^2.

Appendix C

Numerical Solution of Value Function and Simulated Statistics

An alternative way to compute both the value function and the optimal policies, as well as the effect of cost shocks of the price level, is to discretize time as well as the state space. We consider there a discrete time and discrete state version of the problem. We let Δ be the length of the time period. We will let $\sqrt{\Delta}\sigma$ be the distance between any two points in the state space. For this purpose, we represent the Brownian motion during the range of inaction as:

$$g(t+\Delta)-g(t)=\begin{cases} +\sigma\sqrt{\Delta} & \textit{with probability } p_u = \frac{1}{2}\left(1-\frac{\pi\sqrt{\Delta}}{\sigma}\right) \\ -\sigma\sqrt{\Delta} & \textit{with probability } p_d = \frac{1}{2}\left(1+\frac{\pi\sqrt{\Delta}}{\sigma}\right) \end{cases}$$

so that $\mathbb{E}[g(t+\Delta)-g(t)] = -\pi\Delta$ and $\mathbb{E}[g(t+\Delta)-g(t)]^2 = \Delta\sigma^2$. Thus we can let the state space be $G=\{k\sqrt{\Delta}\sigma\}_{k=-M}^{k=M}$ with typical element $g_k \in G$. The value function is just a vector $V \in \mathbb{R}^{(2M+1)}$. The discretized version solves the following Bellman equation:

$$V_k = \max\left\{\psi + \max_j V_j, \Delta F(g_k) + \frac{1}{1+\Delta\rho}[p_u V_{\max\{k+1,K\}} + p_d V_{\min\{k-1,-K\}}]\right\}$$

for $k = -K, -K+1, \ldots, K-1, K$. The choice of K should be large enough so that the range on inaction is well in the interior of G. The choice of Δ should be small enough so that it approximates well the continuous time limit. The solution of this problem will give the thresholds $\underline{g}$ and $\overline{g}$, as well as $k^* = \operatorname{argmax}_k\{V_k\}$ so that $g^* = g_{k^*}$.

We let f be a $2K+1$ positive vector containing the fraction of firms with different price gaps. For instance, f_k will be the fraction of firms with price gap g_k. The law of motion for the vector with the distribution of firms $f(t+1)$ can be represented by a square stochastic matrix L of $2K+1 \times 2K+1$ dimensions, so $f(t+\Delta) = Lf(t)$, where we use $f(t)$ for the column vector of $2K+1$ values of the fractions of firms $f_k(t)$ at time t with price gap g_k. The matrix L can be thought as the sum of two matrices. The first matrix has zeros in all the entries, except in the entries next to the diagonal where it has either p_u or p_d, keeping tract of the mass of firms within the inaction region. The second matrix is also sparse—it has zeros and ones, with ones indicating that the firms that are outside the range of inaction will transit to optimal return point k^* with probability one. Recall that we have: $f_j(t+\Delta) = \sum_{k=-K,K} L_{j.k} f_k(t)$. Let $\underline{k}$, $k^*$and $\overline{k}$ be the indices of the elements of f that correspond to $\underline{g}$, $g^*$and $\overline{g}$, respectively.

The elements $L_{i,j}$ are zero, except in the following cases.

1. For $k = -K,\ldots, \underline{k} - 1)$, then $L_{k^*,k} = 1$,
2. For $k = \overline{k} + 1,\ldots, K$, then $L_{k^*,k} = 1$,
3. For $k = \underline{k} + 1,\ldots, \overline{k} - 1$, then $L_{k,k+1} = p_d$ and $L_{k,k-1} = p_u$,
4. $L_{\underline{k},\underline{k}+1} = p_d$ and $L_{\overline{k},\overline{k}-1} = p_u$, and
5. $L_{k^*,\underline{k}} = p_d$ and $L_{k^*,\overline{k}} = p_u$.

As it is well known, the invariant distribution can be obtained by computing the powers of L^T for a large value of T. Alternatively, the invariant distribution is the eigenvector associated with the eigenvalue equal to one of matrix L.

The fraction of price changes per period of length Δ is given a $\lambda_a^+(t) = \frac{1}{\Delta}\, l^+ f(t)$, where l^+ is a vector that adds the components of the vector $f(t)$ with values below $\underline{g}$. In particular, the row vector l^+ has:

1. $l_k^+ = 1$ for $k = -K,\ldots, \underline{k}-1$,
2. $l_{\underline{k}}^+ = p_d$, and
3. $l_k^+ = 0$ for $k = \overline{k} + 1,\ldots, K$.

We can define $\lambda_a^-(t) = \frac{1}{\Delta}\, l^- f(t)$ in an analogous way. The row vector l^- has:

1. $l_k^- = 1$ for $k = \overline{k} + 1,\ldots, K$,
2. $l_{\overline{k}}^- = p_u$, and
3. $l_k^- = 0$ for $k = -K,\ldots, \overline{k} - 1$.

Thus, the fraction of firms that change prices between t and $t + \Delta$ per unit of time is $\lambda_a(t) = \lambda_a^+(t) + \lambda_a^-(t)$. Note that the definition of λ_a, λ_a^+ and λ_a^-, divides the fraction of firms that adjust prices by Δ so that it is comparable, at least for small Δ, to the continuous time expressions. Hence, if we multiply the expressions for $\lambda_a(t)$, $\lambda_a^+(t)$, or $\lambda_a^-(t)$ by Δ, we will obtain fraction of firms that change prices during the interval t and $t + \Delta$.

Discrete-time – discrete-state impulse response. We can write the equivalent of the impulse response on prices using the matrix L.

We start with $\bar{f}$ given by the invariant distribution implied by L. Then we let $\hat{f}$ (0; δ) be the $2K$+1 vector of the displaced initial conditions. This is a shift of the mass to the left by an amount equal δ. This vector is given by:

$$\hat{f}(0;\delta) = \bar{f}_{k+v} \text{ for } k = -K,\ldots,K-v \text{ and } \hat{f}_k(0;\delta) = 0 \text{ otherwise,} \tag{12}$$

and where $v = \delta / \left[\sigma\sqrt{\Delta}\right]$ which, for simplicity, we assume to be an integer. Now we can compute the impact effect Θ:

$$\Theta = \sum_{k=-K}^{\bar{k}-1}\left(g_{k^*} - g_k\right)\hat{f}_k(0;\delta) - \sum_{k=\bar{k}+1}^{K}\left(g_k - g_{k^*}\right)\hat{f}_k(0;\delta). \tag{13}$$

Then we compute the sequence of distribution of firms indexed by their price gap as:

$f(\Delta) = [L]^j \hat{f}$ (0; δ) for j = 1, 2,...

With these elements we compute the rate of increase in prices θ's:

$$\theta(j\Delta) = \Delta_p^+\lambda_a^+(j\Delta) - \Delta_p^-\lambda_a^-(j\Delta) \text{for } j = 1,2,\ldots \tag{14}$$

where

$$\lambda_a^+(j\Delta) = \frac{1}{\Delta}l^+ f(j\Delta) \quad \text{for } j = 1,2,\ldots, \tag{15}$$

and

$$\lambda_a^-(j\Delta) = \frac{1}{\Delta}l^- f(j\Delta) \quad \text{for } j = 1,2,\ldots \tag{16}$$

The analog of the continuous time impulse response is:

$$P(j\Delta) = \bar{P} + \pi\Delta + \Theta + \sum_{k=1}^{j}\theta(k\Delta)\Delta \quad \text{for} \quad j = 1,2,\ldots$$

where $\bar{P}$ is the price level in steady state just before the cost shock, or $\bar{P} = P(-\Delta)$.

To simplify the exposition we have not taken into account the regular idiosyncratic shocks that also occur during the period between t = 0 and t = Δ. To correct for this effect we have included πΔ to the initial value of $\bar{P}$, which is the correct value, since we are starting with the invariant distribution. This has the interpretation that the shock δ occurs at the end the discrete period, but before the next period. As Δ → 0, these effects can be neglected.

The Nonpuzzling Behavior of Median Inflation

Laurence Ball
Johns Hopkins University

Sandeep Mazumder
Wake Forest University

For decades, textbooks have explained inflation behavior with Friedman (1968)'s Phillips curve: the inflation rate depends on expected inflation and the deviation of unemployment from its natural rate. Yet this theory has always been controversial, and skepticism has been rampant in the decade since the 2008 financial crisis. For several years following the crisis, researchers such as Stock (2011) and Coibion and Gorodnichenko (2015) puzzled over a "missing deflation:" inflation did not fall much despite a sharp rise in the unemployment rate. More recently, as the economy has approached full employment, economists have puzzled over the failure of inflation to rise toward the Federal Reserve's target of 2 percent. According to Bernstein (2017), recent low inflation is "puzzle #1 in economics."

Some observers, such as Summers (2017) and The Economist (2017), have lost patience with the Phillips curve and suggested it is "broken." Blinder (2018) wonders "whether the Phillips curve has died or has just taken an extended vacation." Blanchard (2016) offers a tepid defense of the theory, by saying the Phillips curve is alive and "at least as well as it has been in the past." Blanchard emphasizes that the residuals in the relationship are large.

This paper argues that inflation behavior is less puzzling if we separate the headline-inflation rate into two components: an

This paper was prepared for the XXII Annual Conference of the Central Bank of Chile in October 2018. Thank you to our discussant, Carlos Carvalho, and to conference attendees for their comments. We also thank Randy Verbrugge and Tristan Young for providing data, and Xu Liu and Jionglin Zheng for research assistance.

Changing Inflation Dynamics, Evolving Monetary Policy edited by Gonzalo Castex, Jordi Galí, and Diego Saravia, Santiago, Chile.

underlying or core level of inflation that the Phillips curve explains, and a transitory component arising from changes in relative prices due to microeconomic factors. A good proxy for the core-inflation rate is a measure proposed by Bryan and Cecchetti (1994): the weighted median of price changes across industries.

Many previous researchers, and the policymakers at the Federal Reserve, examine core inflation in an effort to filter out transitory shocks. However, the usual measure of core inflation is the inflation rate excluding the prices of food and energy. This variable filters out shocks in the food and energy industries, but many other industries also experience large price changes that materially influence the headline-inflation rate. The weighted median filters out all of these shocks and produces a less noisy measure of core inflation whose movements are easier to understand.

Section 1 of this paper briefly reviews the theoretical case for measuring core inflation with the weighted median, and the previous empirical literature. Section 2 begins our empirical work by examining the univariate behavior of alternative measures of core inflation. We show that, for both the CPI and PCE (Personal Consumption Expenditure)-deflator versions of inflation, the weighted median of industry price changes is less volatile than inflation excluding food and energy.

Section 3 illustrates the usefulness of weighted-median inflation with a careful study of inflation over 2017 and early 2018. Some observers believe that inflation behavior was especially puzzling during that period: In particular, despite low unemployment rates, the Fed's preferred measure of core inflation—the twelve-month percentage change in the PCE deflator less food and energy—fell from 1.9 in December 2016 to 1.3 in August 2017 and to 1.5 in December. At a September press conference, Fed Chair Yellen said that low inflation before 2017 was consistent with the Fed's specification of the Phillips curve, but: "This year, the shortfall of inflation from 2 percent [...] is more of a mystery, and I will not say that the [Federal Open Market] Committee clearly understands what the causes are of that."

We show that this mystery disappears if we measure core inflation with the weighted-median inflation rate rather than inflation less food and energy. The weighted median does not fall significantly over 2017 because it filters out price decreases in a number of industries that pushed down the Fed's core-inflation measure. Examining the weighted median also helps resolve confusion among policymakers about an apparent uptick in core inflation in early 2018.

Section 4 turns to the Phillips curve. We examine the fit of a simple specification in which quarterly core inflation depends on expected inflation (as measured by long-term forecasts from the Survey of Professional Forecasters) and the cyclical component of unemployment (as measured by the Hodrick-Prescott filter). We first measure core inflation with inflation less food and energy and see the source of recent skepticism about the Phillips curve: the equation fits the data poorly, especially for inflation in the PCE deflator and especially since 2008. We then see that the Phillips curve shows up clearly when core inflation is measured more precisely with weighted median inflation.

All in all, our results suggest that economists should use the weighted-median or related variables (such as trimmed means of industry price changes) as their primary measures of core inflation. Researchers should also work on refining these measures. Section 5 concludes this paper by discussing directions for future research.

1. Background

According to the Phillips curve, the inflation rate depends on expected inflation and the level of slack in the economy. Economists often suggest, however, that inflation movements are also influenced by price changes in certain industries. We will discuss, for example, Chair Yellen's view that large price decreases for cell-phone services and prescription drugs reduced inflation during 2017. In earlier episodes, economists have explained high inflation with rising medical costs, and low inflation with falling prices of imported goods.

The practice of explaining aggregate inflation with industry price changes can, however, be dangerous. There are always some prices that rise by significantly more than the aggregate inflation rate and others that rise by less or fall; that is, there are always changes in relative prices. If the inflation rate is higher than the Phillips curve predicts, one can always find a cheap "explanation" by citing industries whose prices have risen by more than average; in turn, low inflation can be explained by industries with price decreases. To avoid such vacuity, we need a theory of which relative-price changes truly affect aggregate inflation.

Ball and Mankiw (1995) present such a theory, one in which relative-price changes matter if they are unusually large. This result arises because, with costs of nominal price adjustment, large shocks to industries' optimal prices induce them to change their actual prices, while prices are sticky in response to smaller shocks.

The disproportionate effects of large shocks imply that inflation is influenced by asymmetries in the distribution of price changes across industries. If there is a tail of unusually large price increases, which skew the distribution to the right, it raises inflation; in turn, a tail of large price decreases does the opposite. Ball and Mankiw find strong support for these predictions in U.S. data.

Measures of core inflation are intended to filter out the effects on headline inflation of unusual relative-price changes, thereby isolating the component of inflation explained by the Phillips curve. In pioneering work, Bryan and Cecchetti (1994) develop a measure of core inflation by extending the reasoning of Ball and Mankiw. If asymmetries in the price-change distribution cause fluctuations in headline inflation, then one can measure core inflation by eliminating the effects of these asymmetries. A simple variable that does so is the median of industry price changes, weighted by industries' relative importance in the aggregate price index.

The traditional measure of core inflation is the inflation rate excluding food and energy prices. In the U.S. economy, many of the large relative-price changes that influence inflation occur in the food and energy industries (especially energy), so dropping those industries is a step toward isolating the core level of inflation. However, large relative-price changes also occur in industries other than food and energy. Based on the disaggregated PCE deflator, Dolmas (2005) reports that large price changes are common in industries such as computers and software, televisions, clothing, airline services, financial services, and auto insurance. As we will see in our empirical work, filtering out large shocks to all industries with the weighted median yields a core-inflation measure that is less volatile and easier to understand than inflation less food and energy.

Weighted-median measures of core inflation—as well as trimmed means of industry price changes, which also filter out large shocks—have gained increasing attention in recent years. In 2016, the Bank of Canada announced that it would include a weighted median and a trimmed mean among its primary measures of core inflation. Yet most researchers still define core inflation as inflation excluding food and energy. Staff at the Federal Reserve produce forecasts of PCE-deflator inflation less food and energy, and this variable is a focus of FOMC meetings and speeches by Fed officials. We hope that this paper helps push economists and policymakers toward changing their measures of core inflation.

This paper studies the behavior of two versions of weighted-median inflation. One is the weighted-median CPI inflation rate published by the Cleveland Fed, which is currently based on dividing the basket of goods in the CPI into 45 industries. The other is a weighted-median PCE-deflator inflation rate that we have constructed from data on 178 industries provided by the Dallas Fed. Researchers at Dallas use these data to construct a trimmed-mean measure of core inflation; we construct a weighted median instead for comparability with the median CPI series. The relative merits of the weighted-median and trimmed-mean measures of core inflation are an important topic for future research.[1]

2. Univariate Evidence

This section examines the univariate behavior of headline inflation, inflation excluding food and energy, and weighted-median inflation. We examine the period 1985–2017. We find that both of the core-inflation measures filter out much of the transitory variation in headline inflation, but that the weighted median filters out more and is therefore less volatile.

Table 1 measures the volatility of each inflation series with the standard deviation of the change in inflation. We compute this statistic for both the CPI and PCE-deflator versions of inflation. We examine annualized monthly inflation rates, annualized quarterly inflation rates, and a monthly series on the inflation rate over the previous twelve months.[2]

The results in the table are consistent across the two price indices and the three data frequencies: the standard deviation of changes in inflation is smaller for inflation less food and energy than for headline inflation, but smaller still for weighted-median inflation. The ratio of the standard deviations of changes in ex-food-energy and median inflation range from 1.4 to 1.6 (except for monthly PCE data, where the ratio is higher because of an outlier discussed below).

1. A number of previous researchers advocate weighted medians or trimmed means as measures of core inflation because these variables are strongly correlated with an underlying trend in headline inflation, or because they are good forecasters of future inflation. Examples include Bryan and others (1997), Clark (2001), Smith (2004), Brischetto and Richards (2006), and Ball and Mazumder (2011). Crone and others (2013) question the value of medians and trimmed means for forecasting.

2. The series for median CPI inflation from the Cleveland Fed is monthly, and our series for median PCE inflation is derived from monthly data on industry inflation rates. We multiply monthly inflation by 12 to produce annualized inflation rates. To derive annualized quarterly inflation rates, we convert monthly inflation to monthly price levels, average over three months to get quarterly price levels, compute the percentage change from the previous to the current quarter, and multiply by four.

Table 1. Volatility of Alternative Inflation Measures

	Monthly	*Quarterly*	*12-Month*
Headline CPI	3.278	2.307	0.387
CPIX	1.403	0.653	0.131
Median CPI	0.916	0.447	0.095
Headline PCE	2.408	1.567	0.268
PCEX	1.633	0.681	0.134
Median PCE	0.868	0.436	0.085

Source: Authors' calculations.
Note: The numbers in the table are standard deviations of the change in the annualized inflation rate over 1985–2017. The monthly numbers for headline PCE, PCEX, and median PCE inflation are 2.36, 1.36, and 0.89, respectively, when September–November 2001 are excluded.

Figure 1. CPI and PCE Core Monthly Inflation

A. CPI

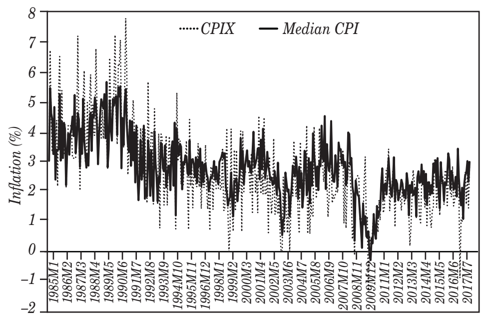

B. PCE

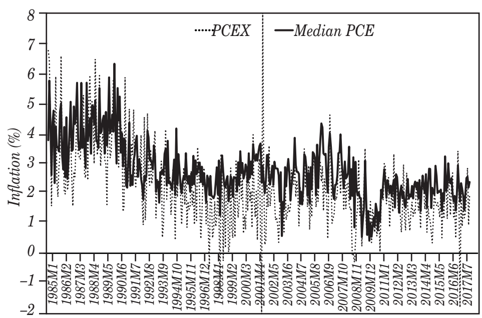

Source: Authors' calculations.

To illustrate these results, figure 1 presents the monthly time series for the two measures of core inflation; in figure 1(a), both are based on the CPI price index, and in figure 1(b) they are based on the PCE deflator. We can see the greater volatility of the ex-food-energy measure of core. In the CPI case, for example, in the late 1980s and early 1990s, median inflation generally fluctuates in a range of about 3–5 percent; inflation less food and energy (CPIX inflation) is often in the same range but spikes up to 6 or 7 percent in a number of months. Stating in the late 1990s, CPIX inflation spikes downward to zero or below in a number of months, whereas median inflation falls that far at only one point (February and March 2010).

The PCE-deflator graph also shows that ex-food-energy inflation is more volatile than median inflation. Some of the months with outliers in PCE inflation less food and energy (PCEX) are also outliers in CPIX inflation (such as March 2017, an observation that we examine closely below). But other times, the outlier months differ for CPIX and PCEX. For example, CPIX inflation spikes down to 0.2 percent in April 2013 and then rises to 2.7 percent in July 2013; PCEX inflation is more stable, with rates of 0.7 percent in April and 1.2 percent in July. Evidently, movements in ex-food-energy measures of core inflation can differ due to differences in the industries covered by the CPI and PCE deflator and/or differences in how industry prices are measured.

One episode produces large outliers in the PCEX data: the annualized inflation rate falls to –6.6 percent in September 2001 and then jumps to 8.6 percent in October. These numbers reflect huge transitory movements in life insurance premiums, which could be related to the September 11 terrorist attacks. Life insurance premiums fell at an annualized rate of 655 percent in September and then rebounded at a rate of 1457 percent in October. These price changes were large enough to strongly influence monthly PCEX inflation rates. Weighted-median inflation, by contrast, filters out this episode along with less dramatic shocks to industry prices.

Figure 2 compares our two versions of median inflation: median CPI inflation and median PCE inflation. Usually, the two medians move together fairly closely: it appears that they isolate more or less the same underlying level of inflation, despite the differences between the CPI and PCE price indices. The standard deviation of the difference between median CPI and median PCE inflation is 0.7, compared to a standard deviation of 1.2 for the difference between CPIX and PCEX inflation.

Figure 2. Median CPI and Median PCE Monthly Inflation

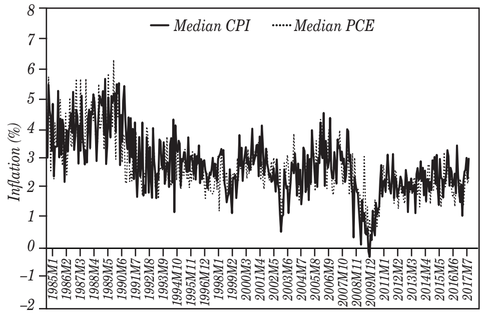

Source: Authors' calculations.

As figure 2 suggests, the average levels over time of the two medians are close. For 1985–2017, median CPI inflation averages 2.8 percent and median PCE inflation averages 2.7 percent. By contrast, it is well known that the average levels of headline and ex-food-energy inflation are higher for the CPI than for the PCE deflator. For 1985–2017, the averages of CPIX and PCEX inflation are 2.6 and 2.2 percent, respectively. For the PCE, the fact that the average of median inflation (2.7 percent) significantly exceeds the average of PCEX inflation (2.2 percent) suggests a tendency toward left skewness in the distribution of industry inflation rates. The reason for such a pattern is unclear and might be a subject for future research.

3. A Case Study: Inflation in 2017–2018

Recent history helps us understand the usefulness of weighted-median inflation. During 2017, the Fed's primary measure of core inflation, the 12-month inflation rate in the PCEX, fell noticeably despite low unemployment, a development that Fed Chair Yellen called a "mystery".[3] In trying to explain this mystery, Yellen stated "there have been some idiosyncratic factors I think that have held down inflation in recent months" including price changes in several industries.[4] We find that inflation in 2017 is less mysterious if we

3. See Yellen (2017a).
4. See Yellen (2017b).

examine the weighted median, which filters out unusual price changes systematically. Examining the weighted median also resolves a puzzle about an uptick in PCEX inflation in early 2018.

Figure 3 shows inflation rates for the PCEX and median PCE from January 2017 through March 2018. Panel (a) shows inflation over the previous 12 months, which is the focus of many discussions by economists and policymakers. We see the behavior of 12-month PCEX inflation that puzzled the Fed: This inflation rate fell from 1.9 percent in January to 1.3 percent in August and 1.5 percent in December, a period when the unemployment rate fell from 4.8 to 4.1 percent. In discussing this experience in September, Chair Yellen said, "I will not say that the [FOMC] clearly understands what the causes are."

The behavior of 12-month median PCE inflation is different. We see that this inflation rate starts three tenths of a percent above PCEX inflation and stays above it, in line with our earlier finding that average median PCE inflation is modestly higher than average PCEX inflation. For our purposes, however, the key fact about 12-month median inflation is that it is stable: it stays in a range from 2.0 to 2.2 throughout 2017. Policymakers would not have perceived a puzzling decline in core inflation if the median were their measure of core.

Figure 3. PCEX and Median PCE Inflation, January 2017–March 2018

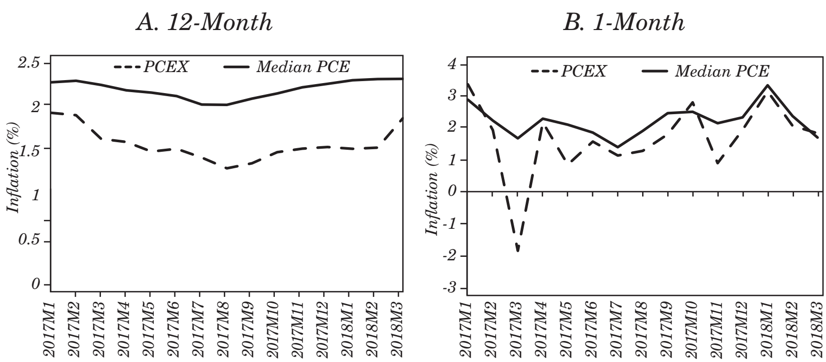

Source: Authors' calculations.

Panel (b) of figure 3 shows the one-month inflation rates underlying the smoother 12-month rates in Panel (a). For PCEX, we see an important outlier: March 2017, when the PCEX inflation rate was –1.8 percent. This rate is 3.8 points below the Fed's inflation target of 2.0, so for 12-month periods including March 2017, that month pushes inflation below the target by approximately (3.8)/12 = 0.32 points. Other months in 2017 that pull down the 12-month rate are May and November, which each have a PCEX inflation rate of 0.9. For median PCE, by contrast, one-month inflation rates in 2017 stay in a relatively narrow range from 1.4 to 2.9, thus leading to a very stable series when these rates are averaged over twelve months.

For the influential month of March 2017, figure 4 shows a histogram of industry price changes within the PCEX index. Each bar in the graph represents an interval of 5 percentage points in annualized inflation rates and shows the total weights of the industries in that range. We see a tail of large price decreases that skews the histogram to the left and pulls down PCEX inflation. Industries with sizable weights in the PCEX and highly negative inflation rates include air transportation (weight of 0.5 percent and annualized inflation rate of –65 percent), communications (weight of 2.5 percent and inflation rate of –38 percent), hotels and motels (weight of 0.9 percent and inflation rate of –34 percent), and men's and boys' clothing (also weight of 0.9 percent and inflation rate of –34 percent). Large price decreases also occur in smaller industries such as watches and videocassettes, and discs.

In a series of speeches and news conferences in 2017, officials from the Federal Reserve sought to explain the low level of PCEX inflation. On several occasions (in May, June, September, and October), Fed officials cited a large decline in the quality-adjusted prices of cell-phone service that occurred when cell-phone companies introduced unlimited data plans.[5] In June, Fed Chair Yellen also mentioned a drop in prescription-drug prices, and in October she mentioned slow growth in medical costs in general. In September, she suggested that "a variety of special factors" had restrained inflation.

In these remarks, Fed officials are trying in a haphazard way to do what the weighted-median inflation rate does more easily and systematically: uncover a stable level of underlying inflation by

5. See Brainard (2017), Yellen (2017b), Evans (2017), and Yellen (2017c).

filtering out unusual industry price changes. Yellen is right about "a variety of factors": many different industries contributed to the negative PCEX inflation of March 2017, and others contributed to the low inflation of May and November. Officials are also on target in specifically mentioning cell phones, which are a significant factor in the March outlier. The March inflation rate was –84 percent for cell-phone services and –38 percent in the broader communications sector.

On the other hand, Yellen's reference to prescription drugs is puzzling. Prices in that industry rose at an annual rate of 4.7 percent in March and 3.4 percent over the 12 months of 2017, numbers that go in the wrong direction for explaining low inflation. Yellen is correct that some medical industries experienced low inflation in 2017—the prices of physician's services, for example, rose by 0.5 percent over the 12 months. However, this inflation rate is only modestly lower than aggregate PCEX inflation, and theory suggests that only large relative-price changes are relevant. In explaining aggregate inflation, it is suspect to point out industries whose inflation rates are modestly higher or lower than average, because there are many such industries at all times.

Figure 4. Histogram of Industry Price Changes in March 2017

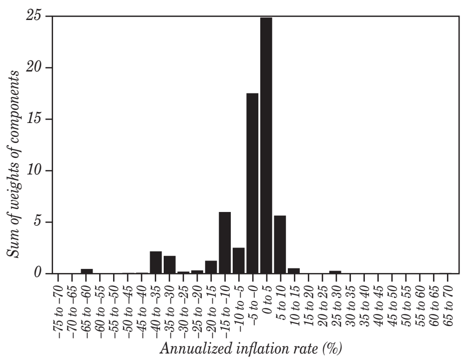

Source: Authors' calculations.

Note: The vertical axis is cut off at 25—the sum of industry weights in the 0 to 5 percent inflation range is 45.6. Food and energy industries are excluded. Industries in tails: air transportation (–65 to –60), watches (–50 to –45), video cassettes and discs (–40 to –35), communication (–40 to –35), and children's and infants' clothing, hotels and motels, and men's and boys' clothing (–35 to –30) in the lower tail, and prerecorded and blank audio discs/tapes/digital files/downloads (–65 to –70) in the upper tail.

We conclude that it would have been easier for the Fed to accurately interpret core-inflation movements in 2017 if its measure of core had been weighted-median inflation rather than PCEX.[6]

A focus on median inflation might also have clarified the Fed's analysis of inflation in early 2018. In the minutes of the FOMC meeting held on May 1, some participants suggest that inflation is likely to overshoot the Fed's 2 percent target, noting "the recent increase in inflation." This increase is presumably the jump in 12-month PCEX inflation from 1.5 percent in February to 1.8 percent in March, the last month for which the Committee had data. Other Committee members question the importance of the increase, saying "it may have represented transitory price changes in some categories of health care and financial services."

This reference to industry price changes, like some of Yellen's remarks in 2017, is questionable. In the first three months of 2018, price changes in health care industries were unremarkable. As the minutes suggest, the prices of financial services rose substantially in March: the annualized inflation rate for financial charges, fees, and commissions was 24 percent. But the effect on aggregate inflation was modest. The weight on financial fees in the PCEX is 2.6 percent, which means the 24 percent inflation rate contributed approximately 0.6 percentage points to PCEX inflation in March, and only 0.05 points to 12-month inflation.

What then explains the March uptick in 12-month PCEX inflation? The answer is that March 2018 is the month when the –1.8 percent inflation rate of March 2017 drops out of the 12-month average and is replaced by the current monthly rate of 1.9 percent. Some journalists, such as Rugaber (2018) and Mutikani (2018), note the role of March 2017 in explaining 12-month inflation a year later, but this point does not appear in the FOMC minutes.

Once again, there is less inflation variability to explain, and potentially become confused about, if we focus on weighted-median inflation. Over the first three months of 2018, there are no outliers in the monthly median inflation rates that enter or exit the 12-month average. The 12-month inflation rate is stable at 2.2 percent.

6. In her September speech, Yellen briefly mentions that trimmed-mean inflation has fallen by less than PCEX inflation, which is some acknowledgment of the usefulness of systematically filtering out large industry price changes.

4. Phillips Curves

Many of the economists who have puzzled over recent inflation behavior emphasize the apparent absence of an unemployment–inflation relationship consistent with a textbook Phillips curve. Here we ask how well a simple Phillips curve fits quarterly data since 1985, and especially whether the relationship has broken down since the onset of the Great Recession in 2008. The answers depend on how inflation is measured. With headline inflation, there is no discernable Phillips curve. With core inflation as measured by the CPIX or PCEX, the evidence is mixed and we can see why many analysts would not find a Phillips curve or would think it has broken down. With weighted-median inflation, by contrast, the data show a clear and robust Phillips curve that remains stable after 2008.

4.1 Specification

We consider a simple version of Milton Friedman (1968)'s expectations-augmented Phillips curve, in which the inflation rate depends on expected inflation and on deviations of unemployment from its natural rate. Specifically, in quarterly data, we assume

$$\pi_t = \pi_t^e + \alpha\overline{(u-u^*)}_t + \varepsilon_t, \tag{1}$$

where π is inflation, π^e is expected inflation, and $\overline{(u-u^*)}_t$ is the average of the unemployment rate, u, minus the natural rate, u^*, from $t-3$ through t. Our inclusion of three unemployment lags follows previous research on the Phillips curve.[7] For parsimony, we assume the coefficients on the current and three lags of $u-u^*$ are all the same, so only the average of these terms appears in the equation (a restriction that the data do not reject).

Again following previous work,[8] we measure expected inflation with long-term inflation forecasts, specifically, the mean of ten-year forecasts from the Survey of Professional Forecasters (SPF). When we measure inflation with any version of the Consumer Price Index (whether headline or one of the core measures), we use ten-year forecasts of CPI inflation. When we measure inflation with the PCE deflator, we have the problem that ten-year SPF forecasts of PCE

7. See Stock and Watson (2010).
8. See Fuhrer and Olivei (2010); Ball and Mazumder (2019).

inflation only started in 2007. We use these PCE forecasts when they are available. As a proxy for PCE expectations before 2007, we use the forecasts of CPI inflation minus the average difference between CPI and PCE forecasts when both are available (which is 0.23).

We measure the natural rate of unemployment, u^*, with the trend in unemployment from the Hodrick-Prescott filter with a smoothing parameter of 1600. We eschew more sophisticated methods, such as Staiger and others (1997), which use inflation and unemployment data to estimate u^* along with the parameters of an assumed Phillips curve. This approach can bias the estimates of u^* in the direction of fitting a Phillips curve relationship even if none exists—a problem that does not arise with our univariate estimation of u^*.

To estimate the Phillips curve, we move expected inflation to the left side of the equation and estimate:

$$\pi_t - \pi_t^e = \alpha\overline{(u - u^*)}_t + \varepsilon_t . \tag{2}$$

This equation does not include a constant term: when $u - u^*$ is zero, Friedman's Phillips curve says $\pi - \pi^e$ should also be zero. However, if we add a constant to the equation, we sometimes find it is statistically significant, so we present estimates both with and without a constant. Arguably, one test of Friedman's theory is whether the estimated constant is close to zero. We do not put too much weight on this test, however, because a constant might reflect measurement error in π^e or u^* with a non-zero mean.[9]

4.2 Estimates for 1985–2017

Table 2 presents Phillips curve estimates with inflation measured with the CPI—panel (a)—and with the PCE deflator—panel (b). For each of these price indexes, we compare results for headline inflation and the two measures of core inflation: inflation less food and energy (CPIX or PCEX) and weighted-median inflation. For the two core measures, figures 5 and 6 present scatterplots of the data underlying our regressions.

9. In particular, the HP filter forces the mean of u^* to equal the mean of u. Other estimates suggest that u^* and u have different means over our sample period of 1985–2017; for example, the mean of the u^* series produced by the Congressional Budget Office exceeds the mean of u by 0.78 percentage points.

Table 2. Phillips Curves for 1985–2017

$$\pi_t - \pi_t^e = \alpha\overline{(u - u^*)}_t + \varepsilon_t$$

a. CPI inflation						
	Headline		*CPIX*		*Median*	
Constant		-0.355		-0.319		-0.167
		(0.173)		(0.065)		(0.061)
α	-0.195	-0.224	-0.424	-0.450	-0.648	-0.661
	(0.312)	(0.331)	(0.181)	(0.128)	(0.117)	(0.093)
$\overline{R}^2$	-0.031	-0.002	-0.052	0.216	0.408	0.480
S.E. of Reg.	1.884	1.857	0.627	0.541	0.468	0.439

b. PCE inflation						
	Headline		*PCEX*		*Median*	
Constant		-0.533		-0.531		0.017
		(0.138)		(0.063)		(0.062)
α	-0.093	-0.136	-0.201	-0.244	-0.478	-0.477
	(0.233)	(0.264)	(0.148)	(0.080)	(0.078)	(0.079)
$\overline{R}^2$	-0.156	-0.003	-0.785	0.066	0.319	0.315
S.E. of Reg.	1.435	1.337	0.768	0.555	0.445	0.446

Source: Authors' calculations.
Note: OLS with robust (HAC) standard errors is used (standard errors in parentheses). The unemployment gap is the deviation of the unemployment rate from the HP filtered series, where the filter is applied over 1948–2017.

These results make it clear, first, that the fit of the Phillips curve is highly sensitive to the choice between headline and core inflation. It is easy to see why someone who focuses on headline inflation would doubt that the Phillips curve exists. For both headline CPI and headline PCE, the Phillips curve slope α is insignificant and the $\overline{R}^2$ of the equations are negative (either with or without a constant). The noise in quarterly headline inflation obscures any underlying Phillips curve.

We can also see that the choice between the two core-inflation measures is important—to a substantial degree for CPI inflation and even more for PCE inflation. For CPIX, the Phillips curve slope is significant at the 5 percent level, but the $\overline{R}^2$ is negative with no constant term and only 0.22 with a constant. The fit is better with weighted-median CPI—the $\overline{R}^2$ is 0.41 without a constant and 0.48 with a constant.

Figure 5. Scatterplots of $\pi - \pi^e$ vs. Unemployment Gap, CPI Inflation, 1985–2017

A. CPIX

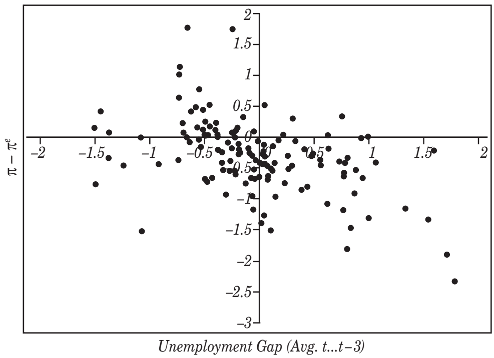

B. Median CPI

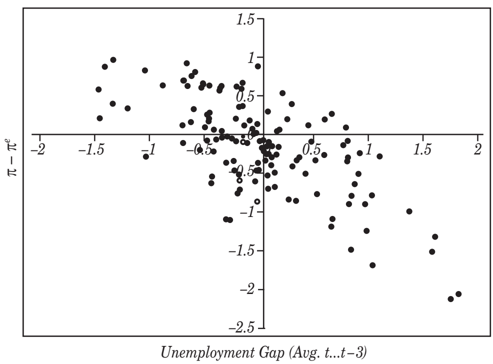

Source: Authors' calculations.

The scatterplots in figure 5 confirm that a Phillips curve appears more clearly for median CPI than for CPIX.

When we turn to PCE inflation, the differences between the results for the two core-inflation measures become larger. For PCEX, the $\overline{R^2}$ for the Phillips curve is negative without a constant and only 0.07 with a constant; for median PCE, the $\overline{R^2}$ is 0.32 in both cases. Figure 6 confirms these big differences in fit.

Some researchers damn the Phillips curve with faint praise, saying that the relationship exists but it is flat—the effect of unemployment on inflation is small—and the residuals are large. Blanchard (2016), for example, reports an unemployment coefficient of about –0.20 since the 1990s and a standard error of the residual of 1.0, indicating a "fairly poor fit." In our results for median inflation, the unemployment coefficients are substantially larger in absolute value: –0.48 for median PCE and –0.65 or –0.66 for median CPI. The standard errors of residuals are between 0.4 and 0.5.

Figure 6. Scatterplots of $\pi - \pi^e$ vs. Unemployment Gap, PCE Inflation, 1985–2017

A. PCEX

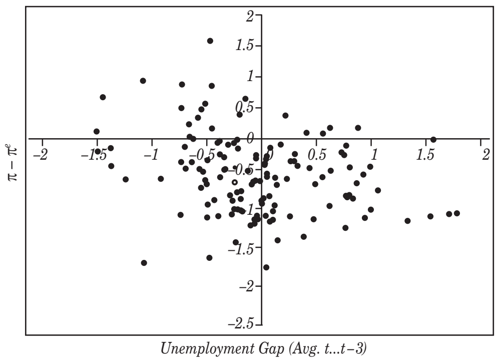

B. Median PCE

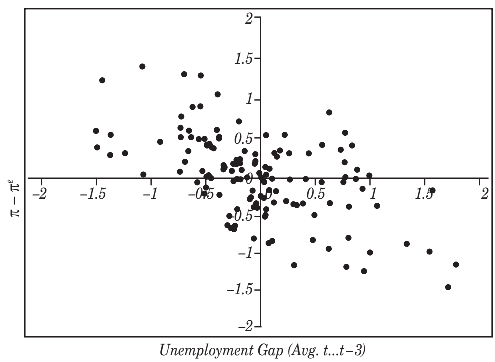

Source: Authors' calculations.

4.3 Has the Phillips Curve Taken a Vacation?

Some economists, such as Blinder (2018), suggest that the Phillips curve once existed but has disappeared since the Great Recession of 2008. Our findings on this issue depend on how core inflation is measured, even more strongly than before. When our sample period is restricted to 2008–2017, the fit of the Phillips curve becomes weaker for inflation less food and energy but stronger for median inflation.

Table 3 and figures 7 and 8 present our results for 2008–2017. Notice first that the Phillips curve always fits well in this period if core inflation is measured with the weighted median. For both median CPI and median PCE, and with and without a constant term, the $\overline{R^2}_s$ range from 0.54 to 0.64. The estimated coefficients on unemployment are close to those for the full sample since 1985. The Phillips curve appears clearly in figures 7(b) and 8(b). Based on these results, we doubt that economists would worry about the demise of the Phillips curve if they examined median inflation.

When core inflation is measured with inflation less food and energy, our results differ somewhat for the CPI and PCE deflator. For the CPIX, the evidence for a post-2008 Phillips curve is borderline. The unemployment coefficient is significant at the 5 percent level when a constant term is included in the equation but not without a constant. We can also see in figure 7(a) that the results depend heavily on two observations in the lower right of the graph: the first two quarters of 2010, which had the highest levels of unemployment in the sample and the lowest levels of CPIX inflation. If we exclude these observations, the Phillips curve slope is far from significant.

For the PCEX, the data since 2008 contain no evidence whatsoever of a Phillips curve. In the regressions, unemployment has no explanatory power for inflation ($\overline{R^2}$ = 0.001 with a constant). Figure 8(a) confirms this result, and we also see that $\pi - \pi^e$ is almost always negative: inflation has persistently fallen short of its expected level. We understand why the behavior of PCEX, the Fed's preferred measure of core inflation, has puzzled economists.

Table 3. Phillips Curves for 1985–2017

$$\pi_t - \pi_t^e = \alpha\overline{(u - u^*)}_t + \varepsilon_t$$

	Headline		CPIX		Median	
a. CPI inflation						
Constant		-0.715		-0.502		-0.183
		(0.397)		(0.100)		(0.102)
α	0.256	0.349	-0.399	-0.334	-0.699	-0.676
	(0.612)	(0.647)	(0.291)	(0.161)	(0.189)	(0.146)
$\overline{R}^2$	-0.068	-0.013	-0.487	0.178	0.543	0.601
S.E. of Reg.	2.628	2.561	0.743	0.553	0.476	0.445
b. PCE inflation						
	Headline		*PCEX*		*Median*	
Constant		-0.669		-0.544		-0.021
		(0.289)		(0.089)		(0.047)
α	0.215	0.303	-0.169	-0.098	-0.451	-0.448
	(0.433)	(0.455)	(0.241)	(0.116)	(0.079)	(0.074)
$\overline{R}^2$	-0.129	-0.006	-1.207	0.001	0.642	0.635
S.E. of Reg.	1.856	1.752	0.733	0.493	0.272	0.275

Source: Authors' calculations.
Note: OLS with robust (HAC) standard errors is used (standard errors in parentheses). The unemployment gap is the deviation of the unemployment rate from the HP filtered series, where the filter is applied over 1948–2017.

Figure 7. Scatterplots of $\pi - \pi^e$ vs. Unemployment Gap, CPI Inflation, 2008–2017

A. CPIX *B. Median CPI*

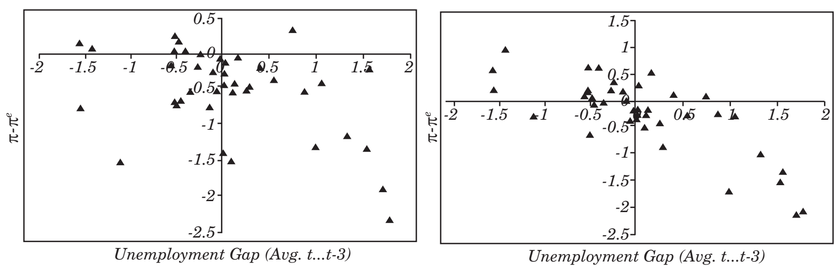

Source: Authors' calculations.

Figure 8. Scatterplots of π – π^e vs. Unemployment Gap, PCE Inflation, 2008–2017

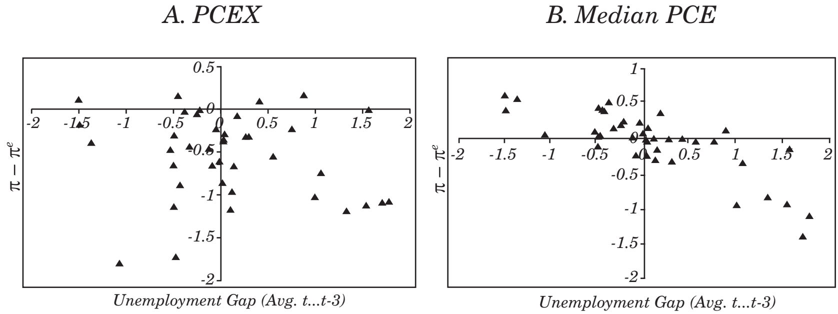

Source: Authors' calculations.

5. Conclusion

The measurement of core inflation might seem like a technical subject of interest to a narrow range of specialists. We have seen, however, that a focus on a sub-optimal core measure, the inflation rate excluding food and energy, has contributed to perplexity about inflation behavior among economists, policymakers, and op-ed writers. The weighted-median measure of core inflation has a stronger theoretical foundation than inflation less food and energy and is empirically less volatile and easier to understand. In particular, we believe that fewer economists would puzzle over a breakdown of the Phillips curve if the weighted median received more attention.

In light of these findings, economists should do more research on the weighted median and related measures of core inflation. There are many open issues. Because the median is a non-linear function of industry inflation rates, it could vary significantly depending on the level of industry disaggregation. The weighted median is also sensitive to time aggregation; for example, a quarterly series computed by averaging monthly median inflation rates differs from the median of industries' quarterly inflation rates. Researchers should ask which version of the weighted median is the most useful measure of core inflation. We should also compare weighted medians to trimmed means of industry inflation rates, which also filter out large price changes.

References

Ball, L. and N.G. Mankiw. 1995. "Relative-Price Changes as Aggregate Supply Shocks." *Quarterly Journal of Economics* 110(1): 161–93.

Ball, L. and S. Mazumder. 2011. "Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity* 42(1): 337–405.

———. 2019. "A Phillips Curve with Anchored Expectations and Short-Term Unemployment." *Journal of Money, Credit and Banking* 51(1): 111-37.

Bernstein, J. 2017. "Two Inflation Puzzles, One Inflation Monster." *The Washington Post*, 11 September.

Blanchard, O. 2016. "The Phillips Curve: Back to the '60s?" *American Economic Review* 106: 31–4.

Blinder, A. 2018. "Is the Phillips Curve Dead? And Other Questions for the Fed." *The Financial Times*, 3 May.

Brainard, L. 2017. "Navigating the Different Signals from Inflation and Unemployment." Remarks at New York Association for Business Economics, New York, NY.

Brischetto, A. and A. Richards. 2006. "The Performance of Trimmed Mean Measures of Underlying Inflation." RBA Research Discussion Papers rdp2006-10, Reserve Bank of Australia.

Bryan, M. and S. Cecchetti.1994. "Measuring Core Inflation." In *Monetary Policy,* edited by N.G. Mankiw. Chicago IL: University of Chicago Press.

Bryan, M., S. Cecchetti, and R. Wiggins. 1997. "Efficient Inflation Estimation." NBER Working Paper 6183.

Clark, T. 2001. "Comparing Measures of Core Inflation." *Federal Reserve Bank of Kansas City Economic Review* 86(2): 5–31.

Coibion, O. and Y. Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics* 7: 197–232.

Crone, T., N. Khettry, L. Mester, and J. Novak. 2013. "Core Measures of Inflation as Predictors of Total Inflation." *Journal of Money, Credit and Banking* 45: 505–19.

Dolmas, J. 2005. "Trimmed Mean PCE Inflation." Working Paper 0506, Federal Reserve Bank of Dallas.

Evans, C. 2017. "The Puzzle of Low Inflation: Implications for Monetary Policy." Remarks at The Economic Club of Grand Rapids, Grand Rapids, MI.

Friedman, M. 1968. “The Role of Monetary Policy.” *American Economic Review* 58(1): 1–17.

Fuhrer, J. and G. Olivei. 2010. “The Role of Expectations and Output in the Inflation Process: An Empirical Assessment.” *Federal Reserve Bank of Boston Public Policy Brief*.

Mutikani, L. 2018. “U.S. Annual Inflation Measures Jump as Consumer Spending Rises.” Reuters.

Rugaber, C. 2018. “U.S. Consumer Prices in March Rose 2.4 pct. from Year Earlier.” Associated Press.

Smith, J. 2004. “Weighted Median Inflation: Is This Core Inflation?” *Journal of Money, Credit and Banking* (36): 253–63.

Staiger, D., J. Stock, and M. Watson. 1997. “How Precise Are Estimates of the Natural Rate of Unemployment?” In *Reducing Inflation: Motivation and Strategy*, edited by C. Romer and D. Romer. Chicago, IL: University of Chicago Press.

Stock, J. 2011. “Discussion of Ball and Mazumder, ‘Inflation Dynamics and the Great Recession’.” *Brookings Papers on Economic Activity* 42(1): 387–402.

Stock, J. and M. Watson. 2010. “Modeling Inflation after the Crisis.” Proceedings of Economic Policy Symposium, Jackson Hole, 173–220.

Summers, L. 2017. “America Needs its Unions More than Ever.” *The Financial Times*, 3 September.

The Economist. 2017. “The Phillips Curve May Be Broken for Good.” The Economist, 1 November.

Yellen, J. 2017a. “Chair Yellen’s FOMC Press Conference June 14, 2017.” FOMC Press Conference Transcript.

———. 2017b. “Inflation, Uncertainty, and Monetary Policy.” Remarks at “Prospects for Growth: Reassessing the Fundamentals,” 59th Annual Meeting of the National Association for Business Economics, Cleveland, OH.

———. 2017c. “The U.S. Economy and Monetary Policy.” Remarks at G-30 International Banking Seminar, Washington, DC.

The Link between Labor Cost Inflation and Price Inflation in the Euro Area

Elena Bobeica
European Central Bank

Matteo Ciccarelli
European Central Bank

Isabel Vansteenkiste
European Central Bank

To gauge inflationary pressures, policymakers generally pay close attention to labor cost developments. A key reason has been the widely held view that labor cost inflation (i.e., wage inflation adjusted for productivity developments) is one of the main causes of price inflation. From a theoretical perspective, this assumption represents the post-Keynesian cost-push/price-markup view of the inflationary process whereby wage increases in excess of productivity are seen as putting upward pressure on prices, and wages are the exogenous variable determining the future direction of inflation.[1]

We would like to thank Gonzalo Castex, Jordi Galí, Giorgio Primiceri, Juan Rubio-Ramirez, Frank Smets, Thomas Westermann, the seminar participants at the University of Maastricht, at Banque de France and at Bundesbank, and the participants of the XXII annual conference of the Central Bank of Chile: 'Changing Inflation Dynamics, Evolving Monetary Policy' held in Santiago, Chile, on 25–26 October 2018, for their useful comments and suggestions on a preliminary version of this paper. The views expressed in this paper are of the authors only and do not necessarily reflect those of the European Central Bank or the Eurosystem.

1. In the paper we will refer to labor cost as compensation per employee developments adjusted for productivity, whereas wages will refer to compensation per employee. In some studies, what we consider as the labor cost is also referred to as unit labor cost (ULC).

Changing Inflation Dynamics, Evolving Monetary Policy edited by Gonzalo Castex, Jordi Galí, and Diego Saravia, Santiago, Chile.

Such a cost-push view of inflation was often invoked in the 1970s to explain inflationary dynamics[2] and to this date often remains the underlying assumption in policy communication on the outlook for inflation. For instance, in the years leading up to the 2008–2009 global financial crisis, labor cost dynamics were closely monitored to sniff out signals of a possible buildup of excessive inflation,[3] in part due to concerns of a return of the 1970s-type wage spiral. In the aftermath of the Great Recession, with concerns having shifted from perceiving inflationary and labor cost pressures from being too high to too low, forecast narratives see a pickup in labor cost growth often as a necessary condition for rising inflation.[4,5]

While these labor-cost-based explanations of inflation dynamics continue to take a prominent place in the policy debate, the academic literature has expressed more skeptical views. Empirical studies, which generally focused on U.S. data, have drawn mixed conclusions on the link between labor cost and price inflation, in particular at shorter horizons. First, it remains unclear whether labor costs tend to precede or follow prices.[6] And second, studies suggest that the relationship between labor cost inflation and price inflation may have weakened over time, potentially due to an improved anchoring of inflation expectations.[7]

However, when looking at the theoretical literature, it is rather unsurprising that empirical studies have not been able to draw any firm conclusion on the link between labor cost inflation and price inflation at shorter horizons. Theoretical models generally do not put into question that, in the long run, labor cost inflation and price inflation are closely interrelated and that we should eventually expect

2. In the 1970s the so-called wage-price spiral was seen as causing inflationary dynamics whereby higher labor-cost growth resulted in higher price inflation which in turn led workers to push for higher wage growth and, subsequently, even faster price increases.

3. See ECB (2004).

4. See ECB (2018).

5. Similar references on the link between labor cost and price inflation developments were made in a Bank of England speech by the external Monetary Policy Committee (MPC) member Saunders on 20 April 2018, who noted that the Committee forecasts a gradual pickup in domestic cost growth that would help keep inflation slight above target two and three years ahead even as currency effects fade. For the Bank of Japan, the Deputy Governor Iwata noted in a recent speech delivered on 31 January 2018 that the inflation rate is projected to rise in line with wage increases.

6. See Knotek and Zaman (2014) and Bidder (2015).

7. See Peneva and Rudd (2017).

wage inflation, adjusted for productivity, to move together with price inflation.[8] However, in the short to medium run, it is not at all obvious that rising labor costs should translate into price inflation.

In the industrial organization literature, an alternative to the cost-push view is that firms will charge whatever the market will bear, regardless of their actual costs. If the markets' acceptance of higher prices is the dominant determinant of inflation, the cost-push model would have less validity.[9] Also, the cost-push view abstracts from any influences that monetary policy may have on the inflation process. For instance, if a central bank is pursuing a contractionary policy trying to keep inflation low, firms might not be able to pass on higher labor costs into prices. In fact, the causality between prices and labor costs might go the other way: in the case of excess demand, firms would be able to increase prices, which would lead to higher demand for wages. Reflecting these differentiations, in New Keynesian models, the correlation and lead-lag relationship between labor cost inflation and price inflation can be expected to depend not only on the degree of the prevailing price and wage rigidities in the economy, but also on the type of shock that hits the economy. As a result, we should in fact expect the link between labor cost inflation and price inflation to vary across time, across countries and also across sectors.

Having a better understanding of the signal that labor cost developments provide for the inflationary process is of key relevance from a policy perspective. In the euro area for instance, it is well-known that the reaction of inflation dynamics to accelerating growth has been atypically slow in the aftermath of the Great Recession.[10] While there are a number of plausible explanations for this, it nevertheless sheds some uncertainty on the inflation outlook. Having a deeper understanding of the drivers of the inflationary process can help reduce this uncertainty. However, to date there exists no study that has

8. In the long run, the real wage is determined by factors such as productivity, bargaining power, and the ability of firms to mark up prices over costs. Consequently, prices and nominal wages must adjust relative to each other to be consistent with these fundamentals. In this case, long-run growth in the real wage can only come from productivity growth. Because of this, if nominal wages grow faster than productivity, they must, in the long run, be associated with price inflation. Otherwise workers would ultimately claim all proceeds of production and business owners would be left with nothing. If wage inflation substantially exceeds productivity growth, then inflation must also be high to be consistent with real wages rising in line with long-run productivity improvements.

9. See Banerji (2005).

10. See Draghi (2018).

systematically documented and analyzed the empirical link between labor cost inflation and price inflation in the euro area.

In this paper, we aim to contribute to the literature by documenting and analyzing the link between labor cost inflation and price inflation for the largest four euro-area countries, by using quarterly data over the period 1985Q1–2018Q1 at the country-wide level and for the three largest sectors in each economy (manufacturing, construction, and services). We argue that the link between labor cost inflation and price inflation is not only state but also shock-dependent. The idea that the relationship between variables is shock-dependent is not new. It has already been more extensively explored in the exchange rate literature,[11] but also for understanding the Phillips curve relationship.[12] However, its relevance for the link between labor cost and price inflation has also already been suggested. Gumiel and Hahn (2018) present evidence based on one of the European Central Bank (ECB) core models for policy simulations that the response of the GDP deflator to wage increases/decreases is different for supply shocks (in this case wage-markup shocks) than for demand shocks.

Our paper builds on these findings. Concretely, we analyze empirically the link between labor cost and price inflation in the euro area over the short- to medium-term horizon, and check if the extent to which the link has changed over time depends on the level of inflation and the type of shocks that hit the economy. We focus on the developments in the total economy and three main sectors of the four largest euro-area economies.[13] As the link between labor cost and price inflation has been less documented for the euro-area countries, we start by presenting some stylized facts and by conducting preliminary analyses that have become commonplace in the U.S. literature on this topic. More specifically, we (i) look at the cross-correlation between labor cost and price inflation, (ii) test Granger causality, and (iii) conduct both a conditional and unconditional forecast evaluation. Subsequently, we consider the link between labor cost and price inflation in a dynamic and conditional setup by estimating a three-variable vector autoregression (VAR) model for each sector of each

11. See Forbes and others (2018), Comunale and Kunovac (2017).

12. See Galí and Gambetti (2018).

13. Note that in this paper we focus on the short- to medium-term horizon, as this is the most relevant horizon from a policy perspective. Moreover, this is also the horizon at which a clear consensus and view is still missing on the link between wage and price inflation.

country. This allows us to answer questions, such as: (i) whether the conditional correlations are different from the unconditional ones; (ii) by how much price inflation rises when labor costs increase, and (iii) the extent to which this "passthrough" has evolved over time or depends, for example, on the level of price inflation. In the final part of the paper we move to a more structural setup and analyze whether and how the link between labor cost and price inflation depends on the type of shocks that hit the economy.

Overall our results show that in the four biggest euro-area countries, contrary to the U.S., there is a clear link between labor cost and price inflation. This result is confirmed across a battery of approaches and tests. The link has also remained overall rather stable over time, albeit with some differences across sectors and countries. However, at the same time, and in line with the findings in the U.S. literature, the link appears to depend on the level of price inflation: when inflation is high, the link becomes stronger. Finally, the link is shock-dependent: when the economy is hit by a demand shock, there is a clear and relatively strong link between labor cost and price inflation. This is not the case for supply shocks, where the link is less conclusive. These findings have important policy implications. In particular, the results suggest that monitoring and analyzing labor cost developments in the euro area is indeed relevant to understanding the evolution of price inflation. However the importance of these developments does depend on the level of price inflation and on the shocks that prevail in the economy. In an environment of expansionary demand, the information contained in labor cost developments is much more relevant for price inflation than when the economy is hit by a supply-type shock. In other words, under circumstances of predominantly demand shocks, one can be confident that labor cost increases will be passed on to prices. However, after a period of low inflation (such as the one between 2012 and 2018), this passthrough could be moderate at least until inflation stably reaches a sustained path.

The remainder of the paper is organized as follows. Section 1 connects the paper to the existing literature. Section 2 discusses some preliminary analysis of the data and presents unconditional stylized facts on the link between labor cost and price inflation. Section 3 analyses the link in a VAR setup and considers to what extent this link has changed over time or depends on the level of price inflation. Section 4 presents results based on a structural VAR model. Section 5 summarizes and concludes.

1. Link to the Existing Literature

Labor markets have been a focus of interest in the study of price inflation ever since Phillips uncovered the negative relationship between the rate of change in wages and the unemployment rate, i.e., the so-called Phillips curve.[14] Since then an extensive literature has developed that studies the interrelationship between labor-market developments and price inflation. An important share of this research has explored how informative labor cost inflation is for price inflation, in particular in the short to medium run.[15]

Studies have taken a number of avenues to analyze this question. A first important strand in the literature has focused on the causal relationship between wage inflation and price inflation. Theoretically, the post-Keynesian view would suggest that the excess of wage gains over productivity gains leads price inflation. Instead, according to the neoclassical theory, the causality between wages and inflation would run in the opposite direction. In this case, the real wage is considered the relevant wage variable in the wage-employment relationship and nominal wages are expected to respond to price changes so as to preserve the real wage for a given productivity level. Empirically, analyses based on in-sample Granger-causality-type of tests have yielded mixed conclusions. A number of studies tend to favor the idea that price inflation causes wage inflation and that the causality can differ across sectors. Hu and Toussaint-Comeau (2010) find that wage growth does not cause price inflation in the Granger causality sense, especially after the mid-1980s. By contrast, price inflation does Granger-cause wage growth. Similarly, Emery and Chang (1996), and Sbordone (2002) find some evidence that rising prices precede the growth in unit labor costs.[16] However, some other studies find actually no causal link between price and wage inflation. For instance, Hess and Schweitzer (2000) find that price and wage changes are best predicted by their own lags, meaning that none Granger-causes the other. Along similar lines, Gordon (1988) and Darrat (1994) conclude that wages and prices are irrelevant to each other and that they *"live*

14. Fisher (1926) had already uncovered the link between price inflation and the unemployment rate earlier, however he saw price inflation as driving the rate of unemployment.

15. In the long run, the relationship between labor-cost inflation (i.e., wage inflation adjusted for productivity) and price inflation is rather uncontroversial, both from a theoretical and empirical point of view.

16. See Bidder (2015).

a life of their own". Finally, Banerji (2005) approaches this changing relationship from a different angle, by looking at cyclical turns. He finds that labor cost inflation leads price inflation at peaks, but lags it at troughs, which would make changes in labor cost a lagging indicator of upturns in price inflation. Finally, Rissman (1995) finds that only in manufacturing and trade services, wages Granger-cause inflation.

A second strand of the literature has investigated whether wages add any information when trying to forecast inflation.[17] Overall, these studies have found that for out-of-sample forecasts, wages do not provide significant additional information beyond what can already be gleaned from other sources, including prices themselves.[18] At the extreme, Stock and Watson (2008) even show that models using common wage measures may perform worse than their preferred benchmark without wages.

A final strand of the literature has examined whether the link between labor cost inflation and price inflation is time-varying. Studies here tend to find that, while in the past (i.e., prior to the mid-1980s) labor cost inflation did provide signals for price inflation, there is little evidence that in recent years movements in average labor cost growth have been an important independent influence on price inflation. Concretely, Knotek and Zaman (2014) show how the correlation between wages and prices has decreased since the mid-1980s. Similarly, Peneva and Rudd (2017) show how the passthrough of labor cost growth to price inflation in the U.S. has declined over the past several decades, to the point where it is currently close to zero. One explanation put forward has been the better anchoring of inflation expectations in recent years. Another one is that low levels of inflation change the wage-price nexus because of downward wage rigidities.[19] Such a view was also empirically uncovered by Mehra (2000), who finds that in periods of low inflation wages do not help to predict inflation, while it does in a high inflationary environment. Zanetti (2007) found similar results when using Swiss data.

From these studies it thus appears generally difficult to ascertain that over shorter horizons wages have an *independent* influence on prices and that the link has weakened over recent years. However, most of them are based on U.S. data. Instead, for the euro area, only few studies have examined this link. IMF (2018) replicates the Peneva and

17. See Stock and Watson (2008), Knotek and Zaman (2014).
18. See Bidder (2015).
19. See Daly and Hobijn (2014).

Rudd (2017) approach for the EU15 panel and they find that there is a statistically significant passthrough from labor cost growth to inflation for these countries. Dees and Guntner (2014) explore the cost-push factors to inflation dynamics from the supply side across four sectors (industry, construction, services, and agriculture) in the four largest euro-area countries over the period 1995–2012. In their analysis the authors find that disaggregated information improves the inflation forecasting performance and that their model, which also accounts for unit labor cost developments, fares comparatively well against common alternatives. Forecast errors however do tend to be larger during the financial crisis period. Jarocinski and Mackowiak (2017) in turn consider whether unit labor cost, among a large set of potential indicators, add information when trying to forecast inflation. They conduct their exercise for both the U.S. and the euro area. The authors find that the unit labor cost ranks low in the U.S. (28th among 38 variables), while ranking somewhat better for the euro area (18th among 38 variables). Using a different approach, Tatierska (2010) finds, by estimating a New Keynesian Phillips curve, that in eight out of eleven euro-area countries there is a plausible relationship between inflation and labor cost growth. Finally, at the micro level, Druant and others (2009) find that wage and price changes feed into each other. Around 40 percent of the firms surveyed acknowledge that there exists a relationship between wages and prices. However, only 15 percent state that this relationship is relatively strong. For half of them, decisions on price changes follow those on wage changes. The opposite holds for another three percent, while decisions are simultaneous in the remaining four percent.

2. A First Exploratory Look at Labor Costs and Inflation in the Euro Area

In our analysis, we concentrate on the link between labor cost and price inflation in the four largest euro-area countries (Germany, France, Italy, and Spain) for the economy as a whole and for the three main economic sectors: services, manufacturing and construction.[20]

For this purpose we collected quarterly data over the period 1981Q1–2018Q1. Details on the data sources and the data series included are provided in appendix A. To measure labor costs, we use nominal compensation per employee adjusted for productivity—in

20. The three economic sectors combined represent between 70% (in Germany) and 80% (in France) of total value added. We did not include the agricultural sector which represents only between 0.7% (in Italy) and 2.9% (in Spain) of total value added.

line, for instance, with Peneva and Rudd (2017)—rather than nominal compensation per employee, as the former is a better proxy of the true cost pressure faced by the firm.[21] For inflation, we use the implicit sectoral gross value added deflator.[22]

Figure 1. Cost Structure of Production of Manufacturing and Services Firms in the Euro Area

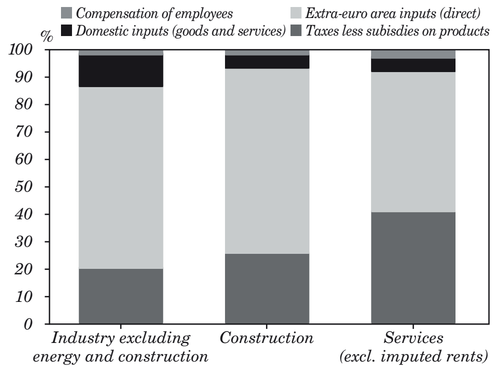

Source: Eurostat, authors' calculations.
Latest observation: Input/output tables 2015.

21. Our wage measure is compensation per employee. Alternative measures of wages across euro-area countries exist, such as compensation per hour or hourly labor cost. The latter encompasses employee compensation (which includes wages, salaries in cash and in kind, employers social security contributions), vocational training costs, and other expenditure (such as recruitment costs, expenditure on work clothes, and employment taxes regarded as labor costs minus any subsidies received). However, these alternatives are generally consistently available across sectors and countries on a quarterly basis only since 1995 and in some cases (in particular Spain) only later. For this reason, our preferred wage proxy is compensation per employee. Moreover, we find that the correlation between our wage measure (i.e., compensation per employee) and the other measures is rather strong in their overlapping sample periods. For compensation per hour, the correlation is on average above 0.8. The only outlier is the Italian manufacturing sector, where the correlation is 0.5. When comparing our measure with Eurostat's labor-cost index during overlapping periods, the average correlation is around 0.6.

22. Note that CPI inflation is not available at sectoral level. The gross value added deflator at sectoral level has been obtained by dividing nominal value added by real value added at sectoral level. The key difference between the implicit gross value added deflator and the consumer price index is that the latter measures price developments from the perspective of the consumer, whereas the former considers price developments from the perspective of domestic production of goods and services. In practice this implies that import prices matter for the consumer price index, but not for the gross value added deflator (where export prices do matter). Appendix B plots the evolution of the GDP deflator, labor cost and CPI inflation for the total economy for the four countries of the analysis. The chart shows that the correlation between the annual growth rate of GDP deflator and of the consumer price index is very high.

We conduct our investigation for each country separately, given the substantial heterogeneity in labor market institutions and in the wage formation process. Moreover, we believe that it is important not just to conduct the analysis at the country level but also to exploit the sectoral dimension. Sectors differ in terms of labor market tightness and many other labor market characteristics that affect the passthrough of labor cost to price inflation. The cost structure of production firms is different, with services having a bigger share of labor costs (see figure 1). At the same time, manufacturing is subject to international competition to a larger degree. Furthermore, other characteristics, such as workers' turnover rates, capital intensity, or the degree of wage bargaining institutions, are also sector-dependent. Finally, sectors differ in terms of their degree of wage rigidity. For instance Du Caju and others (2009) show (by using a Belgian firm-level dataset) that wages in construction are particularly sticky, less so in services and even less so in manufacturing. Tatierska (2010) also argues that the sensitivity of price to labor cost inflation differs across sectors, reflecting the different degree of price stickiness; the services sector exhibits stickier prices, so she finds that for most analyzed countries (out of 11 euro-area countries), labor cost inflation Granger-causes price inflation for the total economy in more instances than for services.

2.1 Data

Figure 2 plots the year-on-year growth rate of the labor cost and our measure of price inflation, for the total economy for each of the four countries. The high correlation (ranging between 0.85 and 0.91) between the two series demonstrates why analysts have paid close attention to labor costs when assessing price inflation.[23] However, what is not clear from the figure is whether movements in labor costs precede movements in price inflation, or *vice versa*.

At the same time, figure 2 does clearly demonstrate that in part the high co-movement between the two data series can be explained by a strong common (downward) trend over an important part of the sample (in particular in the 1980s and early 1990s) which can be attributed to the convergence process in the run-up to the European Economic and Monetary Union (EMU) and the improvements in the anchoring of

23. These high correlations are generally also confirmed at the sectoral level. The correlation is however somewhat lower for the construction sector, where it ranges between 0.31 (for Spain) and 0.64 (for Italy).

inflation expectations towards lower levels. These common patterns are visible across all countries and sectors (not reported). Therefore, before choosing the appropriate level of aggregation at which to remove the trend, we compute a single common factor across all price and labor cost inflation series as well as within-country factors (common to labor cost and price inflation series of all sectors belonging to the same country), and check the variance explained by these factors. It turns out that the variance of the two variables of interest explained by country factors is not only higher on average (60 vs. 50 percent) but also consistently higher across countries than the variance explained by a single common factor. The latter would explain a high variance of the two variables in France, Italy, and Spain (and not in all sectors) but not in Germany.

Figure 2. Unit Labor Cost and GDP Deflator
(year-on-year percentage change)

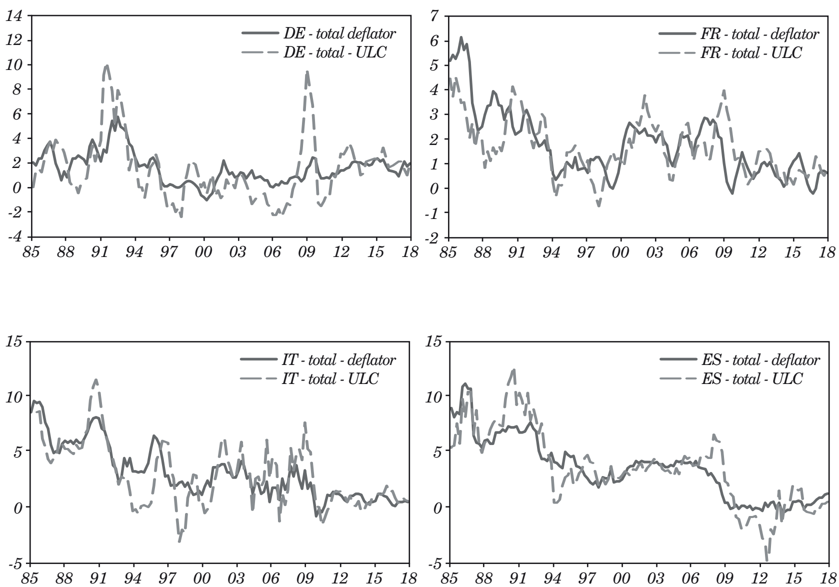

Sources: Various sources, authors' calculations.

Based on this evidence and on the fact that this common movement is related to the improvements in the anchoring of inflation expectations to lower levels, we decided to adjust the series for the common movements at the country rather than sectoral level. To do so, we follow Knotek and Zaman (2014) which are in turn inspired by the forecasting literature that has found gains in inflation forecasting accuracy by specifying inflation in gap form as the deviation from a slow-moving long-run trend.[24] Concretely, we construct labor cost and price inflation gaps as the year-on-year growth rates in these variables minus the consensus survey-based long-run inflation expectations. As inflation expectations for the countries in our sample are only available since 1989 (and for Spain even only since 1995), we rely on an unobserved component model to create labor cost and price inflation gaps in the period prior to that.[25] The adjusted series are shown in figure 3. This adjustment also implies that the series are stationary, according to a standard augmented Dickey-Fuller (ADF) unit root test.[26]

The common trend is crucial in understanding the link between labor cost and price inflation. As shown in figure 4, the correlation between price and labor cost inflation appears to have changed after the crisis when looking at unadjusted data, but there is no striking difference when considering the adjusted series. For the remainder of the paper, we will base our analysis on the adjusted series of labor cost and price inflation.

24. See Kozicki and Tinsley (2001), and Zaman (2013).

25. The unobserved component model is estimated on the price inflation series and assumes that the inflation trend follows a random walk. This trend estimate from the unobserved component model is then applied to both the labor cost and price inflation series.

26. To ensure that our results do not depend on the approach taken, we also construct alternative price inflation and labor-cost inflation gaps as year-on-year growth in these variables minus a series-specific or shared long-run trend. Specifically we use a Hodrick-Prescott filter to adjust the series for the time span where inflation expectations were not available. The results in the paper were unchanged when applying this approach.

Figure 3. Adjusted Labor Cost and GDP Deflator
(year-on-year percentage change)

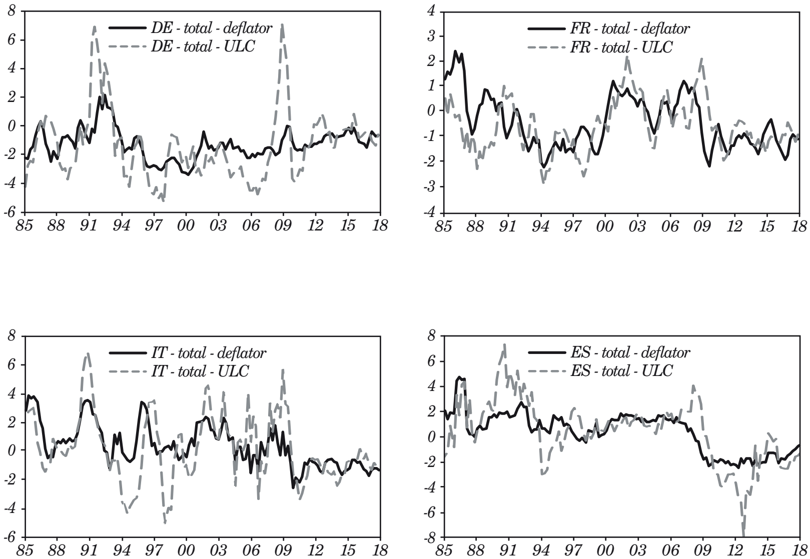

Sources: Various sources, authors' calculations.
Latest observation: 2018Q1.

2.2 Cross-correlations

In this subsection, we analyze the unconditional connection between labor cost inflation and price inflation by looking at cross-correlations, which allow for a simple examination of the lead-lag structure of the correlation and the strength of the connection between the series.

If labor cost inflation reliably comes ahead of price inflation in the data, then the strongest cross-correlation should be between labor cost inflation in quarter t and price inflation in some k-th quarter after t.

The unconditional cross-correlations (figure 5) of the adjusted series continue to show a high contemporaneous correlation (albeit lower than on the non-adjusted series) ranging between 0.4 (France) and close to 0.8 (Spain).

Figure 4. Adjusted and Unadjusted Labor Cost Growth (6 Months Prior) and Price Inflation in the Euro Area

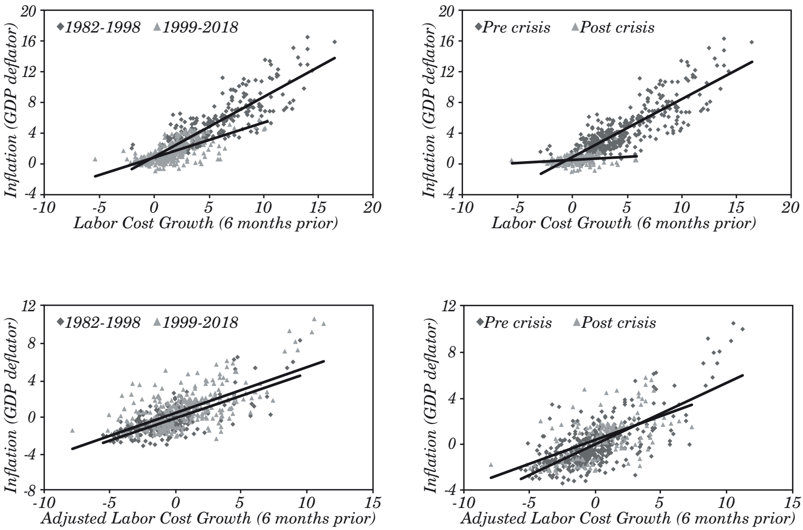

Sources: Various sources, authors' calculations.
Latest observation: 2018Q1.

At the same time, we do not observe a systematic lead/lag pattern across countries or sectors. While in Italy the highest correlations occur mostly contemporaneously, in the German total economy and service sector, labor costs seem to lead prices. In France, except for the service sector, prices lead labor costs. Similarly, in the Spanish service sector and the total economy, prices lead labor costs, while labor costs are clearly leading prices in the construction sector.

When examining the same cross-correlations on a rolling sample, we notice only small changes over time, though in the post-crisis period the correlations have tended to become more contemporaneous (except in the Spanish construction and the French service sector).[27]

27. See figure C1 in appendix C.

Figure 5. Cross-Correlation between Adjusted Labor Cost and Price Inflation

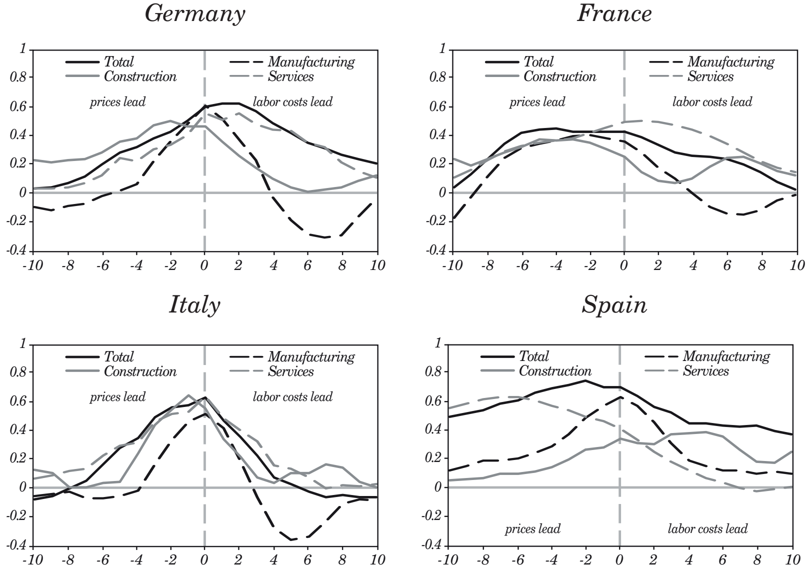

Sources: Various sources, authors' calculations.
Note: The charts show the cross-correlation between price inflation gaps at time t and labor cost inflation gaps at time t-k. Sample period: 1985Q1–2018Q1.

2.3 Granger Causality and Forecast Evaluation

Another angle to look at the link between labor cost and price inflation is to ask whether past changes in labor costs contain useful information for predicting future changes in prices. We consider here two commonly adopted approaches to analyze this question from an in- and out-of-sample perspective, namely Granger causality and a pseudo out-of-sample forecast evaluation.

As regards the Granger causality test, we adopt here the classical approach whereby, in a single equation model, price inflation is regressed on p lags of price and labor cost inflation and the exclusion of the labor cost inflation lags is tested. The test is performed on a recursive basis, starting by estimating the equation over the period 1985Q1–1998Q4 and subsequently adding one quarter at a time. Lags are optimally chosen with a grid search to minimize the p-values of the Granger causality test. In other words, we look for the best specification which is the most likely to result in labor cost inflation being Granger causally prior to price inflation.

Results (figure 6) show that, contrary to U.S. data (section 1), we can find Granger causality from labor cost to price inflation at ten and five percent significance. Moreover, and thus confirming the conclusions from the unconditional cross-correlations, we see that the labor cost and price inflation link has not weakened in the recent period (the notable exceptions are the Italian construction and Spanish service sectors). In fact, in most cases the causality from labor cost to price inflation has strengthened over time. France is the only country where this causality has been less evident throughout the sample, except the construction sector and the service sector until the financial crisis.

Figure 6. Recursive Granger Causality Test Results (*p*-values)

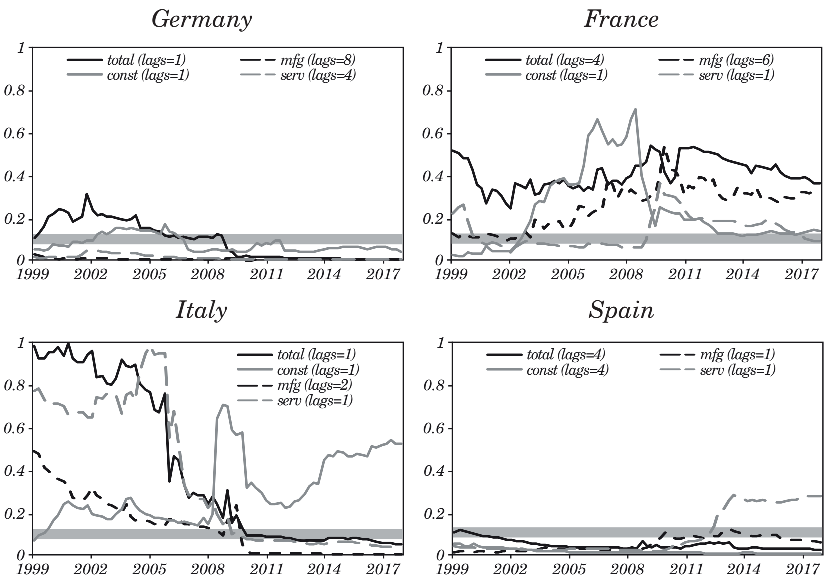

Source: Authors' calculations.
Note: mfg stands for manufacturing, const for construction, and serv for services. The lags for the Granger causality test were optimally chosen. The horizontal dark blue line represents the threshold for the significance of the test at a ten percent level.

In the second approach, we focus on the out-of-sample forecasting power of labor cost inflation for price inflation. For this purpose we estimate a simple trivariate VAR model for each sector which includes: real value added growth, labor cost inflation and price inflation. We subsequently perform two exercises: an unconditional and a conditional forecast. In the first case, we compare the unconditional forecast of price inflation from the trivariate VAR with a bivariate VAR (i.e., a model which only includes activity and prices). Our benchmark evaluation period is 1999–2018 but we also checked the results for the periods 1999–2007 and 2008–2018. Besides the unconditional forecast, we also consider a conditional forecast exercise. In this case, we compare the inflation forecast from the trivariate VAR conditional on the true path of labor costs with the forecast of price inflation from the same model where we condition on a constant path for labor costs (i.e., we assume a random walk).[28]

The results from the unconditional and conditional forecasts are shown in the tables in appendix D. Overall, while the unconditional forecast presents mixed results and would seem to suggest that labor costs can, in our exercise, add some useful but limited information to the price inflation forecasts across samples, the conditional forecasts strikingly show that labor cost inflation has indeed some forecasting power for price inflation in this setup. This result appears consistently across sectors and countries with the exceptions of the construction and service sectors in Spain. When evaluating the forecast before and after (the beginning of the) global financial crisis, we observe a tendency to improve the forecasting over the latter part of the sample in case of the total economy for all countries except Italy (where we do not see a change). When checking the opposite direction (from prices to labor costs), overall we observe many more ratios bigger than one and a better forecasting performance over the last part of the sample for Germany and Spain.

2.4 Summary

This section can easily be summarized: labor cost and price inflation show a consistent and strong (unconditional) link across

28. Concretely, the strategy is the following: (i) we run an initial estimate of the model until 1998Q4; (ii) we do a rolling estimate thereafter and project inflation (for each sector) eight steps ahead conditional on the true path of labor-cost inflation and conditional on a constant labor-cost inflation; and (iii) we evaluate the ratios of root-mean-square error (RMSE) obtained in both cases.

euro-area countries and sectors at a cyclical frequency, i.e., even after removing a common trend. In fact, without removing a common trend, the correlation between labor cost and price inflation would have spuriously changed after the real and financial crisis, as found for the U.S. data by Peneva and Rudd (2017). The direction of causality is not obvious to ascertain but, contrary to the evidence typically based on U.S. data, it is possible to find some in- and out-of-sample forecasting power of labor costs for price inflation. No obvious country- or sector-specific pattern emerges from this preliminary analysis.

3. A Simple VAR Analysis

3.1 Empirical Approach

To examine in a dynamic and more conditional manner the relationship between labor cost and price inflation, we use VAR models for each sector of each country, in total 16 VARs. We do not exclude the possibility of cross-countries/sectors interrelationships, which could be accounted for in a panel VAR approach as in Canova and Ciccarelli (2009), but the sparse number of dynamic interrelationships among countries and sectors can make a fully-fledged panel VAR setup inefficient for our aim. Moreover, the heterogeneity in the data makes the approach used here preferable to approaches that restrict the dynamics of the endogenous variables to homogeneity in a pooling panel. Estimating sector by sector allows us to look at average results, if needed, by simply using consistent mean group estimators on the disaggregated results.

Our baseline VAR system contains three variables: the growth rates of (i) real value added, (ii) unit labor cost, and (iii) the value added deflator. The latter two variables have been adjusted as explained in section 2 to remove a common trend. The baseline estimation period ranges from 1985Q1 to 2018Q1. The VARs are estimated with four lags and Bayesian techniques assuming a normal-diffuse prior with a Minnesota prior on the matrix of coefficients to deal with the curse of dimensionality.[29] We also conduct a robustness check of our results by adding the spread between a long- and a short-term interest rate to the VAR system. The included variable is intended to proxy for monetary policy. Our findings are largely unaffected by this extension, as shown in the figures in appendices H and K.

29. See Kadyiala and Karlsson (1998).

In this simple setup we use the estimates of the 16 VARs to evaluate impulse response functions of inflation to a shock in unit labor cost inflation by means of a Choleski orthogonalization with the variables ordered as listed above. The dynamic responses are used to answer the question: how much does inflation rise when labor costs rise by one-standard deviation? Standardized multipliers are computed mimicking the fiscal literature[30] as the ratio of the cumulative responses of price and labor cost inflation over the horizons 1 (impact) through 28 quarters. With such standardization, the multipliers are comparable across countries and sectors.

3.2 Main Findings: Baseline VAR Specification

We first report the estimated contemporaneous correlations between labor cost and price inflation computed from the moving-average representation of the VARs (i.e., the impulse response estimates) truncated to 40 lags.

Table 1. VAR-Based Correlation between Labor Cost and Price Inflation

	Conditional on	*Total*	*Manufacturing*	*Construction*	*Service*
DE	All shocks	0.62	0.62	0.50	0.57
	Shock to y	0.78	0.91	0.84	0.79
	Shock to ulc	0.88	0.77	0.39	0.89
	Shock to p	0.33	0.06	0.56	−0.18
FR	All shocks	0.40	0.35	0.27	0.48
	Shock to y	0.49	0.39	0.02	0.52
	Shock to ulc	0.82	0.83	0.83	0.70
	Shock to p	−0.04	0.28	0.35	0.29
IT	All shocks	0.63	0.52	0.55	0.63
	Shock to y	0.74	0.88	0.61	0.68
	Shock to ulc	0.90	0.27	0.74	0.85
	Shock to p	0.34	0.03	0.58	0.45
ES	All shocks	0.75	0.65	0.37	0.41
	Shock to y	0.85	0.92	0.53	0.77
	Shock to ulc	0.96	0.90	0.50	0.42
	Shock to p	0.63	0.65	0.31	0.54

Source: Authors' calculations.
Notes: Table 1 reports estimates of conditional correlations between labor cost and price inflation. Significance (values in bold) is based on 68% Bayesian credible intervals.

30. See Mountford and Uhlig (2009).

Table 1 reports the correlation estimates between the two variables of interest conditional on all shocks (which is equivalent to the unconditional correlation discussed above in section 2) and conditional on shocks to real value added growth, labor cost inflation and price inflation. In most cases, the estimates point to relatively large, positive, and significant correlations, thus confirming the previous results that, over the sample of analysis, the link between labor cost and price inflation across euro-area countries and sectors is quite strong, also after controlling for the own dynamics and for the dynamic relationships with a real activity indicator. The only exception is the correlation conditional on shocks to price inflation which in several occasions is insignificant or negative, and in any event almost consistently lower than the correlations conditional on other shocks. The same correlation conditional on shocks to labor cost inflation is instead always positive and significant and can be as high as 0.96 (Spain, total economy).

An interesting result based on the same estimates is given by the forecast error variance decomposition (FEVD),[31] which indicates that almost systematically (with the exception of Italian construction) the variance of inflation explained by shocks to labor cost inflation is bigger than the variance of labor cost inflation explained by price inflation. These percentages are not very high on average, but can reach values as high as 70 percent (France).

In order to better understand these results, figure 7 plots the impulse response functions of price inflation to a shock to labor cost inflation, standardized as explained above in subsection 3.1. The estimates can be interpreted as passthrough multipliers from labor cost to price inflation. The full set of results can be found in figure F1 in appendix F, where we also report the recursive estimates of the steady-state passthrough distributions (median and 68 percent credible interval) for all sectors and countries.

These charts show that the steady-state passthrough values are almost always significantly different from zero. Moreover, they confirm the finding from the unconditional cross-correlations (appendix C) that there is no apparent structural break or significant change in the link between labor cost and price inflation over time and that there are important heterogeneities across countries and sectors.

31. See figure F2 in appendix F.

Figure 7. Choleski Decomposition-Based Passthrough from Labor Cost to Price Inflation

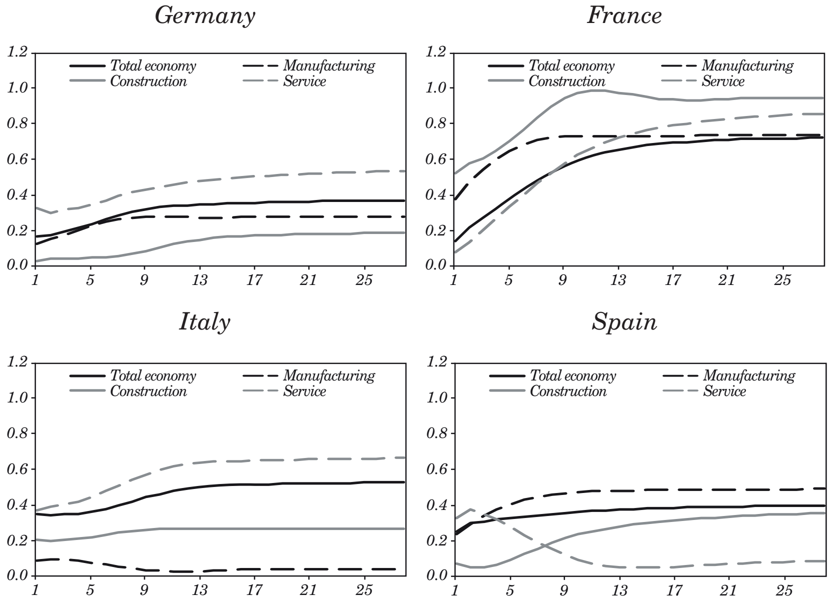

Source: Authors' calculations.

3.2.1 How Does the Passthrough Differ Across Countries?

Another aspect worth considering is how and why our passthrough results differ across countries. In this regard, figure 7 shows that France exhibits the highest passthrough values across all sectors. A cross-check of the conditional and the unconditional cross-correlations would confirm that the construction and manufacturing sectors in France drive up the passthrough across the economy. A strong passthrough from wage growth adjusted for productivity to price inflation was also found in Quevat and Vignolles (2018), based on a model for core inflation where also changes in VAT are accounted for. Charsonville and others (2017) confirm the pattern that we find across sectors in France, namely an initial higher passthrough in manufacturing and a subsequently more important one for services.

One reason for such a relatively high passthrough for France could be the presence of stronger second-round effects.[32,33]

The passthrough in Germany is lower and clearly driven by services. Nevertheless, a 0.4 passthrough suggests that labor costs are being passed through to prices in a noticeable manner. The Bundesbank also acknowledges the importance of wage developments for consumer prices and confirms that the passthrough from wages to prices is below 50 percent.[34] Why would the passthrough be lower in Germany than in France? Following the line of thought of Kuegler and others (2018), the wage setting process in the two countries differs substantially. Germany has witnessed an unprecedented decentralization of the wage formation process since the mid-1990s and a fall in union coverage rates; trade unions were responsible for a prolonged period of wage restraint. In France there was no similar decentralization of the wage setting process and labor unions play a more prominent role. In a situation of similar productivity growth[35] and an increased convergence in inflation rates across countries, the wage moderation process which occurred in Germany would imply, mechanically, a lower passthrough to inflation.

Also in Italy the passthrough of labor costs to prices is driven by services, thus confirming the results based on unconditional contemporaneous correlations. The relatively strong passthrough of labor costs to Italian prices is supported by findings based on firm-level data, whereby firms' inflation expectations are significantly affected by wage changes, particularly in high inflation regimes.[36]

Spain stands out with a low steady-state passthrough in the services sector. This is unsurprising in light of previous findings, such as the fact that in this sector it is price inflation which appears to lead labor cost inflation, as reported in figure 5.

In order to put these findings in perspective, we cross-checked our findings against two main results of the euro-area Wage Dynamics Network (WDN), bearing in mind that those results are based on firm-level (survey) data that do not cover the post-crisis sample.[37] First,

32. See also Gautier and others (2016).

33. In France the indexation of the minimum wage to harmonized indices of consumer prices (HICP) inflation feeds through to a large part of base wages and thereby leads to an informal wage indexation; the minimum wage also acts as a benchmark for wage agreements.

34. See Kohns (2018).

35. See Kuegler and others (2018).

36. See Conflitti and Zizza (2018).

37. See ECB (2009).

our general result that on average across sectors and countries the passthrough from labor cost to price inflation is positive and significant is consistent with the WDN result that a large percentage of firms surveyed declare that they use a strategy of increasing prices when faced with a (permanent) unexpected increase in wages, especially if firms produce intermediate goods. Second, the WDN finds that, at the micro level, the strength of the link between prices and labor cost depends on the labor share. In particular, firms with a high labor cost share report more frequently that there is a tight link between price and wage change. If we check the sectors that drive the highest passthrough across countries we are not able to confirm this result. With the exception of France, where the construction sector has the highest passthrough and the highest labor share, in the other countries the highest passthrough happens in sectors that have had the lowest labor share over the sample of the analysis (services in Italy and Germany, and manufacturing in Spain)[38].

These results, together with the findings in section 2, would suggest that, contrary to the results of the empirical literature based on U.S. data[39], there is no evident or systematic decline in passthrough across euro-area countries or sectors. One possible explanation for this divergent finding can simply be the consequence of the different detrending strategy that we adopt, i.e., by imposing a theory-based long-run restriction that the gap between productivity-adjusted nominal wage growth and price inflation disappears in the long run because the two variables share a common trend.[40]

3.2.2 Implications for the Behavior of the Price-Cost Markup

From a theoretical perspective, the markup should be measured by the price-marginal cost fraction. Empirically, however, measuring the marginal cost is often fraught with important difficulties.[41] For this reason, marginal cost is often proxied by average cost and, more precisely, by average labor cost. Theoretically, a number of conditions exist under which the marginal cost equals the average cost. For

38. See charts in Appendix E.

39. See Peneva and Rudd (2017) and references therein.

40. We have computed a time-varying passthrough for the U.S. data by using the same specification as in Peneva and Rudd (2017), removing a common trend from adjusted labor cost and price inflation, and the results confirm this intuition.

41. For a detailed discussion on existing approaches to measuring the price-cost markup, see Nekarda and Ramey (2013).

instance with a Cobb-Douglas technology and no labor adjustment costs, the marginal wage would equal the average wage, and hence the price-average labor cost fraction would represent the markup. With a constant elasticity of substitution (CES) production function and perfect substitution of labor *vis-a-vis* other non-labor inputs, it is also possible to show that the difference between price and labor cost inflation is the price-cost markup. Since we find an incomplete passthrough from labor costs to prices, our results have thus implications for the price-cost markup.[42]

The implication from our estimation results for the price-cost markup is shown in figure G1 in appendix G. The figure shows the evolution of the price-cost markup as the difference between the impulse response of price inflation and labor cost inflation. Moreover, it also shows the cumulative response on the price-cost markup for the total economy. Overall, the figure confirms, by looking at the results through a different lens, the incomplete passthrough with price-cost markups being squeezed following a positive labor cost shock. Concretely, following a one percent shock to labor cost inflation, the price-cost markup instantaneously declines in the total economy by around 0.8 percent across countries.

3.3 Main Findings: State-dependent VAR Specification

Another important dimension in the context of the passthrough from labor cost to price inflation is to test the empirical proposition that this passthrough could depend on the level of price inflation. We look at this particular variable because reduced-form estimates of the passthrough from labor costs to price inflation capture the underlying nominal rigidities, and the literature has highlighted that these rigidities may, *inter alia*, depend on the level of inflation.

A low passthrough can be associated to a low inflation environment either because low inflation and low expected inflation persistence cause a low passthrough,[43] or because low levels of price inflation could be expected to reduce the passthrough due to downward wage

42. We acknowledge that other costs might make up part of the difference between price and labor-costs growth, in particular the cost of capital. Nevertheless, grasping the cost of capital is a complicated problem beyond the scope of this paper and encompasses issues such as the price of intangible assets or quality-adjusted prices of information and communications technology goods.

43. See Taylor (2000).

rigidities.[44] Another argument that has been suggested as to why the passthrough from costs to inflation could increase with the level of inflation is linked to the search intensity of consumers. Concretely, at low levels of inflation, a large fraction of buyers observe a single price. In that case, any given shock would increase price dispersion sharply, which would increase the search intensity of consumers, thereby reducing firm market power, which limits the ability of firms to pass on the cost increase to prices. At higher levels of inflation, price dispersion is higher and hence any given shock has only a limited impact on price dispersion and the search intensity of consumers. As a result, prices are, at higher levels of inflation, more responsive to shocks.[45]

Finally, in a high inflation environment profits might act less as a buffer than in a low inflation regime due to an intertemporal smoothing of the profit path. For instance, when inflation is high and wages increase, firms may expect an increase in interest rates, which worsens their borrowing conditions and squeezes their future profit margins; hence, they will maintain their profits in the present, which would favor the passthrough from labor costs to prices.

Conversely, the opposite might hold in a lower inflation regime, where decreases in interest rates are expected. Another explanation could relate to the higher degree of economic uncertainty associated with a high inflation regime: in such a regime firms may simply not be prepared to buffer a labor cost increase with margins. Overall, the implicit margin responses in the high and low inflation regimes, as shown in appendix J, confirm this intuition, i.e., that margins act less as a buffer under high compared to a low inflation regime.

Our sample is not long enough to test this proposition on two regimes. However, in our VAR analysis, we can directly test whether this is also the case for euro-area countries as the reduced-form estimates of the passthrough from labor costs to price inflation would capture the underlying nominal rigidities. Therefore, we repeat the above exercise by estimating the VAR over two sets of observations using a dummy variable approach, with the level of inflation in one subset above and in the other below the corresponding historical averages, respectively. Country results for the total economy are reported in figure 8—the results for the other sectors can be found in appendix I.

44. See Daly and Hobijn (2014).
45. See Head and others (2010).

Figure 8. Choleski Decomposition-Based Passthrough from Labor Cost to Price Inflation under Low versus High Price Inflation

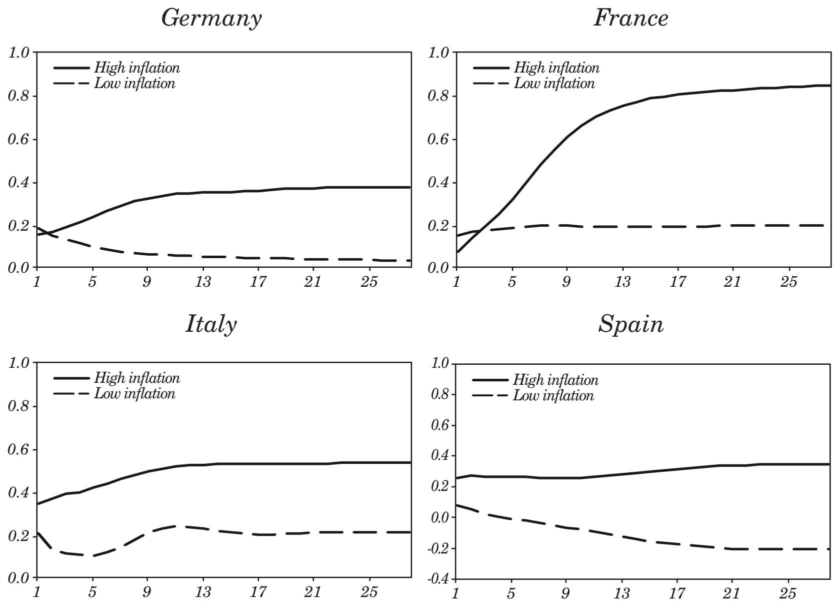

Source: Authors' calculations.

The findings support the theoretical and the U.S.-based empirical literature. Across euro-area sectors and countries (with the exception of the construction sector in Italy), the passthrough is systematically higher if it is estimated over samples when the inflation rate of the corresponding sector is higher than the historical average. The finding also supports the view that a pickup in labor cost inflation is a necessary condition for rising inflation, to the extent that higher inflation expectations associated with a change from lower to higher inflation rates could raise the passthrough, which in turn could speed up the inflationary process again.

4. Is the Link between Labor Costs and Price Inflation Shock-Dependent?

One of the challenges in empirically grasping the link between labor costs and prices arises from the fact that the passthrough may simultaneously depend on several factors. The previous sections allowed us to obtain a preliminary indication of the size of the passthrough from labor cost to price inflation and of the extent to which it has changed over time or has been dependent on the state of the economy (e.g., the level of inflation).

This analysis, however, does not allow us to identify the source of the correlation between labor cost and price fluctuations or the nature of the exogenous shocks that move labor cost inflation and are subsequently being passed on to price inflation. In this section, we want to take a step further and argue that the passthrough is not a deep parameter underlying the economy, but a shock-dependent coefficient that reflects the mechanisms underlying macro fluctuations.

We know, for instance, that in a New Keynesian model, the conditional correlation between labor cost and prices is different for demand and for supply shocks. The idea of the relationship between variables being shock-dependent has also recently been advocated in the exchange rate empirical literature ,[46] but also for understanding the Phillips curve relationship.[47] The same idea, translated to the labor cost passthrough to inflation, has recently become popular in policy circles.[48] Gumiel and Hahn (2018) present evidence based on the new area-wide model (NAWM), where the response of the GDP deflator to wages is stronger for demand than for supply shocks, where the latter capture frictions in the wage setting such as the impact of structural reforms or downward wage rigidity.

46. See Forbes and others (2018), Comunale and Kunovac (2017), and references therein.

47. See Galí and Gambetti (2018).

48. The shock dependency of the passthrough should depend on the degree of both price and labor-cost stickiness. The theoretical literature analyzing this issue is however scant. Most studies have focused on the impact of shocks on both labor cost and price inflation rather than on the passthrough of labor costs to price inflation following such shock. For instance, Bils and Chang (2000) did put forward a theoretical framework in which price rigidity differs with the nature of shocks, with prices being more responsive to increases in costs generated by factor prices than to an increase in marginal costs generated by an expansion of output. Model-based results show that prices react more to a technology (supply) shock than to a preference (demand) shock. Although this paper spells out clearly that it is important to disentangle between the nature of the shocks in seeing how prices react, it does not speak precisely to the question we are interested in, i.e., the passthrough from wages to prices.

Table 2. The Two-Shock VAR: Identification Scheme

Variables	*Shocks*		
	Demand	*Supply*	*Other*
Real value added	+	+	•
Prices	+	–	•
Labor cost	•	•	•

Source: Authors' calculations.
Notes: • = unconstrained, + = positive sign, – = negative sign

4.1 A Structural VAR Analysis

We address the question of the passthrough shock-dependence in the same three-variable VARs and identify a supply- and a demand-type shock for all countries and sectors by using the most parsimonious set of sign restrictions as reported in table 2.

Specifically, a positive demand shock is a shock that increases output growth and price inflation, whereas a supply shock increases output growth but reduces price inflation. Labor costs are left unrestricted, as a certain shock can affect wages and productivity in the same direction and the relative impact is not straightforward. A third shock in the model is left unidentified. The restrictions are imposed only for the first period and as inequality restrictions. The VAR is estimated as in the previous section with Bayesian techniques and a normal-diffuse prior with a Minnesota prior for the mean and the variance of the VAR parameters. Impulse responses are computed based on 5000 draws from the posterior simulators.

The baseline results from our estimation are reported in appendix L. By construction, we find that output and price inflation rise after a positive demand shock, but that output rises and price inflation decreases after a positive supply shock. Labor cost growth tends to decrease immediately after a positive demand shock and rise thereafter—which can be due to the fact that the increase in wages is smaller than the one in productivity, as the output tends to grow more than employment, as suggested by Gumiel and Hahn (2018). After a positive supply shock, labor cost inflation increases.

Equipped with these estimates we run two counterfactual experiments. In the first experiment we compute the counterfactual labor cost and price inflation that would be generated by a demand or a supply shock, and check how the correlation structure between the counterfactual variables changes according to the shock. In the second experiment we compute the counterfactual responses of price inflation to demand or supply shocks, and check how much amplification we give up by shutting down the labor cost channel, i.e., the response of labor cost inflation to the same shock.

Figure 9. Maximum Correlation between Price Inflation at (time *t*) and Labor Cost Inflation (time *t* – *k*) and the Lag for Maximum Correlation

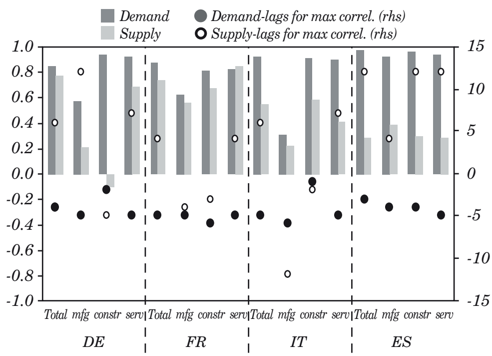

Source: Authors' calculations.
Note: The chart shows the cross-correlation between counterfactual price inflation at time t and labor cost inflation at time *t*–*k*.
Sample period: 1985Q1–2018Q1.

4.2 The Correlation between Labor Cost and Price Inflation Conditional on Demand and Supply Shocks

The first experiment consists in computing a historical decomposition and isolating the counterfactual labor cost inflation and price inflation that would have been generated by demand or supply shocks only. The correlation structure between these counterfactual series is then checked as in Galí (1999). We compute the maximum correlation over a wider lead/lag structure. Results are reported in figure 9, which shows the cross-correlation between the counterfactual price inflation at time *t* and labor cost inflation at time *t*–*k*. From the figure one can see that demand shocks affect prices and labor costs in a similar manner and prices appear to lead labor costs in their response to demand shocks. Conversely, supply shocks appear to affect prices and labor costs differently, with in most cases labor costs leading price inflation. The figure also shows that the correlation between labor cost inflation and price inflation tends to be higher for demand than for supply shocks. This simple fact can help to shed some light on the lack of consensus in the empirical literature that has tried to disentangle the direction of causality in the wage-price inflation nexus:[49] results are likely to depend

49. See Knotek and Zaman, 2014.

on the sample and on the combination of shocks hitting the economy over that particular sample.

4.3 The Amplification Due to the Labor Cost Channel

In the second experiment, we check the importance of the labor cost channel as an amplifier for the response of price inflation. In this case, we identify the same demand- and supply-type of shocks and then compare the response of price inflation in a system where all variables endogenously react to the initial shock with the response of price inflation in another system where the response of labor costs has been shut down. This will tell us how much of the shock is passed on to prices via labor costs.

To give an intuition for this approach, consider a positive demand shock which boosts prices as firms have a higher pricing power and their demand for inputs of production also increases. Of all the mechanisms through which demand shocks affect prices, one particular channel relates to labor costs. We would like to isolate this channel by gauging the impact of demand shocks on prices through labor costs. We will compute an impulse response function (IRF) where the response of labor costs to a demand shock is zero, and check the difference between the unrestricted IRF for price inflation and the IRF for the same variable when labor costs do not react to demand shocks. This difference is an indication of how much of the impact of demand shocks on inflation is driven by labor costs.

The idea of studying amplification mechanisms in a VAR by building a counterfactual scenario in which a certain variable does not react to a particular shock has been previously explored for other purposes. The impact of oil price shocks has been assessed via the reaction of inflation expectations by Wong (2015), or via the reaction of monetary policy by Kilian and Lewis (2011) and Bernanke and others (1997), who took inspiration from an early version of Sims and Zha (2006). Bachmann and Sims (2012) apply the same methodology to isolate the role of confidence in the transmission of government spending shocks, while Ciccarelli and others (2015) identify the effects of monetary policy shocks via the credit channel. What these papers have in common is that they operate with a VAR framework identified with contemporaneous zero restrictions. In a Choleski framework, each variable has a corresponding shock; it is straightforward to shut down the IRF of a variable by constructing a sequence of hypothetical shocks in that variable in a recursive manner, such that its IRF is zero at all times.

Figure 10. Amplification of Price Inflation Response due to the Labor Cost Channel

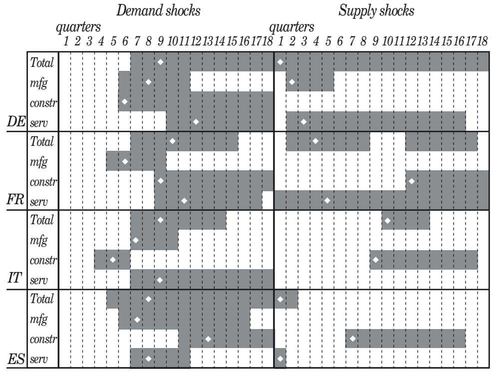

Source: Authors' calculations.
Note: This chart indicates, in grey, the quarters following a demand or a supply shock where the median counterfactual IRF lies outside the 68 percent posterior uncertainty band of the unrestricted IRF; borderline cases were left out. The white diamond indicates the quarter for maximum impact of the price inflation response.

When we move away from the Choleski identification scheme to one based on sign restrictions, as we propose in this paper, things get more complicated. Let's say we want labor costs not to react to demand shocks; there is no *labor cost shock* to offset the response of labor costs to demand shocks. One would have to make certain assumptions on which other shocks are doing the offsetting (i.e., is it the technology, other shocks, or—our preferred version—a combination of all shocks hitting the economy). Appendix N shows how we derive the counterfactual IRFs.

Results of the counterfactual exercise are summarized in figure 10 (and appendix O). This figure shows, in a synthetic manner, the quarters when we find a notable difference between the impulse responses with and without the labor cost channel, by marking these quarters with blue cells. The white diamond shows the quarter for which this amplification reaches its peak.

The striking feature is that, for all countries and sectors, there is a notable amplification under demand shocks, whereas the amplification under supply shocks occurs in fewer instances. In other words, when the economy is predominantly hit by demand-type shocks, it is more likely that the increase in wages above productivity is passed on to inflation than when the economy is predominantly hit by supply-type shocks. It is worth noting that the peaks of this passthrough tend to occur at a higher lag for demand-type shocks than for supply-type shocks.

The question that arises from these results is: why would labor costs tend to be passed through to prices when the economy is hit by a demand rather than by a supply shock? This analysis cannot provide a definite answer. However, it can be reconciled with previous findings whereby the willingness of firms to increase prices after labor cost increases is larger when positive demand shocks dominate. In such an environment, the share of higher income consumers with lower demand elasticity increases, which in turn raises firms' ability and power to pass through cost increases to prices.[50] This has implications for the differentiated behavior of the markup. In an environment where labor costs increase due to demand shocks, the price-cost markup would act as a buffer to a smaller extent than when the increase occurs due to supply shocks. The literature has stressed that the cyclicality of the markup is conditional on various types of shocks.[51] We find that under a positive demand shock, margins are procyclical, as seen in appendix P. Initially, the price-cost markup increases as price inflation increases, while labor costs growth increases by less or even declines in some instances. In a second stage, margins start to decline, as labor cost growth starts to increase (e.g., employment increases with delay) and they subsequently stabilize. Under a positive supply shock, margins appear to be countercyclical. They decrease because price inflation falls, while labor costs growth increases. The evolution of the price-cost markup is similar in the unrestricted and counterfactual scenario. What differs is the magnitudes of adjustment. In the medium term, the price-cost markup tends to stabilize at lower levels in the unrestricted world compared to the counterfactual, which reflects a positive passthrough of labor costs to prices; also, this difference on the medium-term is more notable for the markup following a demand shock, which reflects the more sizable passthrough in case of demand shocks.

We also acknowledge the caveat that the trivariate VAR is insufficient to properly identify supply-type shocks which in our parsimonious representation are identified based on the negative co-movement between output and prices. This simple identification scheme can in fact hide various types of supply shocks. One can imagine three types of such shocks, all of them increasing output and reducing prices: (i) a positive technology/productivity shock, which increases wages; (ii) a negative wage mark-up shock, which reduces

50. See Dornbusch (1987), and Bergin and Feenstra (2001).
51. See Galí and others (2007), and Nekarda and Ramey (2013).

wages; and (iii) a positive labor supply shock, which also reduces wages. The next subsection deals with this idea.

4.4 Robustness Check: A Structural VAR with Labor Market Shocks

In this subsection we check the robustness of the results obtained above along two dimensions: first, we enrich the identification scheme with more shocks on the labor market and, second, instead of a VAR including labor cost inflation, we consider a VAR including both wage and productivity growth separately, and we construct counterfactual IRFs where we impose that the difference in wage and productivity growth is shut down after a certain shock hits.[52]

The first issue we address is particularly important because what we identify as a 'supply shock' based on the negative co-movement between output and prices could in fact bundle together various types of shocks, as said above, and this complicates the assessment of the passthrough following a certain shock.

The VAR is now composed of five variables, namely: real value added, GDP deflator, nominal compensation per employee, labor productivity, and unemployment rate. All variables except the unemployment rate are expressed in annual growth rates, with the GDP deflator and nominal compensation adjusted by long-term expectations, as previously discussed. The system can only be estimated on the total economies since unemployment rate data does not exist at the sectoral level.

Table 3. The 4-shock VAR: Identification Scheme

Variables	*Shocks*				
	Demand	*Supply*	*Labor supply*	*Wage mark-up*	*Other*
Real value added	+	+	+	+	•
Prices	+	–	–	–	•
Wages	+	+	–	–	•
Productivity	+	+	•	•	•
Unemployment rate	–	•	+	–	•

Source: Authors' calculations.

Notes: • = unconstrained, + = positive sign, – = negative sign

52. See details in appendix N.

Figure 11. Amplification of Price Inflation Response due to the Labor Cost Channel in the 4-sShock VAR

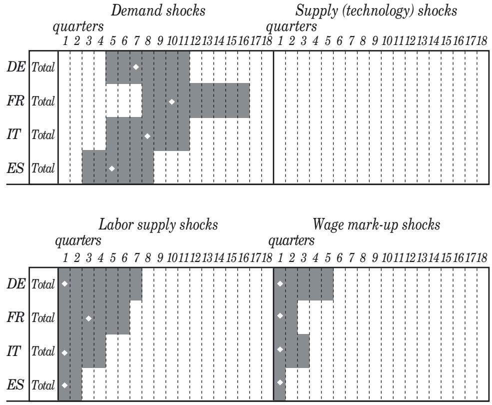

Source: Authors' calculations.
Note: This chart indicates, in gray, the quarters following a demand or a supply shock where the median counterfactual IRF lies outside the 68 percent posterior uncertainty band of the unrestricted IRF; borderline cases were left out. The white diamond indicates the quarter for maximum impact of the price inflation response.

Besides the *classical* demand and supply shocks, this system allows us to identify two more labor market shocks, as shown in table 3. A positive labor supply would increase the labor force participation, which translates into a positive impact on output and on the unemployment rate. Wage growth falls, and so does inflation; the different wages response allows disentangling labor supply from technological shocks, as explained in Peersman and Straub (2009). A wage mark-up shock, or a wage bargaining shock, is a shock that allows firms to capture a larger share of the bargaining surplus, which contributes to lower marginal costs, wage growth, and inflation. Output increases and the unemployment rate decreases, as detailed in Foroni and others (2018).[53]

Results are reported in figure 11 and in appendix Q. Overall, the results from the larger VAR model confirm the findings in the previous subsection, namely that labor costs are being passed through to price inflation in an environment where demand shock are predominant. When it comes to supply shocks, it turns out that the 'classical' supply (technology) shocks play a negligible role in the passthrough of labor costs to price inflation, but supply shocks originating from the labor

53. The estimation has been performed by using the BEAR toolbox, see Dieppe and others (2016).

market, namely labor supply and wage mark-up shocks, do matter and they trigger a fast transmission (in line with an identified smaller lag of maximum impact in case of supply shocks in figure 10). These results hold also when controlling for monetary policy. In appendix R we identify an additional monetary policy shock by including in the VAR model the spread between the long- and the short-term interest rates prevailing in each country.[54]

5. Summary and Conclusions

Understanding the signal labor cost developments are providing for the euro-area inflationary process is of key relevance from a policy perspective. For instance, the projections for euro-area inflation are based on the expectation that increasing labor market tightness will push up wage growth and, given a rather flat outlook for labor productivity, the resulting higher unit labor cost increases should be passed on, at least partly, to prices. However, to date, there does not exist a study which systematically analyses the empirical link between labor cost inflation and price inflation for the euro area and the euro-area countries. In this paper we document this link for the first time.

When using country and sector quarterly data over the period 1985Q1–2018Q1, we uncover a number of facts. First, we find that the cost-push view of inflation found in the economic theory can have some support in the data. We document a strong link between labor cost and price inflation in the four major economies of the euro area and across three sectors (manufacturing, construction and service).

Second, the analysis supports an imperfect but relatively high passthrough on average from costs to prices, in line with available firm-level evidence which documents a statistically significant relationship from the frequency of wage changes to that of prices, and a common strategy by several firms of increasing prices when faced with unexpected increases in wages.[55]

Third, the link between price and labor cost growth is quite heterogeneous across countries and sectors. France is the country where this passthrough is higher, with the link being strongest in the construction sector. In Germany and Italy the driving sector is

54. This measure could reflect the monetary policy stance also in the unconventional monetary policy period—see Baumeister and Benati, 2013—, but admittedly also non-policy factors affecting the term structure, such as sovereign debt issues.

55. See Druant and others (2009).

services, while in Spain the manufacturing sector shows the highest passthrough. Hence, with the exception of France, this evidence contrasts with the idea that the passthrough of wages into prices should be particularly strong in firms/sectors with a high labor share, i.e., sectors which should also be characterized by a higher degree of price stickiness.[56]

Fourth, the dynamic interaction between prices and wages is time-varying and depends on the state of the economy. In particular, the passthrough is systematically lower in periods of low inflation as compared to periods of high inflation. This result would be in line with an expectation theory as proposed, e.g., by Taylor (2000), whereby *a decline in the degree to which firms pass through changes in costs to prices is frequently characterized as a reduction in the pricing power of firms.*

Fifth, the wage-price passthrough also depends on the shocks hitting the economy. The results presented show that it is more likely that the labor costs are passed on to price inflation with demand shocks than with supply shocks. This result holds also when we augment the dynamic system to disentangle more clearly various types of supply shocks, e.g., to capture frictions in the wage setting such as the impact of structural reforms or downward wage rigidity. Rationalizing this result is not simple, as there is no comprehensive theoretical literature which focuses on the difference in the wage passthrough to inflation according to different shocks. Some limited theoretical frameworks are available where price rigidity differs with the nature of shocks, with prices being more responsive to increases in costs generated by factor prices driven by technology than to increases in marginal costs generated by an expansion of output driven by preferences,[57] but nothing can be inferred about the passthrough from wages to prices.

These results have clear implications for the behavior of profit margins or price-cost markups. In an environment where labor costs increase due to demand shocks, the price-cost markup would act as a buffer to a smaller extent than when the increase occurs due to supply shocks.

Finally, our results support the view that a pickup in labor cost growth can drive underlying inflation and confirm the idea that, under circumstances of predominantly demand shocks, labor cost increases

56. See Druant and others (2009).
57. See Bils and Chang (2000).

will be passed on to prices.[58] After a period of low inflation, however, this passthrough could be moderate at least until inflation stably reaches a sustained path.

58. See Gumiel and Hahn (2018).

References

Bachmann, R. and E.R. Sims. 2012. "Confidence and the Transmission of Government Spending Shocks." *Journal of Monetary Economics* 59(3): 235–49.

Banerji, A. 2005. "The Relationship between Labor Costs and Inflation: A Cyclical Viewpoint." Economic Cycle Research Institute, U.S. Bureau of Labor Statistics.

Baumeister, C. and L. Benati. 2013. "Unconventional Monetary Policy and the Great Recession: Estimating the Macroeconomic Effects of a Spread Compression at the Zero Lower Bound." *International Journal of Central Banking* 9(2): 165–212.

Bergin, P.R. and R.C. Feenstra. 2001. "Pricing-to-Market, Staggered Contracts, and Real Exchange Rate Persistence." *Journal of International Economics* 54(2): 333–59.

Bernanke, B.S., M. Gertler, and M. Watson. 1997. "Systematic Monetary Policy and the Effects of Oil Price Shocks." *Brookings Papers on Economic Activity* 28(1): 91–157.

Bidder, R. 2015. "Are Wages Useful in Forecasting Price Inflation?" *FRBSF Economic Letter*.

Bils, M. and Y. Chang. 2000. "Understanding How Price Responds to Costs and Production." *Carnegie-Rochester Conference Series on Public Policy* 52(1): 33–77.

Canova, F. and M. Ciccarelli. 2009. "Estimating Multicountry VAR Models." *International Economic Review* 50: 921–69.

Charsonville, L.D., F. Ferrire, and C. Jardet. 2017. "MAPI: Model for Analysis and Projection of Inflation in France." Technical report.

Ciccarelli, M., A. Maddaloni, and J.L. Peydro. 2015. "Trusting the Bankers: A New Look at the Credit Channel of Monetary Policy." *Review of Economic Dynamics* 18(4): 979–1002.

Comunale, M. and D. Kunovac. 2017. "Exchange Rate Passthrough in the Euro Area." Working Paper Series 2003. European Central Bank.

Conflitti, C. and R. Zizza. 2018. "What's behind Firms' Inflation Forecasts?" Questioni di Economia e Finanza (Occasional Papers) 465. Bank of Italy, Economic Research and International Relations Area.

Daly, M.C. and B. Hobijn. 2014. "Downward Nominal Wage Rigidities Bend the Phillips Curve." Working Paper Series 2013–08. Federal Reserve Bank of San Francisco.

Darrat, A. 1994. "Wage Growth and the Inflationary Process: A Reexamination." *Southern Economic Journal* 61(1): 181–90.

Dees, S. and J. Guntner. 2014. "Analyzing and Forecasting Price Dynamics across Euro Area Countries and Sectors: a Panel VAR Approach." Working Paper Series 1724. European Central Bank.

Dieppe, A., B. van Roye, and R. Legrand. 2016. "The BEAR Toolbox." Working Paper Series 1934. European Central Bank.

Dornbusch, R. 1987. "Exchange Rates and Prices." *American Economic Review* 77(1): 93–106.

Draghi, M. 2018. "Monetary Policy in the Euro Area." Speech delivered at the ECB Forum on Central Banking, Sintra. 19 June 2018, European Central Bank.

Druant, M., S. Fabiani, G. Kezdi, A. Lamo, F. Martins, and R. Sabbatici. 2009. "How Are Firms' Wages and Prices Linked: Survey Evidence in Europe." Working Paper Series 1084, European Central Bank.

Du Caju, P., C. Fuss, and L. Wintr. 2009. "Understanding Sectoral Differences in Downward Real Wage Rigidity: Workforce Composition, Institutions, Technology and Competition." Working Paper Series 1006. European Central Bank.

European Central Bank, ECB. 2004. "Recent Developments in Unit Labour Costs and their Implications for Euro Area Inflation." *Monthly Bulletin, September*.

European Central Bank, ECB. 2009. "Wage Dynamics in Europe. Final Report of the Wage Dynamics Network December 2009." European Central Bank.

European Central Bank, ECB. 2018. "ECB Staff Macroeconomic Projections for the Euro Area." Projections article, European Central Bank, December.

Emery, K. and C.P. Chang. 1996. "Do Wages Help Predict Inflation?" Economic Review First Quarter 1996, Dallas Fed.

Fisher, I. 1926. "A Statistical Relation between Unemployment and Price Changes." *International Labour Review* 13(6): 85–792.

Forbes, K., I. Hjortsoe, and T. Nenova. 2018. "The Shocks Matter: Improving our Estimates of Exchange Rate Passthrough." Working Paper 24773, National Bureau of Economic Research.

Foroni, C., F. Furlanetto, and A. Lepetit. 2018. "Labor Supply Factors and Economic Fluctuations." *International Economic Review* 59(3): 1491–510.

Galí, J. 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89(1): 249–71.

Galí, J. and L. Gambetti. 2018. "Has the U.S. Wage Phillips Curve Flattened?" A Semi-Structural Exploration. CREI Working Papers October 2018, CREI.

Galí, J., M. Gertler, and J.D. López-Salido. 2007. "Markups, Gaps, and the Welfare Costs of Business Fluctuations." *Review of Economics and Statistics* 89(1): 44–59.

Gautier, E., D. Fougère, and S. Roux. 2016. "The Impact of the National Minimum Wage on Industry-Level Wage Bargaining in France." Technical report.

Gordon, R.J. 1988. "The Role of Wages in the Inflation Process." *American Economic Review* 78(2): 276–83.

Gumiel, J.E. and E. Hahn. 2018. "The Role of Wages for the Pick-Up in Inflation." Box in the *Economic Bulletin Issue 5*, European Central Bank.

Head, A., A. Kumar, and B. Lapham. 2010. "Market Power, Price Adjustment, and Inflation." *International Economic Review* 51(1): 73–98.

Hess, G. and M. Schweitzer. 2000. "Does Wage Inflation Cause Price Inflation?" Federal Reserve Bank of Cleveland, Policy Discussion Paper No. 00–01.

Hu, L. and M. Toussaint-Comeau. 2010. "Do Labor Market Activities Help Predict Inflation?" *Economic Perspectives* (Q II): 52–63.

International Monetary Fund (IMF). 2018. "2. European Wage Dynamics and Labor Market Integration." IMF eLibrary, U.S.A.

Jarocinski, M. and B. Mackowiak. 2017. "Granger Causal Priority and Choice of Variables in Vector Autogressions." *Review of Economics and Statistics* 99(2): 319–29.

Kadyiala, R. and S. Karlsson. 1998. "Numerical Methods for Estimation and Inference in Bayesian VAR Models." *Journal of Applied Econometrics* 12: 99–132.

Kilian, L. and L.T. Lewis. 2011. "Does the Fed Respond to Oil Price Shocks?" *Economic Journal* 121(555): 1047–72.

Knotek, E.S. and S. Zaman. 2014. "On the Relationships between Wages, Prices, and Economic Activity." *Economic Commentary* (August).

Kohns, S. 2018. "Introductory Statement to the Bundesbank Conference 'Wages and Inflation in Germany'." Presentation, Bundesbank.

Kozicki, S. and P. Tinsley. 2001. "Shifting Endpoints in the Term Structure of Interest Rates." *Journal of Monetary Economics* 47(3): 613–52.

Kuegler, A., U. Schoenberg, and R. Schreiner. 2018. "Productivity Growth, Wage Growth and Unions." Paper presented at the ECB Forum on Central Banking Sintra, Portugal, 18–20 June 2018.

Mehra, Y.P. 2000. "Wage-price Dynamics: Are They Consistent with Cost Push?" *Economic Quarterly* (Sum): 27–43.

Mountford, A. and H. Uhlig. 2009. "What Are the Effects of Fiscal Policy Shocks?" *Journal of Applied Economics* 24(6): 960–92.

Nekarda, C. and V. Ramey. 2013. "The Cyclical Behavior of the Price-Cost Markup." Working Paper 19099, National Bureau of Economic Research.

Peersman, G. and R. Straub. 2009. "Technology Shocks and Robust Sign Restrictions in a Euro Area SVAR." *International Economic Review* 50(3): 727–50.

Peneva, E.V. and J.B. Rudd. 2017. "The Passthrough of Labor Costs to Price Inflation." *Journal of Money, Credit and Banking* 49(8): 1777–802.

Quevat, B. and B. Vignolles. 2018. "The Relationships between Inflation, Wages and Unemployment have not Disappeared. A Comparative Study of the French and American Economies." Conjuncture in France, March 2018, INSEE.

Rissman, E.R. 1995. "Sectoral Wage Growth and Inflation." *Federal Reserve Bank of Chicago Economic Review*, pages 16–28.

Sbordone, A. 2002. "Prices and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49: 265–92.

Sims, C.A. and T. Zha. 2006. "Does Monetary Policy Generate Recessions?" *Macroeconomic Dynamics* 10(02): 231–72.

Stock, J.H. and M.W. Watson. 2008. "Phillips Curve Inflation Forecasts." NBER Working Papers 14322.

Tatierska, S. 2010. "Do Unit Labor Costs Drive Inflation in the Euro Area?" Working Paper Series 2/2010, National Bank of Slovakia.

Taylor, J. 2000. "Low Inflation, Passthrough, and the Pricing Power of Firms." *European Economic Review* 44: 1389–408.

Wong, B. 2015. "Do Inflation Expectations Propagate the Inflationary Impact of Real Oil Price Shocks?: Evidence from the Michigan Survey." *Journal of Money, Credit and Banking* 47(8): 1673–89.

Zaman, S. 2013. "Improving Inflation Forecasts in the Medium to Long Term." Economic commentary, Federal Reserve Bank of Cleveland.

Zanetti, A. 2007. "Do Wages Lead Inflation?" Swiss Evidence. *Swiss Journal of Economics and Statistics* 143(1).

Appendices

Appendix A

Data Documentation

Most standard data (i.e., nominal and real value added, compensation of employees, total employees) were obtained as seasonally and working-day adjusted series from national accounts over the period 1985Q1–2018Q1 for the four biggest euro-area countries. All series were obtained for the aggregate economy and three sectors: manufacturing, construction, and services. Short- and long-term interest rates come from the ECB Statistical Data Warehouse (https://sdw.ecb.europa.eu/home.do). Unemployment rates were obtained from Eurostat (and back-casted with seasonally adjusted data from the IMF International Financial Statistics (IFS) in the case of Germany and with data from national sources through Haver Analytics in the case of Spain). A number of series were derived on the basis of the national accounts data. The value added deflator was calculated as the ratio of the nominal to real value added. Labor productivity was measured as the ratio of real value added to total employees, while compensation per employee was calculated as the ratio of compensation of employees to total employees. Finally, unit labor costs were calculated as the ratio of compensation per employee to labor productivity. More details on the country-specific national accounts data are listed below:

Germany: Official aggregate and sectoral data on real value added, nominal value added, compensation of employees and total employees were obtained from the Federal Statistical Office through Haver Analytics. In the case of the services sector and total employees, all long time series were constructed by chain linking the ESA2010 (NACE2) and ESA1995 (NACE1) databases. The series were adjusted for the structural break due to unification. Data prior to 1991 is for West Germany only. For services, data prior to 1991 is the sum of hotels and transport, finance and business services, and public and personal services.

France: Official aggregate and sectoral data on real value added, nominal value added, compensation of employees and total employees were obtained from the French National Institute of Statistics and Economic Studies (INSEE) through Haver Analytics. Services sector data were calculated as the sum of market and non-market services.

Italy: Official aggregate and sectoral data on real value added, nominal value added, compensation of employees and total employees were obtained from the Italian National Statistics Institute (ISTAT) through Haver Analytics. In the case of the services sector, all long time series were constructed by chain linking the ESA2010 (NACE2) and ESA1995 (NACE1) databases.

Spain: Official aggregate and sectoral data on real and value added, compensation of employees and total employees were obtained from Spanish National Statistics Institute (INE) through Haver Analytics. With the exception of the total economy data, long series were constructed by chain linking the ESA2010 (NACE2) and ESA1995 (NACE1) databases. For services, data prior to 1995 is the sum of market and non-market services series. Historical data on real value added and compensation of employees was obtained from the INE website. Long historical data on the manufacturing sector was not available, the data used is for industry.

Appendix B

GDP Deflator and CPI Series

Figure B1. Labor Cost, GDP Deflator and CPI

(year-on-year percentage change)

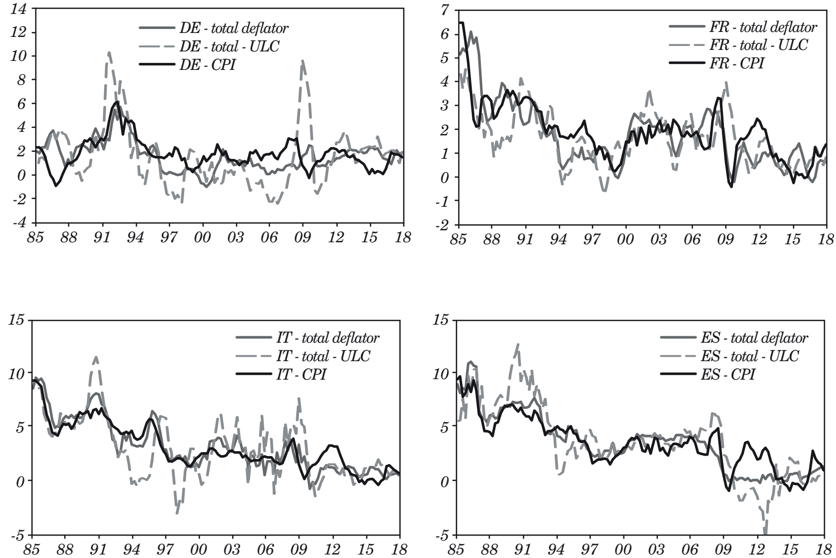

Sources: Various sources, authors' calculations.
Latest observation: 2018Q1.

Appendix C

Cross-Correlations by Sectors and across Time

Figure C1. Cross-Correlation between Adjusted Labor Cost and Price Inflation since 2008

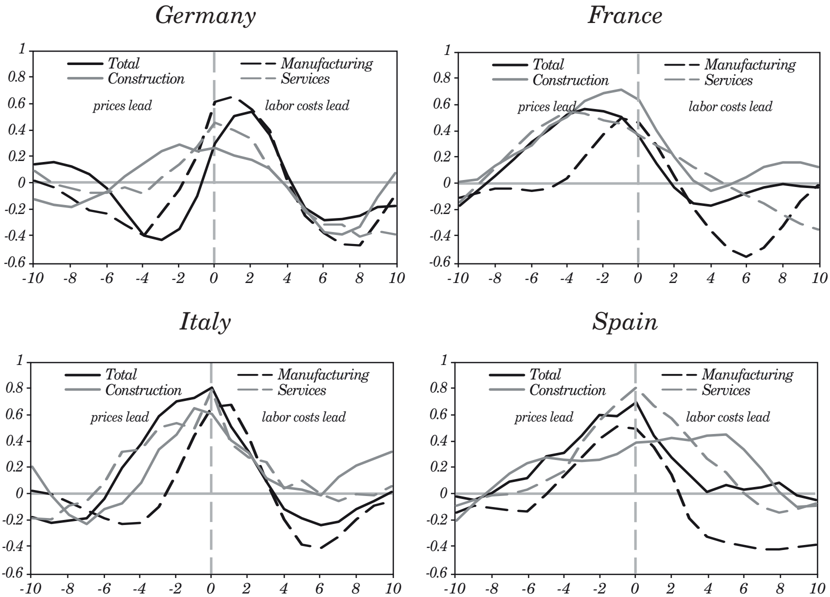

Sources: Various sources, authors' calculations.
Note: The charts show the cross-correlation between price inflation gaps at time t and labor cost inflation gaps at time t-k. Sample period: 2008Q1–2018Q1.

Appendix D

Forecasting Power of Labor Costs for Price Inflation

Table D1. Ratio of RMSE of Inflation Forecasts of Models with to Models without Labor Lost

	199Q1-2018Q1				*1999Q1-2007Q4*				*2008Q1-2018Q1*			
	Germany											
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	1.01	0.96	1.04	0.99	1.12	0.93	1.02	1.01	0.90	0.97	1.07	0.95
2	0.97	0.97	1.02	0.98	1.13	0.91	0.99	1.02	0.83	1.01	1.04	0.93
3	0.93	0.90	1.01	0.99	1.10	0.85	0.99	1.02	0.81	0.94	1.02	0.95
4	0.91	0.90	1.01	0.99	1.05	0.82	0.99	1.01	0.80	0.96	1.02	0.96
5	0.89	0.87	1.00	0.99	1.05	0.82	0.99	1.01	0.78	0.93	1.00	0.98
6	0.87	0.84	0.99	0.99	1.04	0.82	0.99	1.00	0.77	0.83	0.99	0.99
7	0.85	0.77	0.99	1.00	0.99	0.89	0.99	0.99	0.77	0.68	0.99	1.00
8	0.87	0.75	0.99	1.00	0.98	0.99	0.99	1.00	0.79	0.62	1.00	1.01
	France											
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	1.00	1.02	1.04	0.96	1.04	1.01	1.06	0.99	0.96	1.03	1.00	0.94
2	0.96	1.01	1.07	0.91	0.95	1.01	1.12	0.90	0.96	1.01	1.01	0.93
3	0.93	1.02	1.07	0.88	0.88	1.01	1.12	0.80	0.96	1.02	1.02	0.92
4	0.91	1.01	1.05	0.85	0.83	0.99	1.09	0.72	0.97	1.01	1.05	0.92
5	0.88	0.99	1.04	0.84	0.77	0.97	1.07	0.69	0.96	0.99	1.06	0.91
6	0.87	0.97	1.02	0.86	0.78	0.95	1.05	0.74	0.95	0.98	1.06	0.91
7	0.88	0.97	1.01	0.89	0.81	0.95	1.05	0.83	0.94	0.97	1.03	0.91
8	0.90	0.98	1.01	0.92	0.86	0.96	1.06	0.93	0.93	0.99	1.00	0.90
	Italy											
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	0.95	1.09	1.00	1.00	0.92	1.06	1.00	1.01	0.99	1.11	1.00	0.99
2	0.99	1.04	1.00	1.01	0.95	1.06	1.00	1.02	1.05	1.04	1.00	0.98
3	1.05	1.01	1.00	1.01	1.05	1.01	1.00	1.03	1.06	1.01	1.00	0.99
4	1.07	1.01	0.99	1.03	1.07	1.02	1.01	1.06	1.08	1.01	0.99	0.99
5	1.06	1.01	0.99	1.02	1.04	1.00	0.99	1.05	1.08	1.01	0.99	1.00
6	1.06	1.01	0.99	1.00	1.03	0.99	0.98	1.01	1.09	1.02	0.99	1.00
7	1.05	1.01	0.98	1.01	1.04	0.98	0.97	1.01	1.07	1.03	0.98	1.00
8	1.04	1.02	0.98	0.99	1.01	0.98	0.95	0.98	1.06	1.04	0.98	1.00

Table D1. (continued)

	199Q1-2018Q1				*1999Q1-2007Q4*				*2008Q1-2018Q1*			
					Spain							
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	0.95	1.03	1.01	1.40	0.97	1.19	1.01	1.26	0.94	0.96	1.00	1.47
2	0.88	1.04	1.01	1.61	0.94	1.23	1.02	1.46	0.87	0.98	0.99	1.66
3	0.81	1.00	1.00	1.70	0.93	1.16	1.01	1.62	0.79	0.96	0.98	1.73
4	0.80	0.97	1.00	1.71	0.96	1.08	1.01	1.58	0.77	0.95	0.97	1.74
5	0.80	0.97	1.00	1.73	1.05	1.04	1.01	1.57	0.76	0.96	0.96	1.77
6	0.81	1.00	1.00	1.71	1.17	1.04	1.01	1.57	0.76	0.99	0.98	1.76
7	0.83	1.03	1.01	1.69	1.21	1.03	1.01	1.54	0.78	1.03	1.00	1.73
8	0.85	1.04	1.01	1.66	1.21	1.06	1.01	1.45	0.79	1.03	1.01	1.71

Table D2. Theil's U of Inflation Forecasts Conditional on Observed Path of Labor Cost

	199Q1-2018Q1				*1999Q1-2007Q4*				*2008Q1-2018Q1*			
						Germany						
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	1.00	0.91	0.89	0.99	1.00	0.86	1.09	1.07	1.03	1.03	1.09	0.98
2	0.98	0.95	0.85	0.97	0.97	0.78	1.04	1.07	1.06	1.13	1.12	0.96
3	0.96	0.95	0.82	0.94	0.94	0.72	0.94	1.11	0.98	1.15	1.08	0.85
4	0.96	0.85	0.81	0.93	0.97	0.66	0.89	1.16	0.79	1.03	1.04	0.81
5	0.93	0.79	0.80	0.86	0.96	0.64	0.82	1.12	0.59	0.86	0.97	0.75
6	0.89	0.77	0.82	0.79	0.94	0.62	0.81	1.07	0.49	0.83	0.93	0.72
7	0.87	0.76	0.87	0.73	0.96	0.66	0.84	1.04	0.44	0.80	0.94	0.72
8	0.84	0.72	0.91	0.66	0.96	0.71	0.88	0.88	0.46	0.73	0.94	0.69
						France						
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	0.78	0.94	0.83	0.99	0.90	0.96	0.95	1.07	0.90	0.95	0.73	0.99
2	0.75	0.88	0.78	1.00	0.88	0.93	0.98	1.07	0.86	0.84	0.63	1.01
3	0.74	0.83	0.77	0.99	0.86	0.90	1.04	1.05	0.88	0.77	0.60	1.03
4	0.77	0.80	0.79	0.98	0.87	0.85	1.09	1.01	0.93	0.74	0.60	1.04
5	0.83	0.80	0.86	0.97	0.96	0.83	1.17	0.98	0.96	0.75	0.63	1.05
6	0.90	0.84	0.94	0.95	1.03	0.89	1.23	0.98	0.99	0.79	0.67	1.04
7	0.95	0.91	0.99	0.92	1.06	0.97	1.18	0.99	1.00	0.85	0.70	1.01
8	0.99	0.94	0.99	0.88	1.03	1.00	1.11	0.99	1.01	0.88	0.70	0.96
						Italy						
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	0.69	0.96	0.82	0.58	0.76	0.87	0.75	0.65	0.72	0.98	0.75	0.60
2	0.73	0.89	0.84	0.69	0.78	0.71	0.76	0.72	0.79	0.89	0.71	0.79
3	0.71	0.83	0.82	0.72	0.75	0.65	0.70	0.75	0.77	0.81	0.68	0.84
4	0.67	0.79	0.78	0.68	0.69	0.61	0.60	0.66	0.74	0.73	0.66	0.80
5	0.71	0.76	0.79	0.73	0.71	0.62	0.64	0.67	0.76	0.67	0.69	0.86
6	0.73	0.78	0.76	0.79	0.77	0.65	0.70	0.75	0.75	0.68	0.67	0.82
7	0.76	0.82	0.80	0.81	0.78	0.71	0.83	0.75	0.76	0.73	0.74	0.84
8	0.77	0.84	0.97	0.82	0.76	0.75	1.23	0.74	0.77	0.75	0.93	0.84
						Spain						
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	0.96	0.93	1.03	1.05	0.80	0.87	0.95	1.25	0.82	0.94	1.16	1.24
2	1.02	0.89	1.06	1.07	0.87	0.81	0.97	1.36	0.76	0.89	1.06	1.33
3	0.97	0.87	1.07	1.07	0.86	0.80	0.96	1.49	0.70	0.84	1.01	1.34
4	0.99	0.88	1.07	1.07	0.86	0.83	0.98	1.55	0.69	0.82	0.99	1.31
5	0.96	0.91	1.02	1.09	0.86	0.92	0.98	1.71	0.67	0.80	0.86	1.22
6	0.94	0.93	0.94	1.09	0.86	1.04	0.97	1.86	0.68	0.76	0.73	1.16
7	0.95	0.93	0.87	1.09	0.90	1.09	0.96	1.89	0.73	0.73	0.67	1.06
8	0.94	0.95	0.82	1.05	0.92	1.08	0.92	1.57	0.77	0.71	0.60	1.02

Table D3. Ratio of RMSE of Labor Cost Inflation Forecasts of Models with to Models without Price Inflation

	199Q1-2018Q1				*1999Q1-2007Q4*				*2008Q1-2018Q1*			
Germany												
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	1.04	1.06	1.04	1.01	0.96	0.90	1.04	1.04	1.07	1.08	1.03	0.98
2	1.04	1.10	1.01	1.01	0.89	0.88	0.99	1.05	1.08	1.12	1.00	0.98
3	1.03	1.12	0.99	1.02	0.83	0.89	0.91	1.05	1.09	1.15	0.98	1.00
4	1.03	1.09	0.96	1.01	0.85	0.93	0.81	1.04	1.08	1.11	0.95	1.00
5	1.00	1.05	0.98	1.01	0.84	1.02	0.84	1.03	1.05	1.06	0.96	1.00
6	0.97	1.01	1.04	1.00	0.87	1.06	0.87	1.01	0.96	1.02	1.07	1.00
7	0.96	1.00	1.02	1.00	0.91	1.07	0.82	1.00	0.84	1.00	1.15	1.01
8	0.96	0.99	0.99	1.00	0.99	1.07	0.84	0.98	0.78	0.97	1.16	1.01
France												
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	1.00	1.01	1.00	1.02	1.01	0.99	1.00	1.02	1.00	1.01	1.00	1.02
2	1.01	1.00	0.98	1.04	1.01	0.99	1.00	1.04	1.00	1.01	0.97	1.04
3	1.01	0.99	0.97	1.05	1.02	0.97	0.99	1.06	1.01	1.01	0.96	1.04
4	1.01	0.97	0.97	1.06	1.01	0.96	0.99	1.07	1.01	0.99	0.95	1.05
5	1.01	0.97	0.97	1.09	1.01	0.99	1.01	1.10	1.01	0.99	0.94	1.07
6	1.01	1.00	0.97	1.11	1.00	1.05	1.00	1.11	1.02	1.02	0.95	1.12
7	1.01	1.02	0.97	1.12	1.00	1.10	0.99	1.10	1.04	1.05	0.96	1.17
8	1.01	1.02	0.97	1.11	1.00	1.09	0.98	1.09	1.05	1.01	0.97	1.20
Italy												
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	0.98	1.04	0.93	1.15	0.96	0.93	0.93	1.12	1.05	1.10	0.92	1.26
2	1.06	1.04	1.00	1.04	1.01	0.90	0.98	1.03	1.26	1.09	1.01	1.14
3	1.13	1.01	1.01	0.90	1.06	0.86	0.97	0.87	1.29	1.05	1.03	0.94
4	1.11	0.98	0.99	0.88	1.04	0.88	0.90	0.86	1.30	1.01	1.03	0.93
5	1.12	0.95	1.02	0.81	1.02	0.93	1.00	0.79	1.39	0.97	1.03	0.87
6	1.14	0.94	0.99	0.77	1.02	1.00	0.99	0.69	1.57	0.92	0.99	0.90
7	1.15	0.96	0.96	0.81	1.07	1.08	0.93	0.77	1.51	0.89	0.96	0.89
8	1.09	0.98	0.97	0.80	1.02	1.10	0.99	0.76	1.32	0.95	0.95	0.93
Spain												
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	0.98	0.96	0.93	1.02	1.00	1.00	0.92	0.81	0.98	0.93	0.93	1.08
2	0.94	0.93	0.96	1.02	0.96	0.99	0.92	0.79	0.92	0.88	0.97	1.08
3	0.92	0.90	1.00	0.97	0.96	0.99	0.93	0.71	0.90	0.83	1.02	1.02
4	0.90	0.87	1.03	1.00	0.95	1.00	0.94	0.64	0.88	0.77	1.05	1.06
5	0.88	0.85	1.04	1.00	0.94	1.01	0.94	0.60	0.87	0.77	1.06	1.06
6	0.89	0.84	1.05	1.00	0.95	1.00	0.94	0.60	0.88	0.78	1.08	1.05
7	0.90	0.83	1.04	0.99	0.95	0.99	0.95	0.63	0.89	0.79	1.08	1.04
8	0.91	0.85	1.03	0.98	0.95	0.99	0.97	0.67	0.90	0.82	1.06	1.01

Table D4. Theil's U of Labor Cost Inflation Forecasts Conditional on Observed Path of Price Inflation

	199Q1-2018Q1				*1999Q1-2007Q4*				*2008Q1-2018Q1*			
						Germany						
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	0.95	0.78	1.02	0.99	1.19	0.98	1.15	1.07	0.86	0.82	0.99	0.90
2	0.95	0.78	1.02	0.98	1.22	0.92	1.16	1.06	0.85	0.83	1.01	0.93
3	0.94	0.77	1.04	0.97	1.20	0.85	1.24	1.04	0.85	0.83	1.00	0.95
4	0.94	0.73	1.07	0.94	1.19	0.80	1.33	0.99	0.85	0.80	0.99	0.95
5	0.94	0.72	1.08	0.95	1.18	0.76	1.47	1.00	0.84	0.78	0.97	0.98
6	0.94	0.71	1.05	0.93	1.12	0.73	1.45	1.01	0.82	0.74	0.90	0.99
7	0.93	0.69	1.05	0.89	1.06	0.77	1.38	0.97	0.82	0.69	0.85	0.99
8	0.92	0.70	1.07	0.85	0.98	0.79	1.32	0.87	0.91	0.69	0.81	1.02
						France						
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	0.92	1.00	0.88	1.02	0.93	1.03	0.94	1.00	1.00	1.00	0.87	1.12
2	0.93	0.98	0.84	1.00	0.93	1.07	0.96	0.98	0.98	0.97	0.83	1.14
3	0.97	0.96	0.85	0.99	0.95	1.13	0.96	0.96	1.01	0.93	0.86	1.14
4	1.01	0.92	0.86	1.01	0.95	1.14	0.99	0.97	1.04	0.89	0.90	1.12
5	1.05	0.87	0.88	1.04	0.97	1.14	0.99	0.98	1.05	0.82	0.94	1.13
6	1.09	0.81	0.90	1.08	0.98	1.04	1.04	1.00	1.07	0.76	0.97	1.16
7	1.11	0.80	0.92	1.11	0.98	0.95	1.07	1.00	1.10	0.83	0.98	1.26
8	1.05	0.79	0.91	1.09	0.95	0.88	1.08	0.97	1.03	0.89	0.96	1.29
						Italy						
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	0.72	0.88	1.01	0.64	0.81	0.99	1.05	0.75	0.66	0.81	0.94	0.54
2	0.79	0.83	0.91	0.80	0.87	0.90	0.96	0.86	0.79	0.76	0.81	0.90
3	0.79	0.80	0.88	0.85	0.87	0.85	1.00	0.91	0.79	0.74	0.75	1.06
4	0.71	0.79	0.83	0.77	0.79	0.80	0.99	0.83	0.73	0.71	0.69	0.95
5	0.74	0.80	0.79	0.84	0.81	0.77	0.89	0.89	0.75	0.70	0.65	1.38
6	0.72	0.82	0.80	0.86	0.85	0.72	1.03	0.99	0.66	0.68	0.65	1.14
7	0.69	0.83	0.76	0.81	0.78	0.68	1.00	0.92	0.60	0.64	0.63	1.18
8	0.70	0.83	0.76	0.86	0.80	0.68	0.88	0.95	0.57	0.67	0.65	1.11
						Spain						
steps	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*	*total*	*mfg*	*const*	*serv*
1	1.01	0.99	0.98	0.98	0.97	1.00	1.00	1.01	0.99	1.06	0.99	0.99
2	1.05	0.99	0.97	0.99	1.03	1.09	1.08	1.01	1.02	1.02	0.99	0.99
3	1.04	0.94	0.92	0.99	1.11	1.04	1.05	1.02	1.01	0.94	0.98	0.99
4	1.03	0.90	0.89	0.96	1.16	0.95	1.08	1.02	1.01	0.87	0.98	0.99
5	1.03	0.84	0.86	0.95	1.22	0.83	1.10	1.01	0.99	0.79	0.94	0.99
6	1.01	0.82	0.82	0.94	1.21	0.77	1.11	1.00	0.95	0.70	0.89	0.98
7	1.02	0.83	0.81	0.93	1.22	0.75	1.12	0.98	0.95	0.71	0.84	0.98
8	1.02	0.86	0.81	0.93	1.28	0.83	1.09	0.97	0.93	0.72	0.80	0.97

Appendix E

Labor-Share Developments

Figure E1. Labor Share across Countries and Sectors

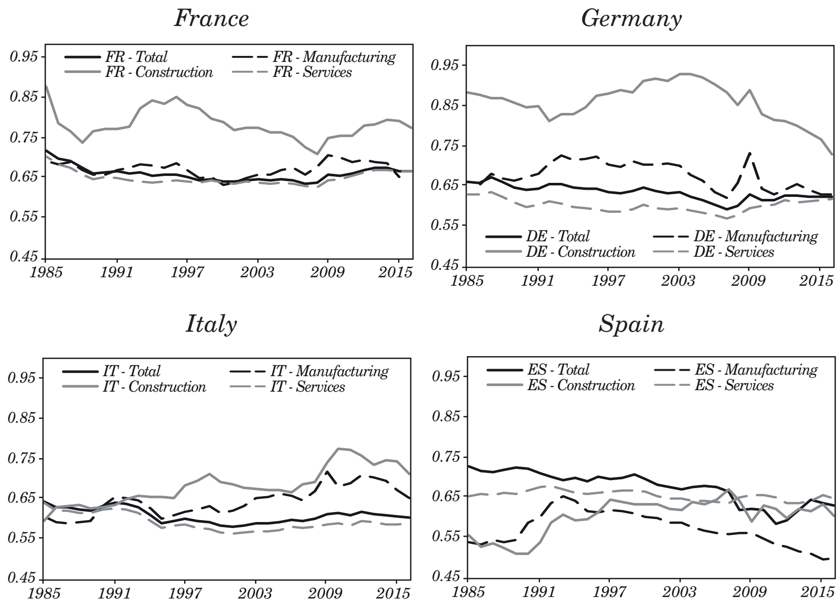

Source: Authors' calculations based on Organization for Economic and Co-operation and Development (OECD) data.

Appendix F

VAR-based Analysis: Impulse Responses from Choleski Orthogonalization and the Forecast Error Variance Decomposition

Figure F1. Steady-state Passthrough from Unit Labor Cost Inflation to Price Inflation

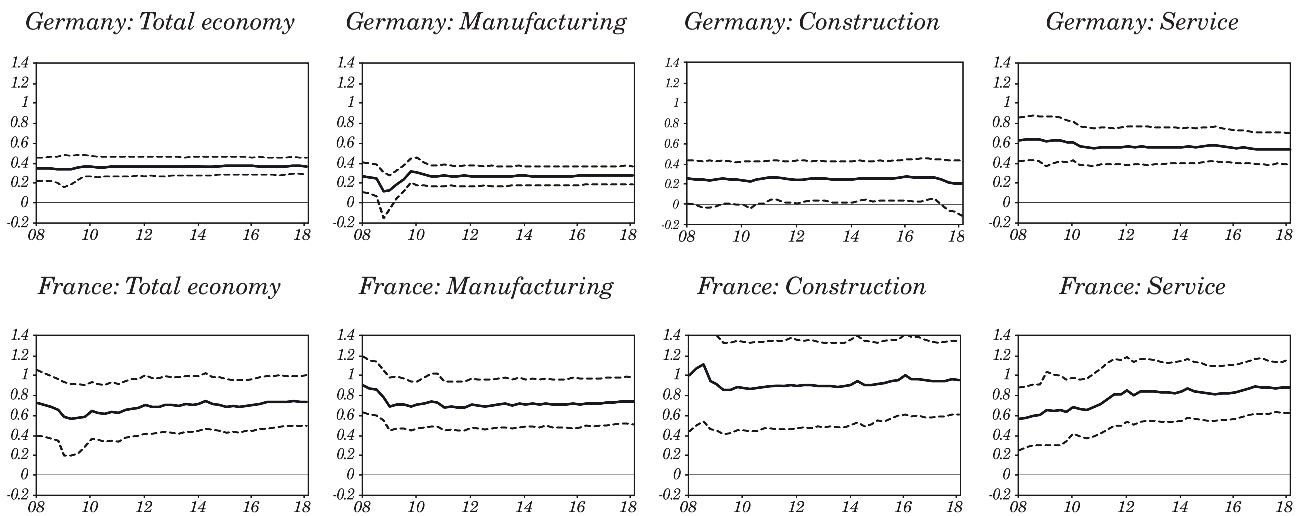

Figure F1. (continued)

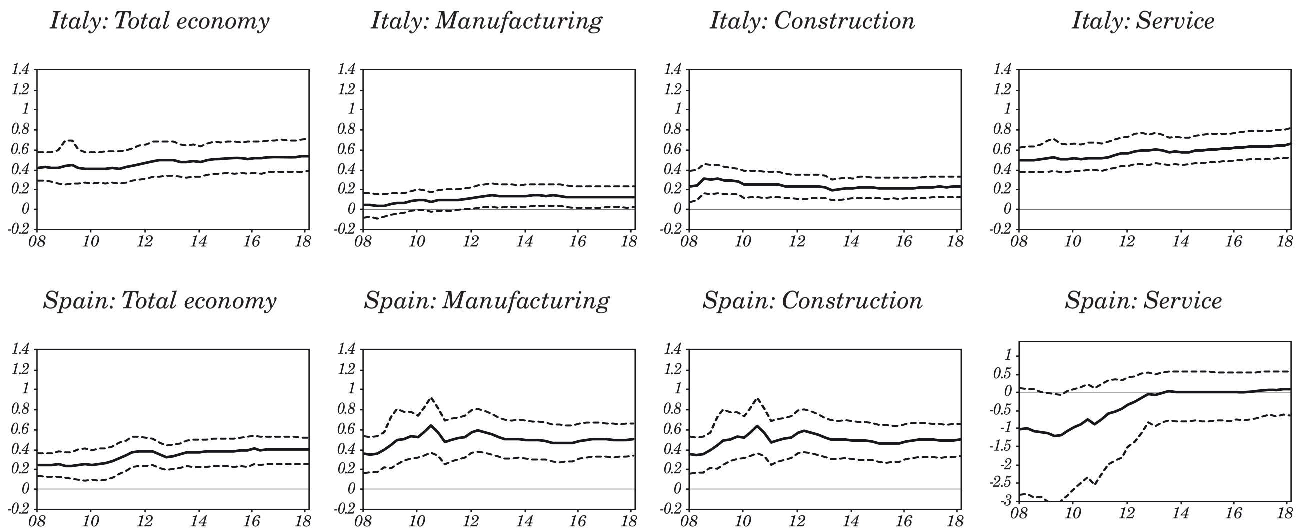

Source: Authors' calculations.
Note: The results show the steady-state impulse response (at quarter 40) from a time-varying approach whereby the first sample covered 1985Q1–2008Q1 and thereafter one quarter at a time was recursively added.
Sample period: 1985Q1–2018Q1.

Figure F2. Forecast Error Variance Decomposition (FEVD)

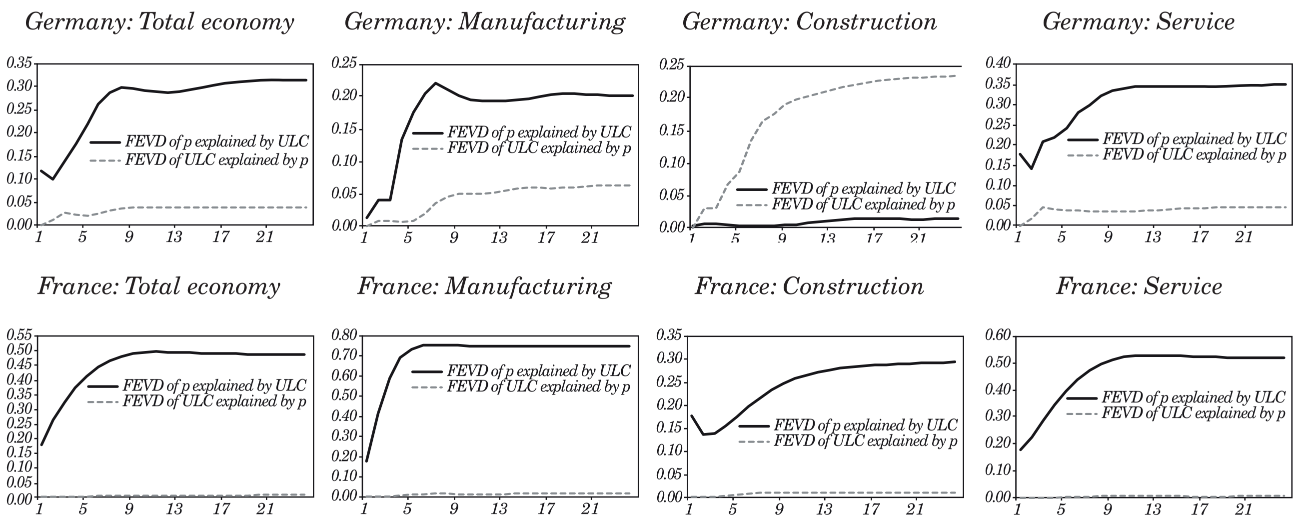

Figure F2. (continued)

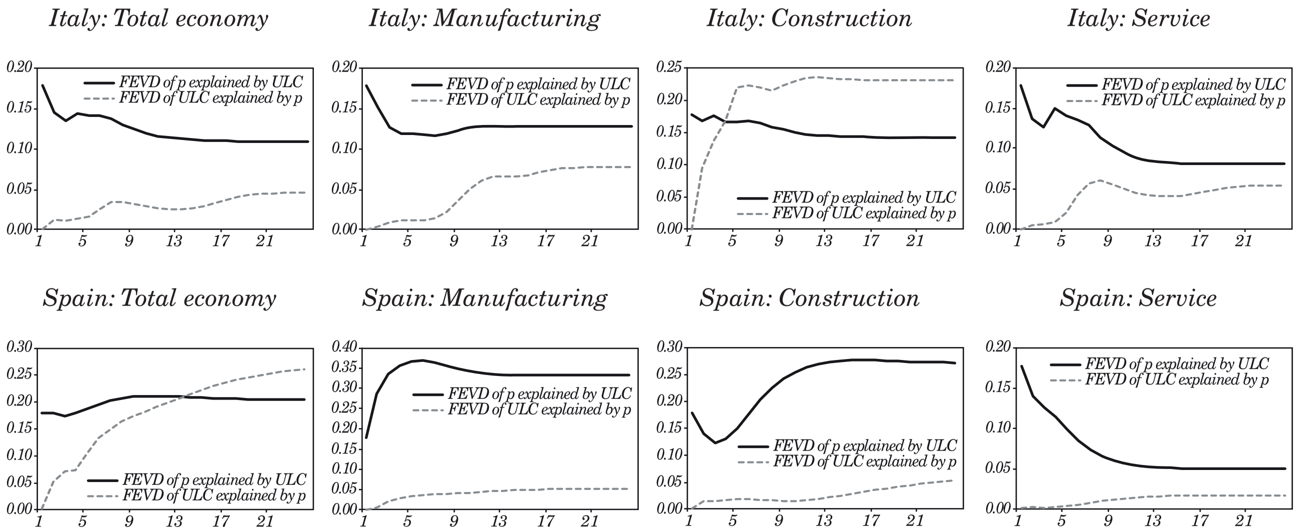

Source: Authors' calculations.
Note: The results show the steady-state impulse response (at quarter 40) from a time-varying approach whereby the first sample covered 1985Q1–2008Q1 and thereafter one quarter at a time was recursively added.
Sample period: 1985Q1–2018Q1.

Appendix G

VAR-based Analysis: Impulse Response from Choleski Orthogonalization —Implications for the Markup

Figure G1. Choleski Decomposition-based Passthrough from Labor Cost to Price-cost Markup

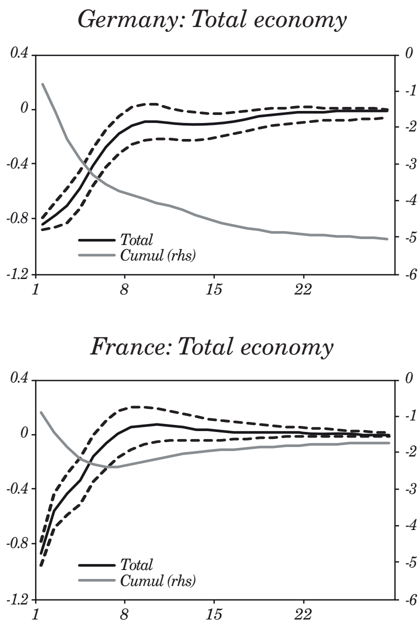

Figure G1. (continued)

Italy: Total economy

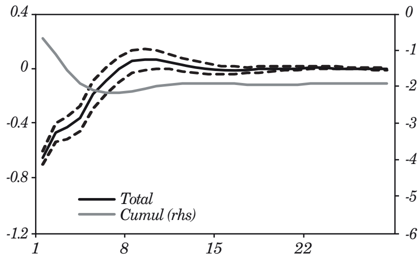

Spain: Total economy

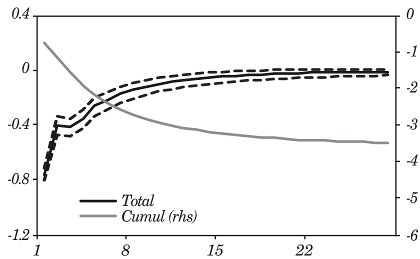

Source: Authors' calculations. The price-cost markup is calculated as the difference between the impulse response of price inflation to a shock to labor cost inflation.

Appendix H

VAR-based Analysis: Impulse Responses from Choleski Orthogonalization—Results from Model which includes Monetary Policy

Figure H1. Steady-state Passthrough from Unit Labor Cost Inflation to Price Inflation

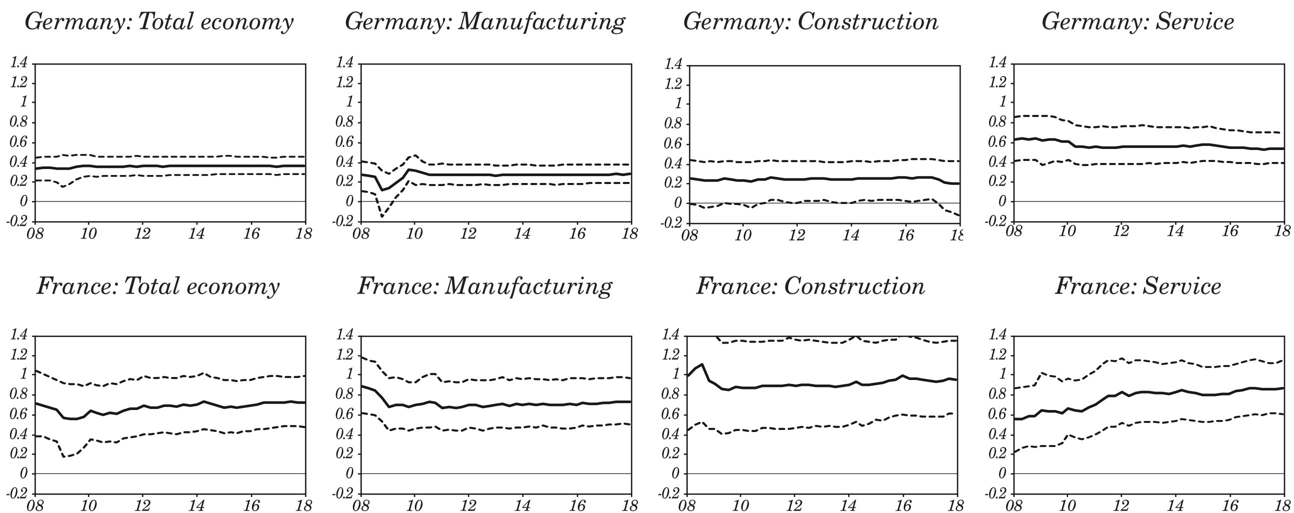

Figure H1. (continued)

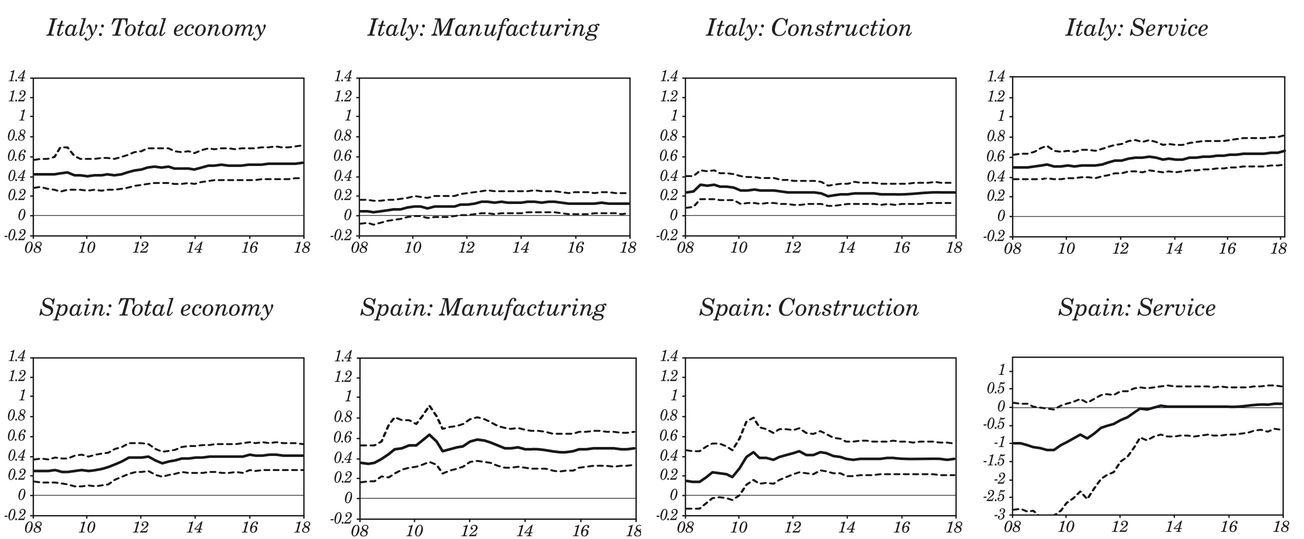

Source: Authors' calculations.
Note: Results from a four-variable VAR model (which includes real value added, the gross value added deflator, labor cost and monetary policy proxied by the spread between a long and a short-term interest rate. The results show the steady-state impulse response (at quarter 40) from a time-varying approach whereby the first sample covered 1985Q1–2008Q1 and thereafter one quarter at a time was recursively added.
Sample period: 1985Q1–2018Q1.

Appendix I

VAR-based Analysis: Impulse Responses from Choleski Orthogonalization under High versus Low Inflation

Figure I1. Passthrough from Unit Labor Cost Inflation to Price Inflation for High versus Low Inflation

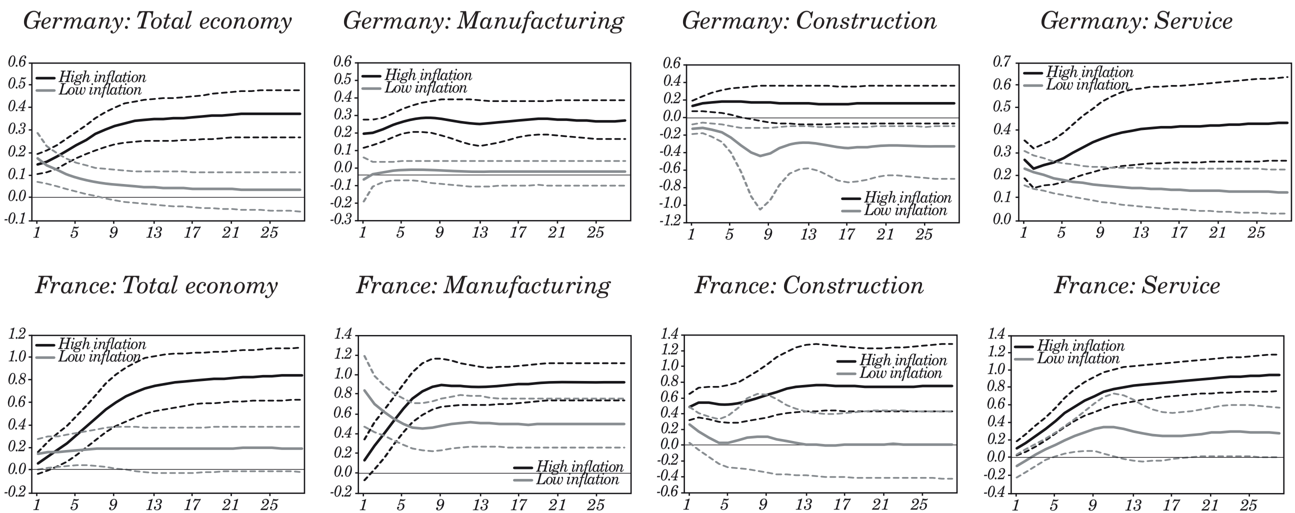

Figure I1. (continued)

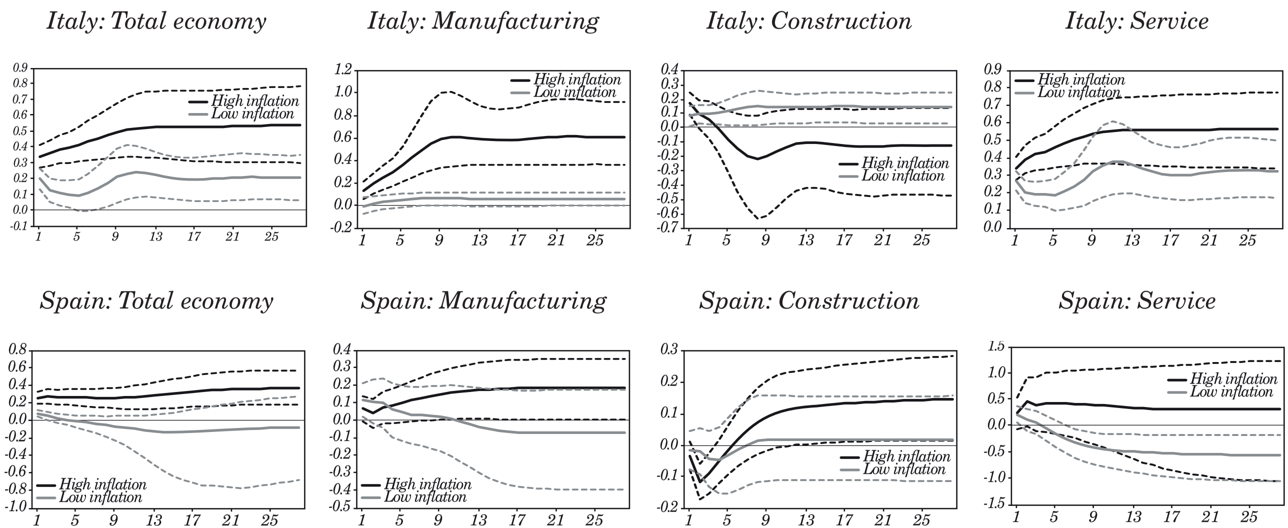

Source: Authors' calculations.
Sample period: 1985Q1–2018Q1.

Appendix J

VAR-based Analysis: Impulse Responses from Choleski Orthogonalization under High and Low Inflation Implications for Margins

Figure J1. Passthrough from Unit Labor Cost Inflation to Price Inflation for High versus Low Inflation

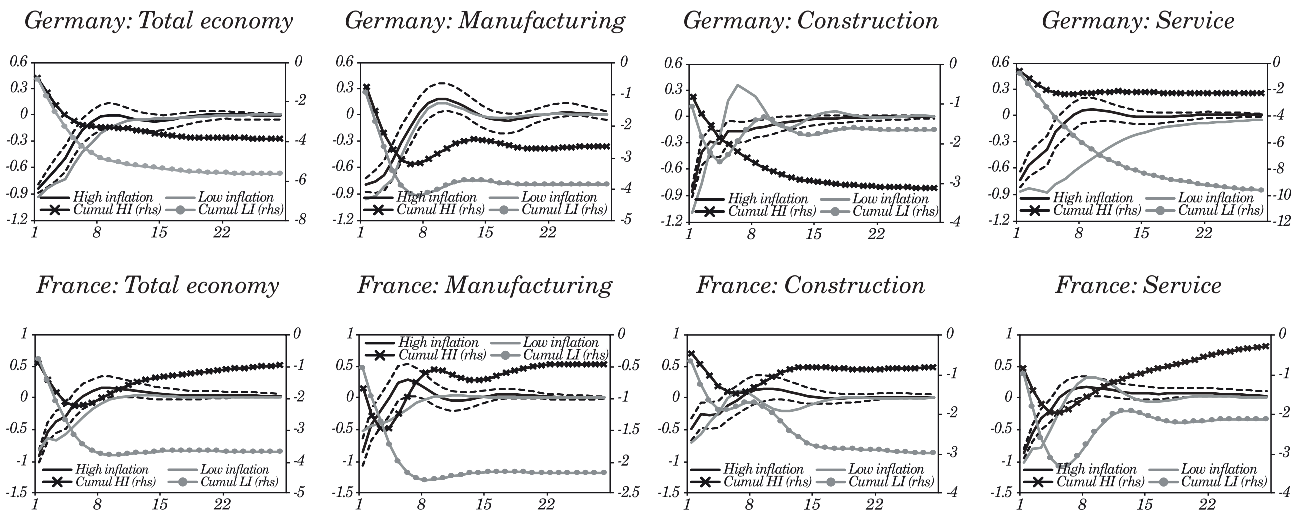

Figure J1. (continued)

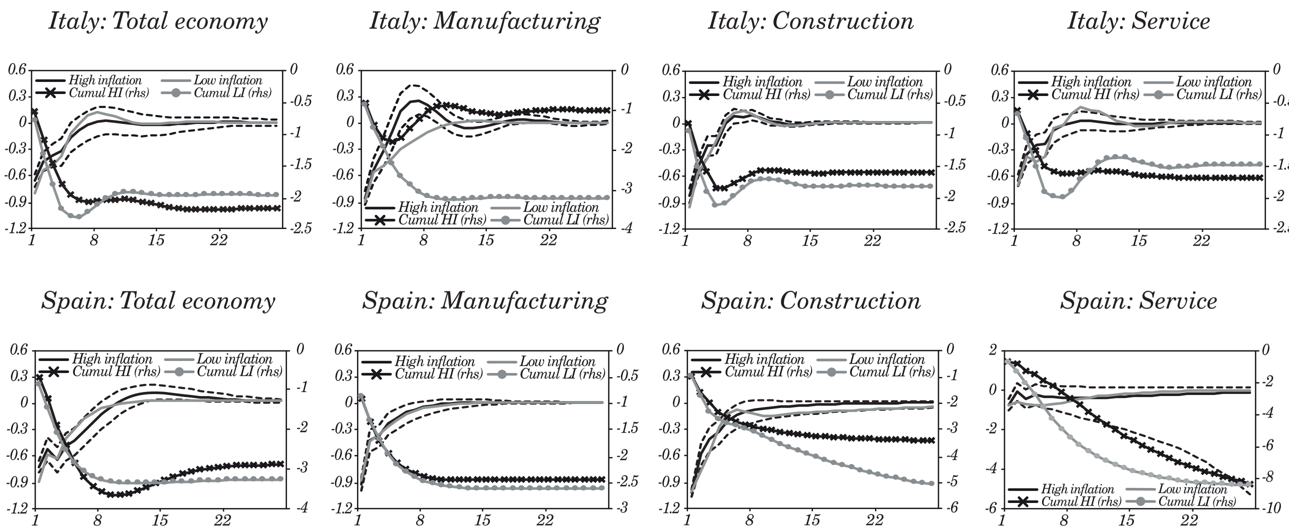

Source: Authors' calculations.
Sample period: 1985Q1–2018Q1.

Appendix K

VAR-based Analysis: Impulse Responses from Choleski Orthogonalization under High versus Low Inflation—Results from Model which Includes Monetary Policy

Figure K1. Passthrough from Unit Labor Cost Inflation to Price Inflation for High versus Low Inflation

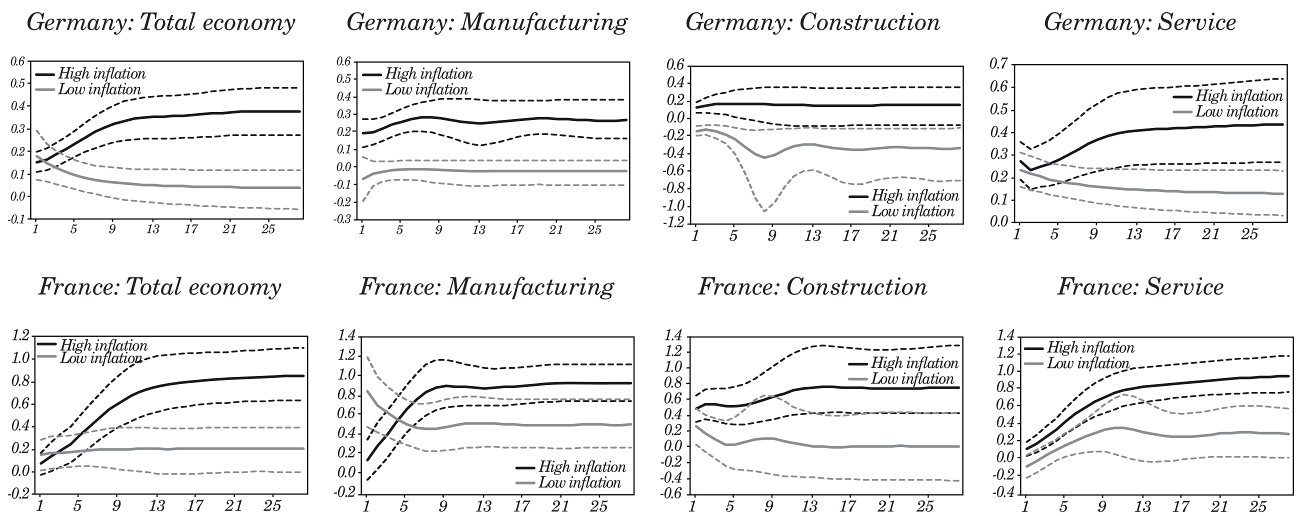

Figure K1. (continued)

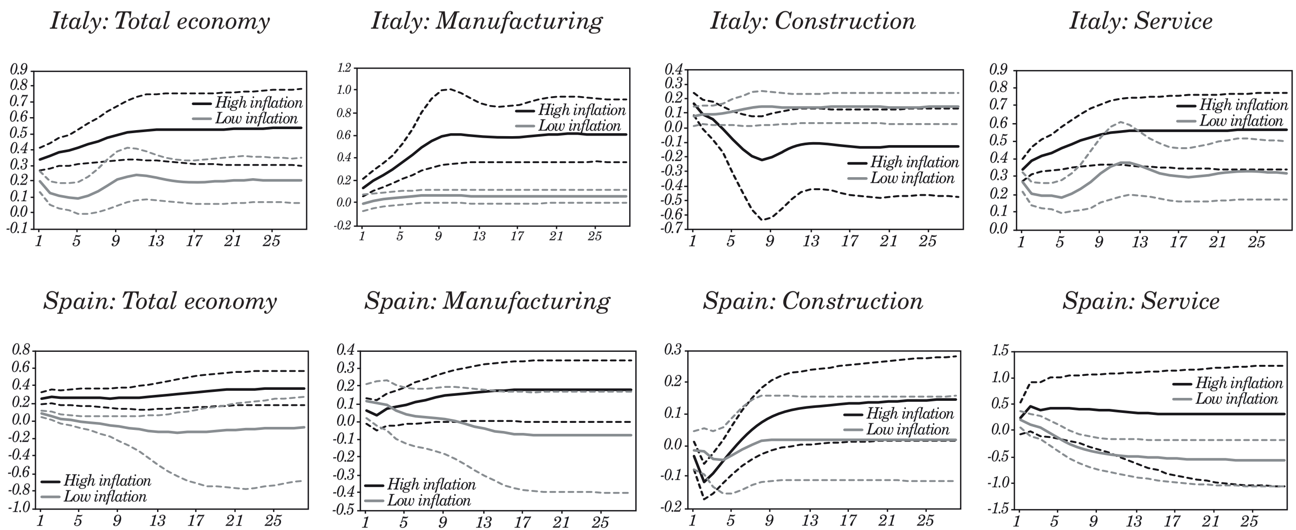

Source: Authors' calculations.
Sample period: 1985Q1–2018Q1.

Appendix L

SVAR with Sign Restrictions: Impulse Response Functions

Figure L1. Impulse Response Functions for the Total Economy

Germany

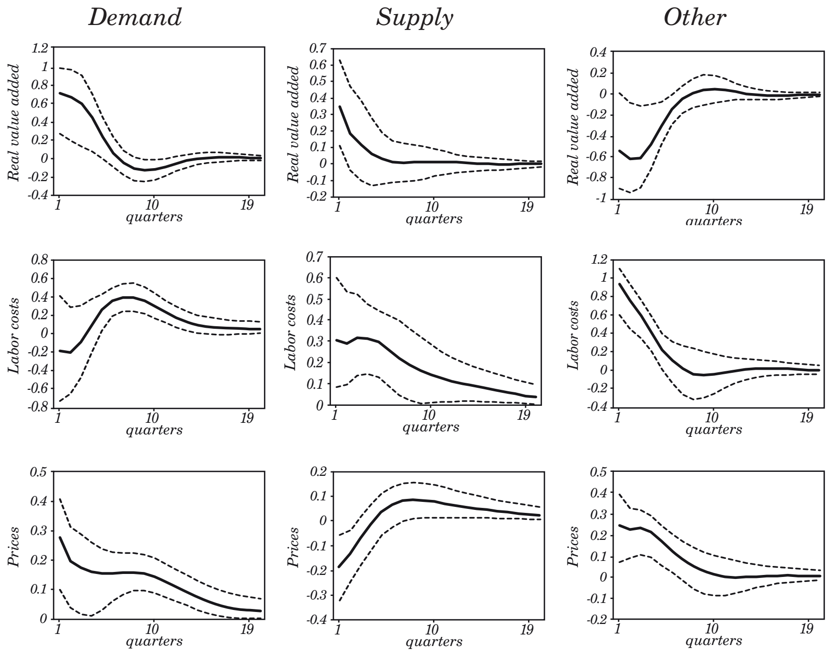

France

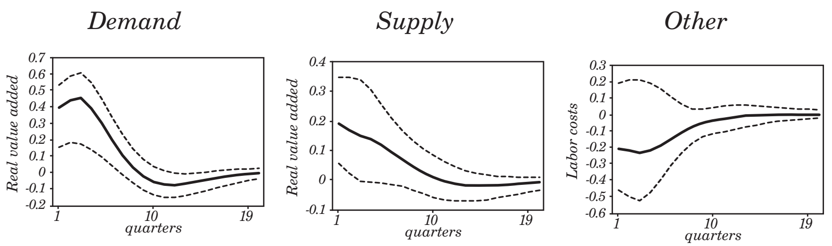

Figure L1. (continued)

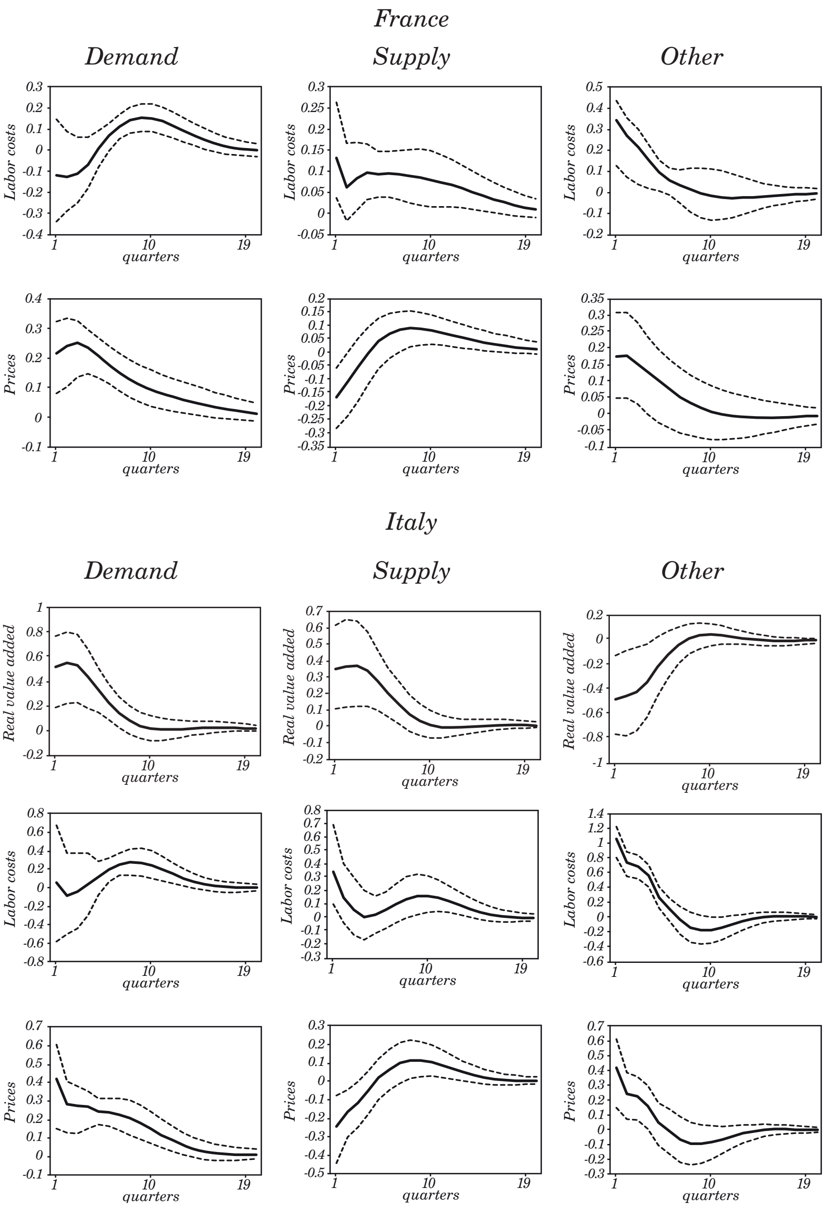

Figure L1. (continued)

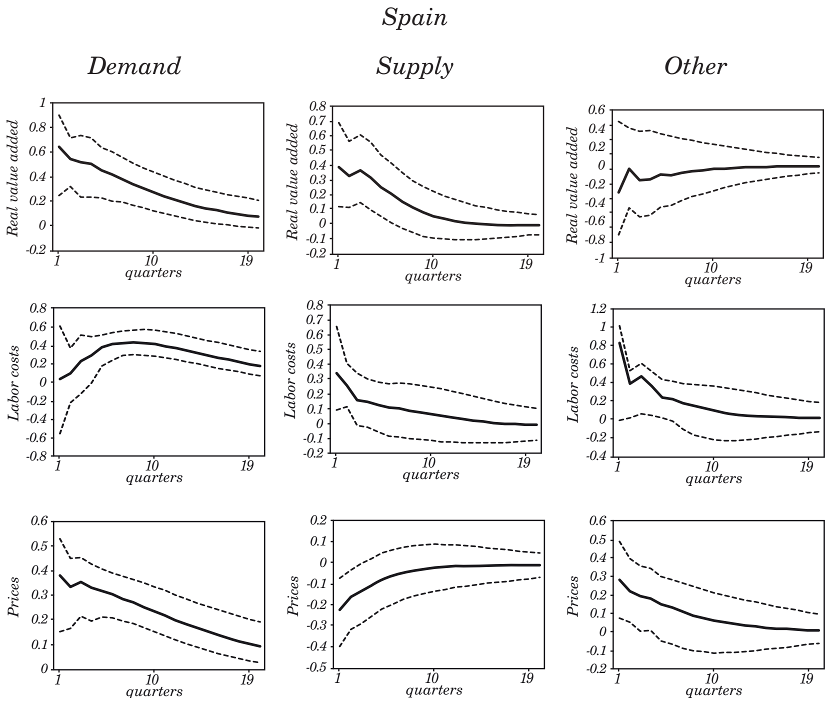

Source: Authors' calculations.

Appendix M

Sign-Restricted SVAR: Historical Contributions

Figure M1. The Contribution of Demand Shocks to Price and Labor Cost Inflation

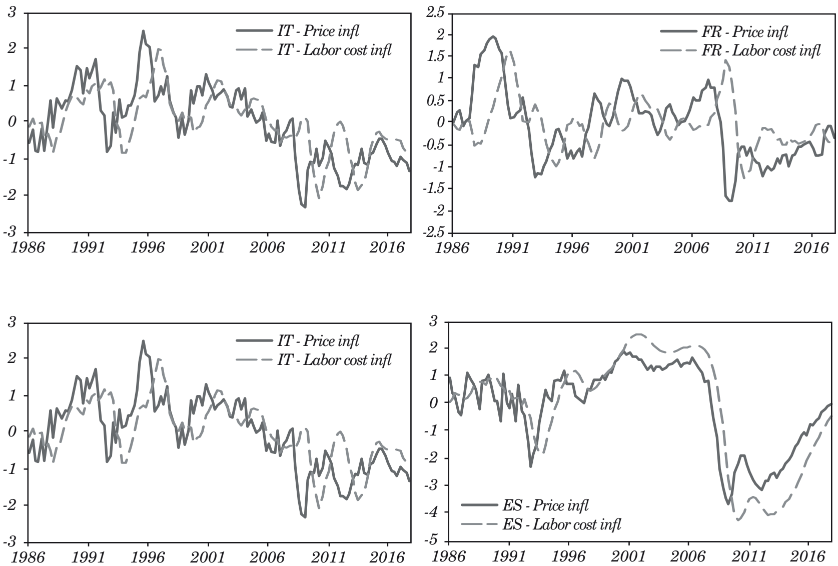

Source: Authors' calculations.
Sample period: 1985Q1–2018Q1.

Figure M2. The contribution of supply shocks to price and labor cost inflation

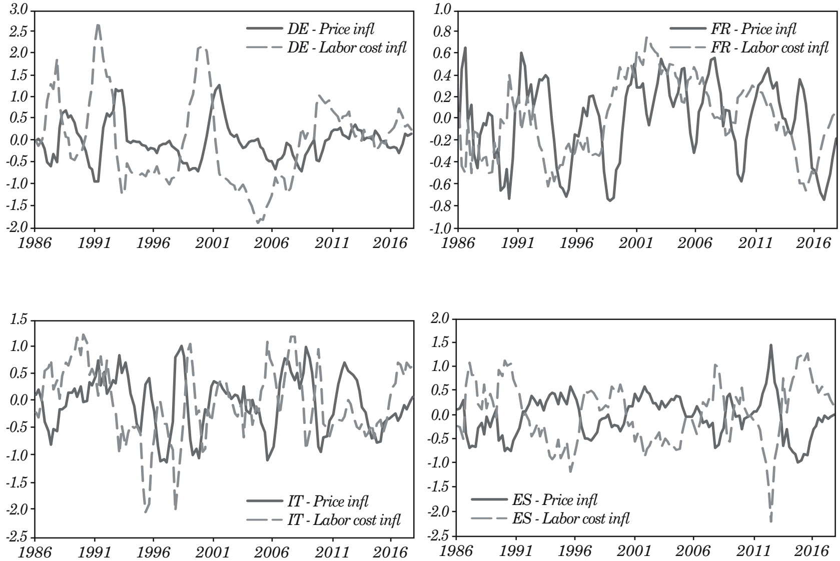

Source: Authors' calculations.
Sample period: 1985Q1–2018Q1.

Appendix N

The Derivation of the Counterfactual IRFs

Consider the following VAR(1)[59]:

$$A_0Y_t = A_1Y_{t-1} + \in_t \qquad \forall t = 1,\ldots,T \tag{1}$$

where Y_t is the vector of endogenous variables, A_0, A_1 the matrices of contemporaneous and lag coefficients, respectively and $\in_t$ are structural shocks.

$$Y_t = (A_0)^{-1}A_1Y_{t-1} + (A_0)^{-1}\in_t \tag{2}$$

$$Y_t = BY_{t-1} + (A_0)^{-1}\in_t,\ B = (A_0)^{-1}A_1 \tag{3}$$

A simple way to calculate IRFs is to iterate starting with $t = 0$.

$$Y_0 = (A_0)^{-1}\in_0 \tag{4}$$

$$Y_1 = B\cdot(A_0)^{-1}\in_0 + (A_0)^{-1}\in_1 \tag{5}$$

$$Y_k = B^k\cdot(A_0)^{-1}\in_0 + B^{k-1}\cdot(A_0)^{-1}\in_1 + \ldots + B\cdot(A_0)^{-1}\in_{k-1} + (A_0)^{-1}\in_k \tag{6}$$

$$Y_k = \sum\nolimits_{h=0}^{k} B^{k-h}\cdot(A_0)^{-1}\in_h \tag{7}$$

The IRF of variable i following a certain shock j at period $h(IRF_{ij}^h)$ is achieved by setting $\in_0 = e_j$, where e_j is an identification column vector with 1 on the j-th position and zero otherwise.

We choose variable i^* for which the counterfactual responses to shock j are set to zero.

In order to offset the IRF of variable i^* to shock j, we produce a set of counterfactual shocks (ε_t). We set:

$$\in_0 = e_j + \widehat{\in_0}\cdot\sum\nolimits_{l=1}^{n} e_l \tag{8}$$

where e_j is a column vector with 1 on the l position and zero otherwise and n is the number of structural shocks.

59. Lag 1 was selected for illustration purposes, the formulas derived for the counterfactual IRFs also hold in the general VAR(p) case.

At this point we depart from similar approaches. For example, if the VAR is identified using a Choleski framework, the impact of shock j on variable i^* is offset by modifying only the shock corresponding to variable i^* in the recursive identification scheme. As in the sign restriction framework each identified structural shock can impact instantaneously all endogenous variables, in deriving the counterfactual IRFs we assume that all the structural shocks contribute to the offset. One assumption we make in order to ensure determinacy is that the shocks have an equal contribution in offsetting the impact of shock j on variable i^*.[60]

$$\epsilon_1 = \widehat{\epsilon_1} \cdot \textstyle\sum_{l=1}^{n} e_l \tag{9}$$

$$\epsilon_k = \widehat{\epsilon_k} \cdot \textstyle\sum_{l=1}^{n} e_l \, . \tag{10}$$

We determine $\widehat{\epsilon_0}, \widehat{\epsilon_1}, \ldots, \widehat{\epsilon_k}$ such that $\widehat{IRF}_{i^*j}^{h} = 0$ for all periods $h = 0,1,\ldots,k$, where i^* is the variable whose IRF is being shut down.

$$IRF_{ij}^{h} = e_i^{'} \cdot Y_h \tag{11}$$

$$Y_0 = (A_0)^{-1} e_j + \widehat{\epsilon_0} \cdot (A_0)^{-1} \cdot \textstyle\sum_{l=1}^{n} e_l \, . \tag{12}$$

The counterfactual IRF of variable i to shock j at the moment 0 is $\widehat{IRF}_{ij}^{0}$:

$$\widehat{IRF}_{ij}^{0} = e_i^{'} \cdot Y_0 = e_i^{'} (A_0)^{-1} e_j + \widehat{\epsilon_0} \; e_i^{'} \cdot (A_0)^{-1} \cdot \textstyle\sum_{l=1}^{n} e_l \tag{13}$$

but $e_i^{'} \cdot (A_0)^{-1} e_l = IRF_{il}^{0}$, therefore:

$$\widehat{IRF}_{ij}^{0} = IRF_{ij}^{0} + \widehat{\epsilon_0} \cdot \textstyle\sum_{l=1}^{n} IRF_{il}^{0} \, . \tag{14}$$

Notation: $IRF_{i\Sigma}^{h} = \sum_{l=1}^{n} IRF_{il}^{h}$ (the sum for the period h of all IRFs of variable i to all other shocks).

60. In this approach the combination of structural shocks that is constructed to offset the response of variable i^* to structural shock j also impacts instantaneously all the other variables. This is consistent with assuming the existence of instantaneous effects, but it may be argued that this instantaneous impact contribute to the difference between the unrestricted and the counterfactual IRFs. We checked therefore an alternative way of constructing the counterfactual IRFs, in which each structural shock can have a different contribution to the offsetting (relaxing the equal weights assumption). The resulting system is identified assuming that the counterfactual shock impacts instantaneously only variable i^*. The results are qualitatively similar as in the baseline approach.

$$\widehat{IRF}^{\,0}_{i^*j} = IRF^{0}_{i^*j} + \widehat{\in_0} \cdot IRF^{0}_{i^*\Sigma} = 0 \tag{15}$$

$$\widehat{\in_0} = -\frac{IRF^{0}_{i^*j}}{IRF^{0}_{i^*\Sigma}} \tag{16}$$

$$Y_1 = B \cdot (A_0)^{-1} e_j + \widehat{\in_0} \cdot B \cdot (A_0)^{-1} \cdot \textstyle\sum_{l=1}^{n} e_l + \widehat{\in_1} \cdot (A_0)^{-1} \cdot \sum_{l=1}^{n} e_l \tag{17}$$

$$\begin{aligned}\widehat{IRF}^{\,1}_{ij} = e_i' \cdot Y_1 = e_i' \cdot B \cdot (A_0)^{-1} e_j + \widehat{\in_0} \cdot e_i' \cdot B \cdot (A_0)^{-1} \\ \cdot \textstyle\sum_{l=1}^{n} e_l + \widehat{\in_1} \cdot e_i' \cdot (A_0)^{-1} \cdot \sum_{l=1}^{n} e_l\end{aligned} \tag{18}$$

$$\widehat{IRF}^{\,1}_{i^*j} = IRF^{1}_{i^*j} + \widehat{\in_0} \cdot \textstyle\sum_{l=1}^{n} IRF^{1}_{i^*l} + \widehat{\in_1} \cdot \sum_{l=1}^{n} IRF^{0}_{i^*l} = 0 \tag{19}$$

$$\widehat{\in_1} = -\frac{IRF^{1}_{i^*j} + \widehat{\in_0} \cdot IRF^{1}_{i^*\Sigma}}{IRF^{0}_{i^*\Sigma}}\,. \tag{20}$$

In general:

$$\begin{aligned}\widehat{IRF}^{\,k}_{i^*j} = IRF^{k}_{i^*j} + \widehat{\in_0} \cdot IRF^{k}_{i\Sigma} + \widehat{\in_1} \cdot IRF^{k-1}_{i\Sigma} + \ldots \\ + \widehat{\in_{k-1}} \cdot IRF^{1}_{i\Sigma} + \widehat{\in_k} \cdot IRF^{0}_{i\Sigma} = 0\end{aligned} \tag{21}$$

$$\widehat{\in_k} = -\frac{IRF^{k}_{i^*j} + \sum_{h=0}^{k-1} \widehat{\in_h} \cdot IRF^{k-h}_{i\Sigma}}{IRF^{0}_{ij}}\,. \tag{22}$$

As shown in equation 21 for the case of $i = i^*$, for a given variable i the counterfactual IRF[61] is the following:

$$\widehat{IRF}^{\,k}_{ij} = IRF^{k}_{ij} + \textstyle\sum_{h=0}^{k} \widehat{\in_h} \cdot IRF^{k-h}_{i\Sigma}. \tag{23}$$

61. The derivation of counterfactual IRFs follows the same principles when setting the difference between the IRF of wages and of productivity to zero after a structural shock j. Additionally, we assume that the two IRFs contribute equally to setting this difference to zero and the weight of shock j is allowed to vary from that of other shocks contributing to the offsetting.

Appendix O

Unrestricted and Counterfactual IRFs in the Two-Shock VAR

Figure O1. Impulse response functions for the total economy

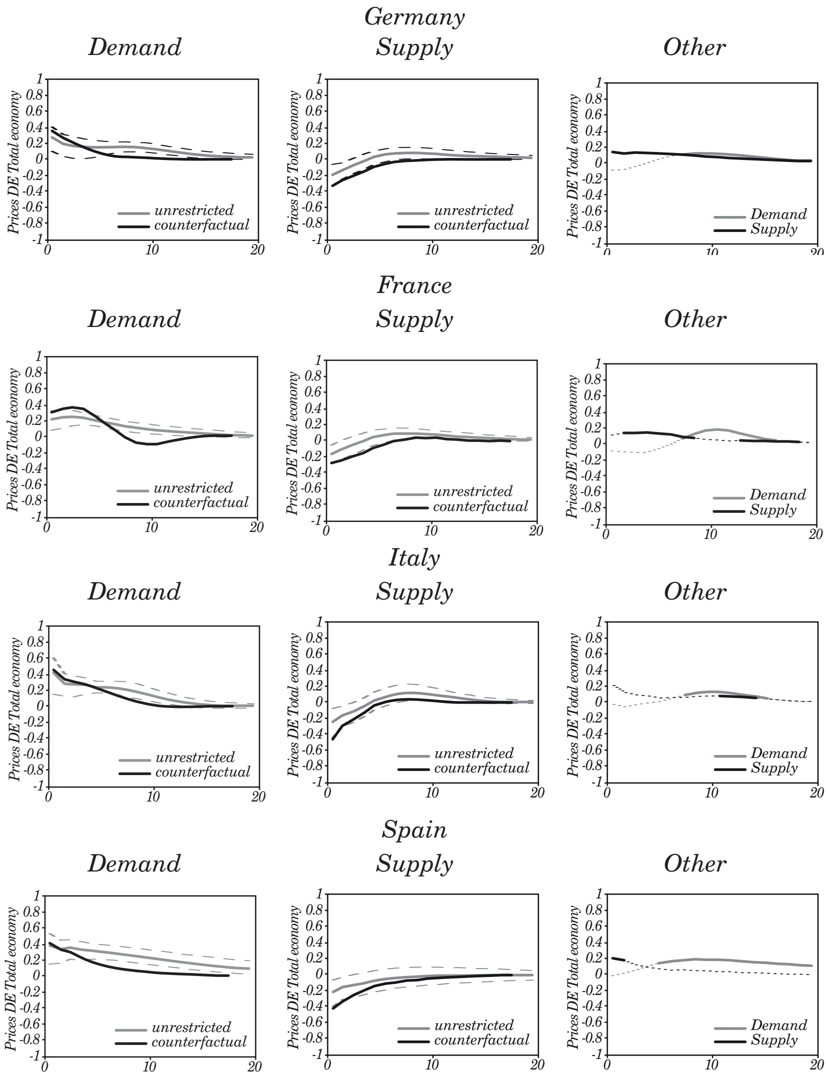

Source: Authors' calculations.

Appendix P

The Response of Margins after a Demand and after a Supply Shock

Figure P1. The cumulated response of margins for the total economy

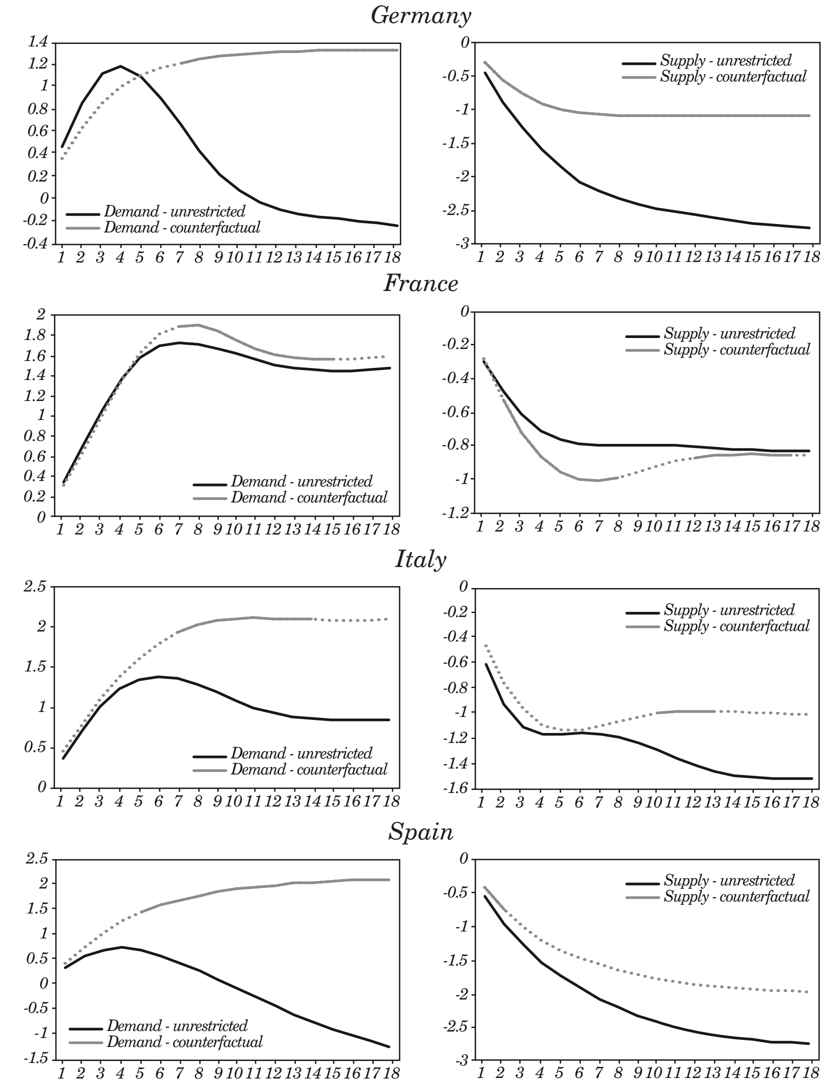

Source: Authors' calculations.
Note: The solid gray lines mark the instances where labor costs act as an amplification channel for price inflation in a significant manner.

Appendix Q

Unrestricted and Counterfactual IRFs in the Four-Shock VAR

Figure Q1. Impulse response functions for the total economy

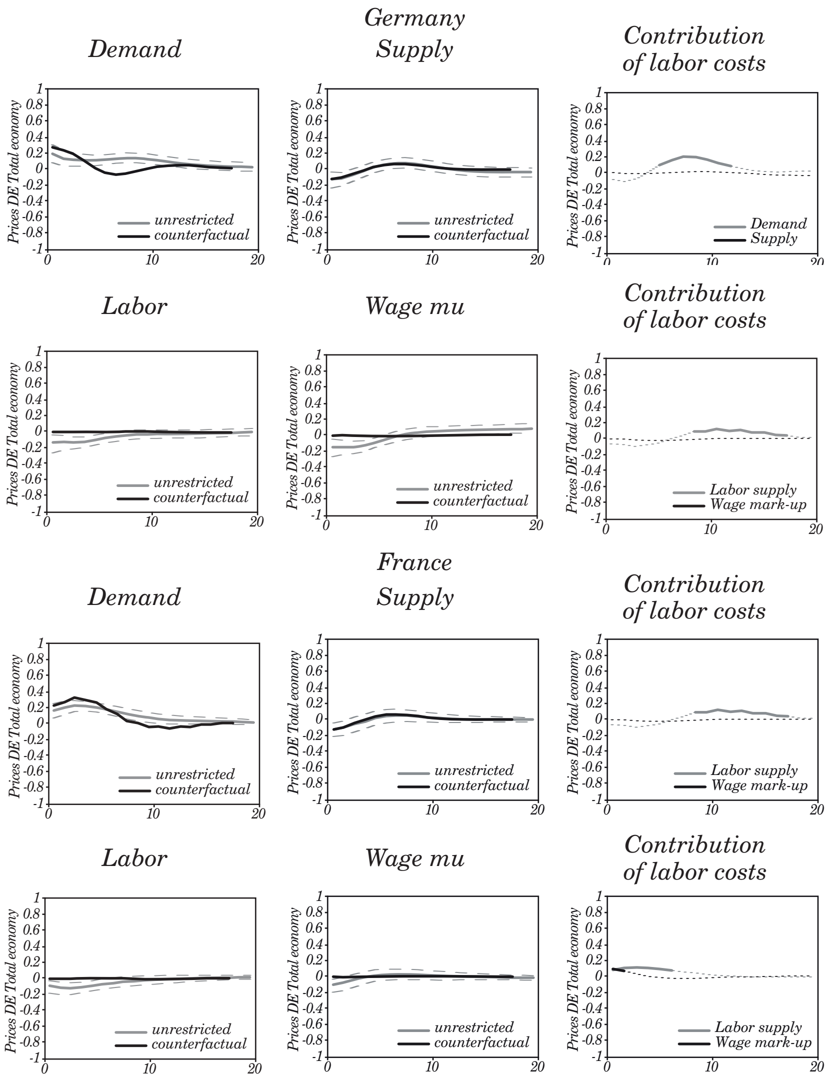

Figure Q2. Impulse response functions for the total economy

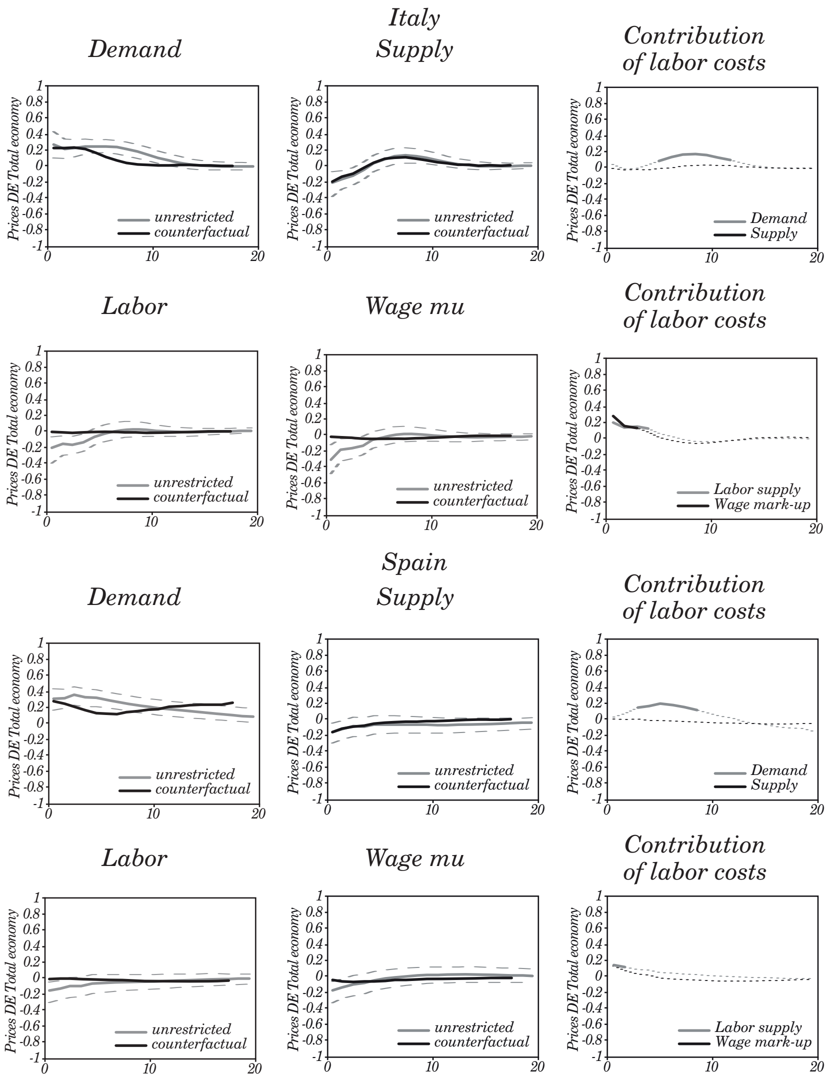

Source: Authors' calculations.

Appendix R

Results Based on a Five-Shock VAR

Table R1. The Five-Shock VAR: Identification Scheme

Variables	*Shocks*					
	Demand	*Supply*	*Labor supply*	*Wage mark-up*	*Monetary pol*	*Other*
Real value added	+	+	+	+	+	•
Prices	+	-	-	-	+	•
Wages	+	+	-	-	+	•
Productivity	+	+	•	•	•	•
Unemployment rate	-	•	+	-	-	•
Spread	+	•	•	•	-	•

Source: Authors' calculations.
Notes: • = unconstrained, + = positive sign, - = negative sign

Figure R1. Amplification of price inflation response due to the labor cost channel in the five-shock

VAR

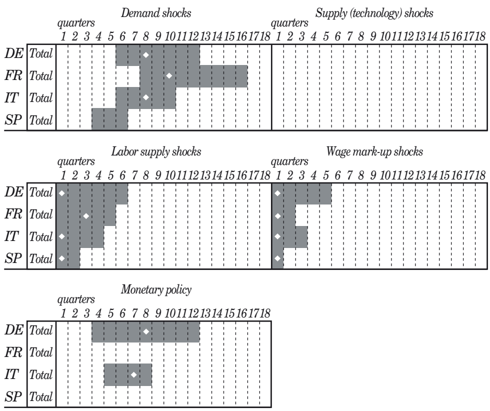

Source: Authors' calculations.
Note: This chart indicates, in blue, the quarters following a certain shock where the median counterfactual IRF lies outside the 68 percent posterior uncertainty band of the unrestricted IRF.

Has the U.S. Wage Phillips Curve Flattened?A Semi-Structural Exploration

Jordi Galí
CREI, Universitat Pompeu Fabra, and Barcelona GSE

Luca Gambetti
Universitat Autònoma de Barcelona, Barcelona GSE, Collegio Carlo Alberto, and Università di Torino

The deep and prolonged recession triggered by the global financial crisis of 2007–2009 led to a large increase in the unemployment rate in most advanced economies. Ten years later, at the time of writing this paper, the recession has long ended, and the subsequent recoveries have brought the unemployment rate to levels close to, and in some cases even below, those at the peak of the previous expansion. In the U.S., the unemployment rate increased from 4.4 percent in May 2007 to 10 percent in November 2009. Since that peak was attained, the unemployment rate has decreased, albeit at a slower pace than in earlier recoveries, down to its current level below 4 percent. Both movements represent, respectively, the largest increase and the largest decrease in the unemployment rate experienced by the U.S. economy during the postwar period. Despite those wide and persistent fluctuations in unemployment, inflation has remained surprisingly stable during the same period, as illustrated in figure 1. The previous phenomenon, often referred to in the literature as the "twin puzzle,"

Prepared for the XXII Conference of the Central Bank of Chile, Santiago, October 25–26. We are grateful for comments to Fernanda Nechio, Régis Barnichon, Geert Mesters as well as participants at the CREI Faculty lunch and the Central Bank of Chile conference.

Changing Inflation Dynamics, Evolving Monetary Policy edited by Gonzalo Castex, Jordi Galí, and Diego Saravia, Santiago, Chile.

appears to be robust to the measure of inflation and economic slack used,[1] and has also been observed in other advanced economies.[2]

Not surprisingly, central banks around the world have sounded the alarm in the face of that development, and with good reason.[3] For one, a flattening of the Phillips curve implies a larger sacrifice ratio and the need for more extreme policy measures in order to eliminate deviations of inflation from target. Furthermore, an outright decoupling of inflation from indicators of economic slack would call into question the inflation targeting framework widely adopted by central banks over the past decades, since that framework hinges critically on the existence of a positive relation between inflation and the level of economic activity. This is because it is only through its ability to influence the latter through an appropriate setting of the interest rate and other policy instruments that central banks can aim at controlling inflation.

In the present paper we revisit a key link of the relation between prices and economic activity, namely, the relation between wage inflation and unemployment. This empirical relation, which was the focus of Phillips's original work (Phillips (1958)), is widely perceived to be at the heart of the “twin puzzle.” In particular, the failure of wage inflation to respond sufficiently to the tightening of the labour market in recent years is generally viewed as one of the main factors behind the extremely accommodating monetary policies at central banks, like the Federal Reserve or the ECB. Figure 2 shows a scatterplot of wage inflation and the unemployment rate to illustrate that phenomenon. We use quarterly data for the period 1964Q1–2017Q4.[4] As discussed in Galí (2011a), the absence of a clear inverse relation between the two variables over the full sample period can be attributed to the large changes in mean price inflation experienced by the U.S. economy in the ’70s (“the Great Inflation”) and in the early ’80s (the Volcker disinflation). When we restrict ourselves to the ‘60s and the Great Moderation period, a clear negative relation between the two variables becomes noticeable. Interestingly, that relation can be seen to become nearly flat in the years of the financial crisis and the subsequent recovery, in a way analogous to the "twin puzzle" phenomenon for price inflation.

1. See Stock and Watson (2018), Coibion and Gorodnichenko (2015), and Daly and Hobijn (2014) for U.S. evidence and possible interpretations

2. See Ciccarelli and Osbat (2017) for euro-area evidence.

3. See Constancio (2017).

4. The wage inflation measure is the (centered) year-on-year change in the (log) average hourly earnings of production and nonsupervisory workers in the private sector, from the Establishment Survey.

Figure 1. Unemployment and Price Inflation

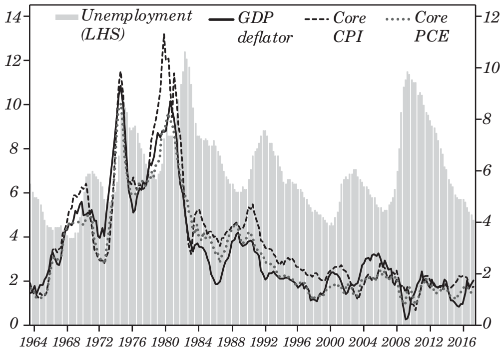

Source: Authors' calculations.

Figure 2. The U.S. Wage Phillips Curve

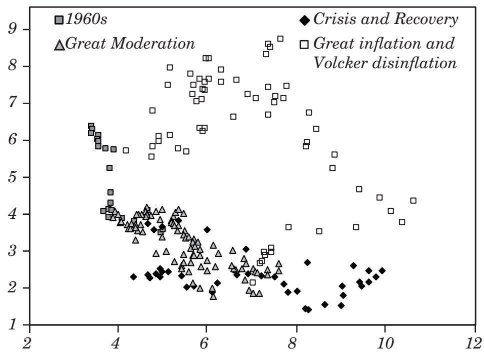

Source: Authors' calculations.

In the present paper, we seek to accomplish two tasks. First we document changes in the U.S. wage Phillips curve by using simple reduced-form regressions. Secondly, and after discussing the limitations of such a reduced-form approach, we report estimates of a *conditional* wage Phillips curve, based on a structural decomposition of wage, price, and unemployment data generated by a VAR with time-varying coefficients and identified by a combination of long-run and sign restrictions.

Our main findings can be summarized as follows. First, our reduced-form estimates point to a substantial decline in the estimated coefficients on both lagged inflation and unemployment in the U.S. wage Phillips curve. Secondly, our estimates of conditional wage Phillips curves display similar qualitative results, thus suggesting that the reduced form evidence is not driven by endogeneity problems or possible changes in the relative importance of wage-markup shocks. The conditional evidence, however, suggests that actual changes in the slope of the wage Phillips curve may not have been as large as implied by the unconditional estimates. Finally, we show that the reduced sensitivity of wage inflation to unemployment is also reflected in the estimated changes in a dynamic-multiplier statistic relating the time-varying joint responses of wage inflation and unemployment to both demand and monetary-policy shocks.

The paper is organized as follows. Section 1 describes the reduced-form estimates of a wage Phillips curve and discusses some of its limitations. Section 2 presents our structural vector autoregressive (VAR) model and describes the "semi-structural" evidence based on it. Section 3 concludes.

1. Wage Inflation and Unemployment: Reduced-form Evidence

In the present section we provide some reduced-form evidence on the changing relation between wage inflation and the unemployment rate in the U.S. economy. The starting point of our empirical analysis is the estimation of a baseline wage inflation Phillips curve given by:

$$\pi_t^w = \alpha + \gamma\overline{\pi}_{t-1}^p + \varphi u_t + \varepsilon_t \tag{1}$$

where $\pi_t^w \equiv 400(w_t - w_{t-1})$ is (annualized) quarterly wage inflation (with w_t denoting the log nominal wage), $\overline{\pi}_t^p$ is a measure of price inflation (also annualized), and u_t is the unemployment rate. A specification similar to (1) has often been proposed and used in empirical applications.[5]

Table 1A reports the OLS estimates of the coefficients on lagged inflation and the unemployment rate in (1) for different sample periods, using our baseline specification. We use average hourly earnings of

5. See Blanchard and Katz (1999).

production and nonsupervisory workers in the private sector from the Establishment Survey to construct our wage inflation measure. Our baseline price inflation measure is the (annualized) quarterly rate of change in the GDP deflator. The civilian unemployment rate is our measure of unemployment. We take 2017Q2, which corresponds to the trough in the unemployment rate, to be the cutoff date splitting the two sample periods.

Several observations stand out. First, for the sample period before the financial crisis, the coefficients on both lagged inflation and unemployment are highly significant and with the expected sign. In particular, the estimated slope coefficient suggests that an increase of 1 percentage point in the unemployment rate is associated with a reduction of (annualized) wage inflation of about 30 basis points, given price inflation.

We uncover substantial changes *within* the precrisis sample period, however. Thus, in the 1986Q1-2007Q2 subsample (roughly corresponding to the Great Moderation), the estimated inflation coefficient becomes much smaller (though still significant), while the negative effect of unemployment on wage inflation is estimated to be about twice as large (and significantly different from the pre-86 period, as reflected in the p-value reported).

Table 1A. Empirical Wage Phillips Curves

	Earnings, GDP deflator (Q)		
	π^p_{t-1}	u_t	R^2
Precrisis			
1964Q1-2007Q2	0.67	-0.29	0.57
	(0.04)	(0.07)	
1986Q1-2007Q2	0.17	-0.58	0.45
	(0.07)	(0.07)	
p-value	0.003	0.016	
Crisis and recovery			
2007Q3-2017Q4	0.01	-0.11	0.10
	(0.12)	(0.04)	
p-value 1	0.002	0.102	
p-value 2	0.056	0.001	

Source: Authors' calculations.

Our estimates for the period of crisis and the subsequent recovery, shown in the bottom panel of the table, point to a large decline in the sensitivity of wage inflation to the unemployment rate, though the coefficient on the latter still remains significant. While the test of equality of that coefficient between the full precrisis and the post-crisis samples can only reject that hypothesis at the 10 percent level (see p-value 1), the equality with the "late" precrisis period is rejected with very low p-values (see p-value 2). In addition, it is worth noting that the estimated coefficient on lagged inflation becomes insignificant in the more recent subsample period, thus suggesting a further decline in the importance of price indexation in wage setting over the past decade.[6]

Figure 3A. Unemployment Coefficient
(unconditional estimate)

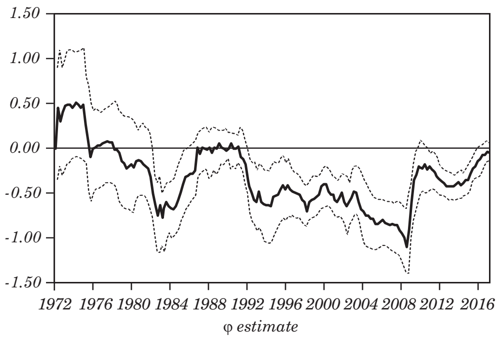

Figure 3B. Inflation Coefficient
(unconditional estimate)

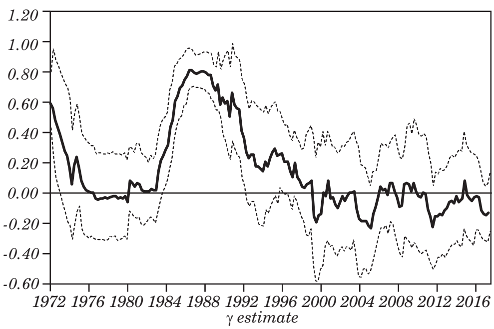

Source: Authors' calculations.

6. The finding of a reduced importance of lagged inflation since the 1980s is consistent with Blanchard (2016)'s estimates of price inflation Phillips curves.

Figures 3A and 3B show the time-varying estimates of γ and, respectively, based on a rolling OLS regression with a 32-observation window. The estimates illustrate, in a flexible way, the evidence reported earlier, namely, the consecutive steepening (during the Great Moderation) and flattening (during the financial crisis and its aftermath) of the wage Phillips curve in the U.S., as well as the seeming irrelevance of lagged price inflation since the mid-1980s.

The qualitative findings discussed above are, for the most part, robust to alternative specifications of equation (1), as shown in tables 1B (using CPI inflation), 1C (using lagged unemployment), and 1D (using year-on-year lagged price inflation).

Table 1B. Empirical Wage Phillips Curves

	Earnings, CPI (Q)		
	π_{t-1}^{p}	u_t	R^2
Precrisis			
1964Q1-2007Q2	0.45	-0.16	0.42
	(0.04)	(0.08)	
1986Q1-2007Q2	0.09	-0.57	0.45
	(0.04)	(0.07)	
p-value	0.001	0.021	
Crisis and recovery			
2007Q3-2017Q4	0.05	-0.11	0.14
	(0.04)	(0.05)	
p-value 1	0.001	0.56	
p-value 2	0.243	0.001	

Source: Authors' calculations.

Table 1C. Empirical Wage Phillips Curves

Earnings, GDP deflator (Q), lagged unemployment			
	π^p_{t-1}	u_{t-1}	R^2
Precrisis			
1964Q1-2007Q2	0.65	-0.22	0.55
	(0.04)	(0.07)	
1986Q1-2007Q2	0.16	-0.56	0.43
	(0.07)	(0.07)	
p-value	0.014	0.367	
Crisis and recovery			
2007Q3-2017Q4	0.01	-0.15	0.18
	(0.11)	(0.05)	
p-value 1	0.002	0.34	
p-value 2	0.020	0.001	

Source: Authors' calculations.

Table 1D. Empirical Wage Phillips Curves

Earnings, GDP deflator (YOY)			
	π^p_{t-1}	u_t	R^2
Precrisis			
1964Q1-2007Q2	0.80	-0.49	0.63
	(0.04)	(0.06)	
1986Q1-2007Q2	0.25	-0.61	0.47
	(0.08)	(0.07)	
p-value	0.045	0.028	
Crisis and recovery			
2007Q3-2017Q4	0.03	-0.11	0.10
	(0.19)	(0.05)	
p-value 1	0.015	0.001	
p-value 2	0.045	0.001	

Source: Authors' calculations.

1.1 Shortcomings

The evidence reported above, based on OLS estimates of equation (1), has several shortcomings as a measure of the sensitivity of wage inflation to variations in unemployment. Firstly, the specification of (1), while frequently found in textbooks and empirical applications, is generally viewed as being largely *ad-hoc*. In appendix A.1 we provide some possible microfoundations for such a specification based on the staggered wage-setting model of Erceg and others (2000) augmented with partial indexation to (lagged) price inflation and reformulated in terms of unemployment, as in Galí (2011a). The resulting microfounded wage equation takes the form:

$$\pi_t^w = \alpha + \gamma\overline{\pi}_{t-1}^p + \varphi\hat{u}_t + \psi\hat{\mu}_{w,t}^n$$

where α, γ, φ, and ψ are functions of structural parameters, and $\hat{\mu}_{w,t}^n$ represents exogenous fluctuations in the *natural* wage markup.[7]

Under the previous microfounded interpretation, the assumption of orthogonality between the right-hand side variables and the disturbance term in (1) that would justify the use of OLS is unlikely to be satisfied in practice. Estimated dynamic stochastic general equilibrium (DSGE) models suggest that natural wage-markup shocks are far from negligible sources of macro fluctuations. In particular they have significant effects on both price inflation and unemployment.[8] The latter observation suggests that reduced-form OLS estimates of γ and φ would likely be inconsistent. Furthermore, changes over time in the volatility of wage-markup shocks could be a source of spurious changes in the OLS estimates of those coefficients across subsample periods, thus giving a misleading impression of a "structural change" in the response of wage inflation to unemployment.

Below we propose and implement an empirical framework that aims at assessing possible changes in the responsiveness of wage inflation to unemployment in a way that overcomes (at least in principle) some of the previous limitations of unconditional reduced-form estimates of the wage Phillips curve.

7. The natural wage markup is the gap between the average real wage and the marginal rate of substitution that would prevail under flexible wages.

8. See Smets and Wouters (2007), and Galí and others (2012).

2. Wage Inflation and Unemployment: Semi-Structural Evidence

In the present section we describe our empirical approach to identifying the different components of unemployment, wage inflation, and price inflation, based on a structural vector autoregressive model with time-varying coefficients (TVC-SVAR). Our empirical model provides a flexible specification which allows for structural changes in the relation between wage inflation and unemployment, as well as other structural changes that the U.S. economy may have experienced over the sample period considered.[9] In addition, our framework makes it possible to overcome the potential endogeneity problem discussed above. More generally, our approach allows us to estimate the joint dynamics of wage inflation and unemployment in response to monetary-policy interventions, and to uncover any changes over time in those dynamics.

2.1 Empirical Model

Let $x_t = [\Delta(y_t - n_t), \pi_t^w, \pi_t^p, u_t, i_t^L]$ where y_t is (log) GDP, n_t denotes (log) hours of all persons in the nonfarm business sector, and i_t^L is the yield on 10-year government bonds. Price inflation and wage inflation, π_t^w and π_t^p, are now defined as quarterly (log) first-differences of wage earnings of production and nonsupervisory workers and of the GDP deflator, respectively, i.e. $\pi_t^w \equiv w_t - w_{t-1}$ and $\pi_t^p \equiv p_t - p_{t-1}$. As above, u_t denotes the civilian unemployment rate. We use a long-term yield to avoid problems related to the binding zero lower bound at the end of our sample. All data are quarterly. The sample period is 1964Q1-2017Q4.

We assume the existence of a TVC-VAR representation

$$x_t = \mathbf{A}_{0,t} + \mathbf{A}_{1,t}\mathrm{x}_{t-1} + \mathbf{A}_{2,t}\mathrm{x}_{t-2} + \ldots + \mathbf{A}_{p,t}\mathrm{x}_{t-p} + \mathbf{u}_t \qquad (2)$$

where $\mathbf{A}_{0,t}$ is a vector of time-varying intercepts, $\mathbf{A}_{i,t}$, for $i = 1,\ldots,p$ are matrices of time-varying coefficients, and $\mathbf{u}_t$ is a Gaussian white-noise vector process with time-varying covariance matrix Σ_t. We assume

9. These may include the change in the cyclical behaviour of productivity emphasized in Galí and Gambetti (2009), or the change in monetary policy starting with Paul Volcker's tenure at the Fed; see Clarida and others (2000).

the reduced-form innovations $\mathbf{u}_t$ are a (possibly time-varying) linear transformation of the underlying structural shocks ε_t given by

$$\mathbf{u}_t \equiv \mathbf{Q}_t \varepsilon_t \tag{3}$$

where $\mathbb{E}\left\{\varepsilon_t \varepsilon_t'\right\} = I$ and $\mathbb{E}\left\{\varepsilon_t \varepsilon_{t-k}'\right\} = 0$ for all t and $k = 1, 2, 3, \ldots$. It follows that $\mathbf{Q}_t \mathbf{Q'}_t = \Sigma_t$. As described in the appendix, our approach assumes all the time-varying coefficients follow random walks with independent innovations.

Estimation is carried out as in Del Negro and Primiceri (2013).[10] Estimates of (2) can be used to obtain the (local) reduced-form moving average (MA) representation:

$$x_t = \mu_t + \mathbf{B}_t(L)\, \mathbf{u}_t.$$

Equation (3) can then be used to recover the structural (local) TVC-MA representation:

$$x_t = \mu_t + \mathbf{C}_t(L)\, \mathbf{u}_t.$$

where $\mathbf{C}_t\,(L) = \mathbf{B}_t\,(L)\mathbf{Q}_t$, and where $x_t^{ij} = \mathbf{C}_t^{ij}\,(L)\varepsilon_t^j$ represents the component of the ith variable associated with the jth shock. Determination of $\mathbf{Q}_t$ requires a set of assumptions to identify the different shocks (i.e. the different elements of ε_t) driving fluctuations in x_t.

We identify technology shocks, following Galí (1999), as the only shocks in vector ε_t to have a long-run effect on labour productivity, implemented by imposing $\mathbf{C}_t^{1j}(1) = 0$ for all t and $j = 2, 3, 4, 5$. In addition to a technology shock, we assume the existence of four additional shocks typically found in estimated DSGE models: (non-monetary) demand shocks, monetary-policy shocks, price-markup shocks, and wage-markup shocks. By construction, those four shocks are restricted to have only transitory effects on labour productivity. We disentangle them through restrictions on the sign of their implied comovements between certain variables over a four-quarter horizon after each shock.[11] Our sign restrictions are motivated by the predictions of the

10. We refer the reader to Galí and Gambetti (2014) for details.

11. The use of sign restrictions for identification purposes in structural VARs was pioneered by Uhlig (2005).

estimated medium-scale New Keynesian model with unemployment in Galí and others (2012).[12] Here are our short-run sign restrictions:

• Demand shocks (to be understood as non-monetary) are assumed to generate a positive comovement among y_t, π_t^p, and i_t^L.

• Monetary-policy shocks imply a positive comovement between y_t and π_t^p, but a negative comovement between each of those variables and i_t^L.

• Price-markup shocks are identified as the only source of fluctuations that generates a positive comovement between π_t^p and the price markup $\mu_t^p \equiv (y_t - n_t) - (w_t - p_t)$.

• Wage-markup shocks are assumed to be the only structural disturbances that generate a positive comovement between π_t^w and the unemployment rate u_t, with the latter variable interpreted as a proxy for the wage markup, following Galí (2011a, 2011b).

Table 2 summarizes our identification strategy in a compact way.

Table 2. Identification

	$y_t - n_t$	π_t^w	u_t	π_t^p	μ_t^p	i_t^L
Technology						
Demand	0∞	+	-	+	-	+
Monetary policy	0∞	+	-	+	-	-
Price markup	0∞	+/-	-/+	+	+	
Wage markup	0∞	+	+	+/-	-/+	

Source: Authors' calculations.

12. That model is itself an extension of those in Smets and Wouters (2007) that introduces an explicit relation between the unemployment rate and the wage markup discussed in Galí (2011b).

2.2 Conditional Wage Phillips Curves

The next step in our approach consists in re-estimating the wage Phillips curve (1) using the time series for wage inflation, price inflation and unemployment, *purged of the component associated with wage-markup shocks* obtained using the TVC-SVAR described above. To the extent that the error term in (1) captures fluctuations in wage-markup shocks, the estimation of such a *conditional* wage Phillips curve should overcome any bias resulting from the correlation between the error term and the regressors.[13]

Table 3A reports estimates of coefficients γ and φ in (1) for different sample periods, using the non-wage-markup components of the three variables involved. As in our TVC-SVAR specification, price inflation is measured as the log first-difference of the GDP deflator. As in our baseline estimates of (1), we annualize both inflation variables before applying OLS. The estimates for the full precrisis period point to a smaller price inflation coefficient (0.32) and a larger (in absolute terms) unemployment coefficient (0.55) than the "unconditional" estimates of Table 1A. Interestingly, when we now restrict ourselves to the Great Moderation period, we still get a smaller inflation coefficient and larger unemployment coefficient than in the full precrisis period, but now the differences in the estimated coefficients are much smaller (and in the case of the inflation coefficient, statistically insignificant). When we turn to the crisis and recovery period, we obtain estimates of the inflation and unemployment coefficients that are smaller (in absolute value) than in the precrisis period. As in our unconditional estimates of table 1A for this period, the coefficient on lagged inflation is now insignificant. The estimated coefficient on unemployment is significantly smaller (in absolute value) than in the precrisis period, though more than twice as large as its unconditional counterpart.

Figures 4A and 4B report time-varying estimates of φ and γ respectively, based on rolling OLS regressions with a 32-observation window, applied to the non-wage-markup components of the time series involved. Note that, relative to figures 3A and 3B, and consistent with the evidence just discussed, the coefficient on unemployment appears to be more stable over time, and to experience a smaller decline in recent years. On the other hand, the conditional rolling estimates

13. Our approach is similar in spirit to that of Barnichon and Mesters (2018a), who estimate a New Keynesian Phillips curve for price inflation by using current and lagged monetary-policy shocks as instruments.

of the inflation coefficient display a pattern which is very similar to their unconditional counterparts, though with slightly lower values in the 1980s.

The previous qualitative findings are largely robust to alternative specifications of conditional wage Phillips curves, as shown in tables 3B (using lagged unemployment) and 3C (using year-on-year lagged price inflation). An exception to this similarity is given by the estimates of the coefficient on lagged year-on-year inflation in table 3C, which do not appear to vary significantly across sample periods.

Figure 4A. Unemployment Coefficient (conditional estimate)

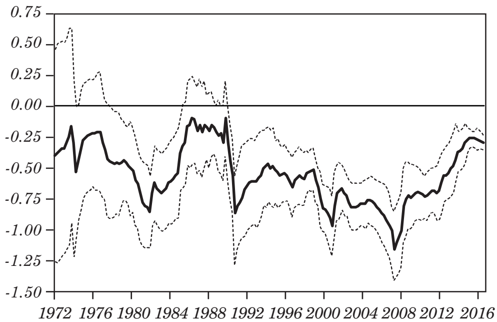

Figure 4B. Inflation Coefficient (conditional estimate)

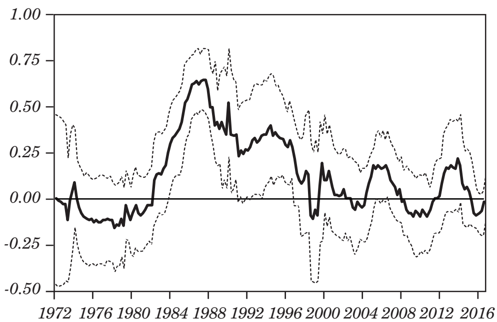

Source: Authors' calculations.

Table 3A. Conditional Wage Phillips Curves

	Earnings, GDP deflator (YOY)		
	π^{p}_{t-1}	u_t	R^2
Precrisis			
1964Q1-2007Q2	0.32	-0.55	0.39
	(0.04)	(0.07)	
1986Q1-2007Q2	0.21	-0.74	0.64
	(0.05)	(0.06)	
p-value	0.226	0.019	
Crisis and recovery			
2007Q3-2017Q4	0.07	-0.29	0.10
	(0.09)	(0.05)	
p-value 1	0.24	0.04	
p-value 2	0.36	0.001	

Source: Authors' calculations.

Table 3B. Conditional Wage Phillips Curves

	Earnings, GDP deflator (Q), lagged unemployment		
	π^{p}_{t-1}	u_{t-1}	R^2
Precrisis			
1964Q1-2007Q2	0.31	-0.40	0.31
	(0.05)	(0.00)	
1986Q1-2007Q2	0.20	-0.70	0.44
	(0.06)	(0.07)	
p-value	0.199	0.009	
Crisis and recovery			
2007Q3-2017Q4	0.07	-0.32	0.56
	(0.08)	(0.05)	
p-value 1	0.28	0.46	
p-value 2	0.34	0.001	

Source: Authors' calculations.

Table 3C. Empirical Wage Phillips Curves

	Earnings, GDP deflator (YOY)		
	π^p_{t-1}	u_t	R^2
Precrisis			
1964Q1-2007Q2	0.43	-0.65	0.47
	(0.05)	(0.07)	
1986Q1-2007Q2	0.33	-0.76	0.68
	(0.06)	(0.06)	
p-value	0.420	0.026	
Crisis and recovery			
2007Q3-2017Q4	0.32	-0.23	0.57
	(0.12)	(0.05)	
p-value 1	0.665	0.001	
p-value 2	0.96	0.001	

Source: Authors' calculations.

2.3 Conditional Dynamic Multipliers

The empirical approach described in the previous subsection should in principle overcome one of the shortcomings of the reduced-form evidence, namely, the potential biases in the OLS estimates of equation (1) resulting from the endogeneity of unemployment and inflation with respect to wage-markup shocks. Yet, the estimates of conditional wage Phillips curves are still subject to another important caveat, namely, the *ad-hoc* specification of (1). In the present subsection we uncover possible changes over time in the relation between unemployment and wage inflation without the straitjacket of any assumed functional relation between the two variables. Instead, we focus on the estimated impulse responses generated by our TVC-SVAR and trace the evolution over time of the dynamic wage inflation–unemployment multiplier, defined as a ratio of the cumulative impulse responses of those two variables to a given shock ε^i_t at different horizons:

$$\Phi^i_t(k) \equiv \frac{\sum_{k=0}^{K} \frac{\partial \pi^w_{t+k}}{\partial \varepsilon^i_t}}{\sum_{k=0}^{K} \frac{\partial u_{t+k}}{\partial \varepsilon^i_t}}$$

for $K = 0, 1, 2, \ldots, 8$. A similar impulse response-based statistic was originally proposed and implemented in Barnichon and Mesters (2018b) in order to measure the sensitivity of price inflation to different slack measures, using a constant coefficient SVAR.

Figure 5A displays the evolution of the above dynamic multiplier conditional on monetary-policy shocks. As expected, the multiplier is always negative, suggesting that a monetary shock tends to move wage inflation and unemployment in opposite directions. We also see that its absolute value declines with the horizon, thus suggesting a more persistent effect of the shock on unemployment than on wage inflation. More interestingly, however, the (absolute) size of the multiplier appears to decrease over time. This is true at all the horizons considered (up to eight quarters) but particularly so at the shortest horizons. Thus, in the early part of the sample, we see how an expansionary monetary-policy shock that drives the unemployment rate down by one percentage point, simultaneously raises (quarterly) wage inflation by about 3 percentage points, thus implying a multiplier of 3. This short-run multiplier decreases over time in absolute value to a level close to 1. This finding is consistent with the evidence in previous sections pointing to a more muted response of wage inflation to fluctuations in unemployment. In contrast with that evidence, however, the present estimates suggest that such a change in responsiveness has been quite gradual, having started well before the financial crisis.

Figure 5B displays analogous evidence for (non-monetary) demand shocks. Many of the qualitative patterns observed in figure 5A are also present here, including the gradual decline in the (absolute) size of the estimated multiplier at all horizons. A different picture, however, emerges in figure 5C, which displays the dynamic multiplier for price-markup shocks and which does not suggest any major changes over time.

The above evidence rules out a change in the relative importance of different shocks as the main or only source of any reduction in the sensitivity of wage inflation to unemployment fluctuations: a change in that relation appears to have occurred even when conditional on specific shocks. Unfortunately, our approach cannot shed direct light on the nature of the structural change(s) that may underlie the lower conditional dynamic multipliers. A greater decoupling of wage inflation from price inflation, possibly due to a stronger anchoring of inflation expectations associated with the adoption of a price stability-oriented monetary policy, and captured in our estimates of the wage Phillips curve (both conditional and unconditional) would limit the so-called second round effects on wage inflation and dampen the response of the latter variable to any given variation in the unemployment rate, thereby providing a possible explanation to the evidence above.[14]

Figure 5A. Dynamic Multiplier: Monetary-Policy Shocks

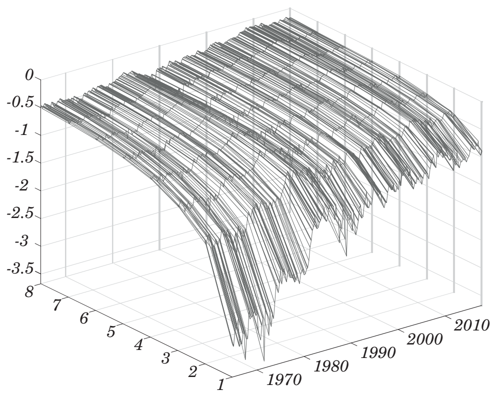

Figure 5B. Dynamic Multiplier: Demand Shocks

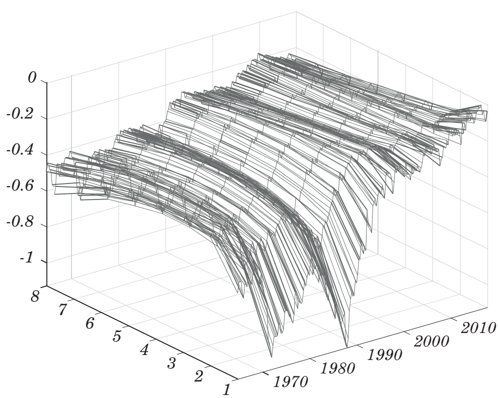

Source: Authors' calculations.

14. Blanchard and Galí (2009) point to that mechanism as an explanation of the smaller macroeconomics effects of oil price shocks in the 2000s relative to the 1970s.

Figure 5C. Dynamic Multiplier: Price-Markup Shocks

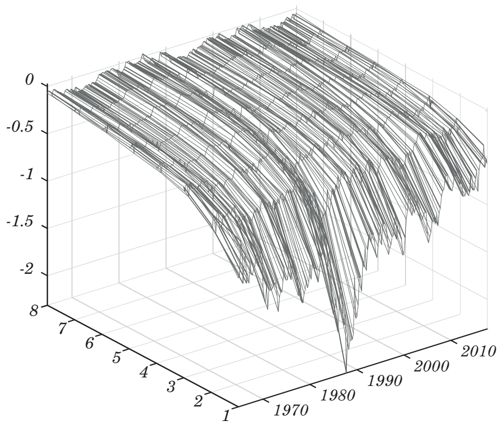

Source: Authors' calculations.

3. Conclusions

We have started the present paper by documenting the changes in the wage Phillips curve, using simple reduced-form regressions applied to aggregate U.S. data. In particular, we have provided evidence of a substantial decline in the estimated coefficients on both lagged inflation and unemployment in our wage Phillips curve. We have also provided estimates of *conditional* wage Phillips curves, based on a structural decomposition of wage, price, and unemployment data generated by a VAR with time-varying coefficients identified by a combination of long-run and sign restrictions. Our estimated conditional wage Phillips curves show that most qualitative findings from the reduced-form evidence are not driven by endogeneity problems or possible changes in the relative importance of shocks, though such factors may have led unconditional reduced-form estimates to overstate some of the actual changes. Finally, we have shown that the reduced sensitivity of wage inflation to unemployment is also reflected in the estimated changes in a dynamic-multiplier statistic, based on the estimated time-varying impulse responses to monetary-policy and demand shocks.

We draw two main conclusions from our findings. Firstly, we confirm the existence of a growing disconnect between wage inflation and unemployment. Secondly, more research is needed in order to understand the nature of that phenomenon.

References

Barnichon, R. and G. Mesters. 2018a. "Identifying the Phillips Curve from Shifts in Demand." Unpublished manuscript.

Barnichon, R. and G. Mesters. 2018b. "Direct Estimation of the Inflation-Unemployment Tradeoff." Unpublished manuscript.

Blanchard, O. and L. Katz. 1999. "Wage Dynamics: Reconciling Theory and Evidence." *American Economic Review* 89(2): 69–74.

Blanchard, O. and J. Galí. 2009. "The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so different from the 1970s?" (with O. Blanchard). In *International Dimensions of Monetary Policy*, edited by J. Galí and M. Gertler. Chicago, IL: University of Chicago Press.

Blanchard, O. 2016. "The Phillips Curve: Back to the 60s?" *American Economic Review* 106(5): 31–4.

Ciccarelli, M. and C. Osbat. 2017. "Low Inflation in the Euro Area: Causes and Consequences." ECB Occasional Paper Series No. 181.

Clarida, R., J. Galí, and M. Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics* 115(1): 147–80.

Coibion, O. and Y. Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics* 7(1): 197–232.

Constâncio, V. 2017. "Understanding and Overcoming Low Inflation." Remarks at the Conference "Understanding Inflation: Lessons from the Past, Lessons for the Future? Frankfurt am Main, 21–22 September.

Daly, M.C. and B. Hobijn. 2014. "Downward Nominal Wage Rigidities Bend the Phillips Curve." *Journal of Money, Credit, and Banking* 46(2): 51–93.

Del Negro, M. and G. Primiceri. 2013. "Time-Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum." Federal Reserve Bank of New York Staff Report No. 619.

Erceg, C.J., D.W. Henderson, and A.T. Levin. 2000. "Optimal Monetary Policy with Staggered Wage and Price Contracts." *Journal of Monetary Economics* 46(2): 281–314.

Galí, J. 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?, *American Economic Review*, 89(1), 249-271.

Galí, J. 2011a. "The Return of the Wage Phillips Curve." *Journal of the European Economic Association* 9(3): 436–61.

Galí, J. 2011b. *Unemployment Fluctuations and Stabilization Policies: A New Keynesian Perspective.* Cambridge , MA: MIT Press.

Galí, J., F. Smets, and R. Wouters. 2012. "Unemployment in an Estimated New Keynesian Model." *NBER Macroeconomics Annual 2011*, 329–60.

Galí, J. and L. Gambetti. 2009. "On the Sources of the Great Moderation," *American Economic Journal: Macroeconomics* 1(1): 26–57.

Galí, J. and L. Gambetti. 2014. "Monetary Policy and Rational Asset Price Bubbles," *American Economic Review*, vol. 104(3), 2014, 721-752.

Phillips, A.W. 1958. "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957." *Economica* 25: 283–99.

Smets, F. and R. Wouters. 2007. "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach," *American Economic Review* 97(3): 586–606.

Stock, J.H. and M.W. Watson. 2018. "Slack and Cyclically Sensitive Inflation." Presented at the ECB Forum on Central Banking, "Price and Wage-Setting in Advanced Economies", 16–18 June.

Uhlig, H. 2005. "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics* 52(2): 381–419.

Appendix

A.1 Some Microfoundations for the Wage Phillips Curve

In this appendix we sketch how one may derive our wage Phillips curve from a microfounded model. As shown in Galí (2011a, 2011b), the aggregation of wage decisions by monopolistically competitive unions reoptimizing the nominal wage with a constant probability every period and with partial indexation to a measure of lagged price inflation $\overline{\pi}_{t-1}^{p}$ in case of no reoptimization implies the following relation between wage inflation and the wage-markup gap $\mu_{w,t}-\mu_{w,t}^{n}$, i.e. the gap between the average wage markup $\mu_{w,t}$, and the natural (or flexible wage) wage markup, $\mu_{w,t}^{n}$

$$\tilde{\pi}_{t}^{w}=\beta\mathbb{E}_{t}\left\{\tilde{\pi}_{t+1}^{w}\right\}-\lambda_{w}(\mu_{w,t}-\mu_{w,t}^{n})$$

where $\tilde{\pi}_{t}^{w}=\pi_{t}^{w}-\left[\gamma\overline{\pi}_{t-1}^{p}+(1-\gamma)\pi\right]$. Assuming $\mu_{w,t}-\mu_{w,t}^{n}\sim AR(1)$ one can write:

$$\tilde{\pi}_{t}^{w}=-\frac{\lambda_{w}}{1-\beta\rho_{w}}(\mu_{w,t}-\mu_{w,t}^{n}).$$

In addition, and as shown in Galí (2011a, 2011b), the following relation between the average wage markup and the unemployment rate obtains:

$$\mu_{w,t}=\vartheta u_{t}.$$

Combining both relations one can derive

$$\pi_{t}^{w}=\left(1-\gamma\right)\pi+\gamma\overline{\pi}_{t-1}^{p}-\frac{\lambda_{w}\vartheta}{1-\beta\rho_{w}}\hat{u}_{t}+\frac{\lambda_{w}}{1-\beta\rho_{w}}\hat{u}_{w,t}^{n}$$

where a "hat" represents deviations from an assumed constant mean. Note that the previous specification is consistent with the estimated wage Phillips curve (1), with the error term in the latter capturing exogenous fluctuations in the natural wage markup.

A.2 Specification and Estimation of the Empirical Model

Let $\theta_t = vec(A'_t)$ where $A_t = [A_{0,t}, A_{1,t}, \ldots, A_{p,t}]$ and $vec(\,)$ is the column stacking operator. We assume θ_t evolves over time according to the following equation:

$$\theta_t = \theta_{t-1} + \omega_t \tag{4}$$

where ω_t is Gaussian white-noise vector process with covariance matrix Ω.

Time variation of Σ_t is modelled as follows. Let $\Sigma_t = \boldsymbol{F}_t \boldsymbol{D}_t \boldsymbol{F'}_t$ where $\boldsymbol{F}_t$ is lower triangular, with 1s on the main diagonal, and $\boldsymbol{D}_t$ a diagonal matrix. The vector containing the diagonal elements of $\boldsymbol{D}_t^{1/2}$, denoted by $\boldsymbol{\sigma}_t$, is assumed to evolve according to the process

$$log\boldsymbol{\sigma}_t = log\boldsymbol{\sigma}_{t-1} + \zeta_t, \tag{5}$$

Moreover, let $\phi_{i,t}$ denote the column vector with the non-zero elements of the $(i+1)$-th row of $\boldsymbol{F}_t^{-1}$. We assume

$$\phi_{i,t} = \phi_{i,t-1} + \nu_{i,t} \tag{6}$$

where ζ_t and $\nu_{i,t}$ are Gaussian white-noise vector processes with zero mean and (constant) covariance matrices Ξ and Ψ_i, respectively. We further assume that $\nu_{i,t}$ is independent of $\nu_{j,t}$ for all $j \neq i$, and that ω_t, ε_t, ζ_t, and $\nu_{i,t}$ (for all i) are mutually independent.

Priors Specification

We make the following assumptions about prior distributions:

$$\theta_0 \sim N(\theta, 4V_\theta)$$

$$log\boldsymbol{\sigma}_0 \sim N(log\boldsymbol{\sigma}_0, \mathrm{I}_n)$$

$$\phi_{i0} \sim N(\phi_i, V_{\phi_i})$$

$$\Omega^{-1} \sim N(\underline{\Omega}^{-1}, \underline{\rho}_1)$$

$$\boldsymbol{\Xi}^{-1} \sim N\left(\underline{\boldsymbol{\Xi}}^{-1}, \underline{\rho}_2\right)$$

$$\Psi^{-1} \sim N(\underline{\Psi}^{-1}, \underline{\rho}_1)$$

where $W(\mathbf{S},d)$ denotes a Wishart distribution with scale matrix $\mathbf{S}$ and degrees of freedom d and I_n is a $n \times n$ identity matrix, with n the number of variables in the VAR.

We use a time invariant VAR for x_t estimated by using the first $\tau = 64$ observations to calibrate prior means and variances. θ and V_θ are set equal to the OLS estimates. Let Σ be the covariance matrix of the residuals $\mathbf{u}_t$ of the initial time-invariant VAR. We apply the decomposition $\Sigma = \boldsymbol{FDF'}$ and set $\log\boldsymbol{\sigma}_0$ equal to the log of the diagonal elements of $\boldsymbol{D}^{1/2}$. ϕ_i is set equal to the OLS estimates of the coefficients of the regression of $\mathbf{u}_{i+1,t}$ the $i+1$-th element of $\mathbf{u}_t$, on $-\mathbf{u}_{1,t}, \ldots, -\mathbf{u}_{i,t}$ and V_{ϕ_i} equal to the estimated variances.

The scale matrices are parametrized as follows: $\underline{\Omega} = \underline{\rho}_1(\lambda_1 V_\theta), \underline{\Xi} = \underline{\rho}_2(\lambda_2 I_n)$ and $\underline{\Psi}\underline{\rho}_{3i}(\lambda_3 V_{\phi_i})$. The degrees of freedom $\underline{\rho}_1$ and $\underline{\rho}_2$ are set equal to the number of rows $\underline{\Omega}^{-1}$ and I_n plus one respectively while $\underline{\rho}_{3i}$ is $i+1$ for $i = 1, \ldots, n-1$. Finally $\lambda_1 = 0.0002$, $\lambda_2 = 0.01$, and $\lambda_3 = 0.01$.

Gibbs sampling algorithm

The Gibbs sampling algorithm is identical to that described in the online appendix of Galí and Gambetti (2014). We use all the data points available from 1964:Q2 to 2017:Q4. We draw 50000 realizations, discard the first 40000, and then take 1 out of 10 draws, therefore collecting a total of 1000 realizations.

Trade Exposure and the Evolution of Inflation Dynamics

Simon Gilchrist
New York University and NBER

Egon Zakrajšek
Federal Reserve Board

The Phillips curve—the relationship between price inflation and fluctuations in economic activity— is a central building block of economic models that allow for nominal rigidities and are relied upon by central banks around the world to gauge cyclical inflationary pressures and forecast inflation. The lack of deflationary pressures during the Great Recession and, more recently, the apparent lack of inflationary pressures during the recovery have brought into the forefront the question of whether this relationship still exists in the data.[1] More generally, the fact that inflation appears to have become less responsive to fluctuations in economic activity during the past couple of decades has been documented for the United States by Atkeson and Ohanian (2001), Roberts (2006), Mavroeidis and others (2001), and Blanchard (2016). This flattening of the Phillips curve appears to have occurred in other advanced economies as well;

We are grateful to Kimberly Bayard for her expert help with the industry-level industrial production data. We also thank Larry Ball, Mark Watson, Eduardo Zilberman (our discussant), Kurt Lewis, and conference participants for numerous helpful comments and suggestions. Tyler Pike and Gerardo Sanz-Maldonado provided excellent research assistance. The views expressed in this chapter are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System.

1. Recent work that studies the unusual inflation dynamics during the Great Recession and its aftermath in the United States and other advanced economies includes Stock and Watson (2010b), Ball and Mazumder (2011, 2018), Gordon (2013), Friedrich (2016), Berganza and others (2016), Miles and others (2017), Blanchard (2018), and Stock and Watson (2018).

Changing Inflation Dynamics, Evolving Monetary Policy edited by Gonzalo Castex, Jordi Galí, and Diego Saravia, Santiago, Chile.

see Beaudry and Doyle (2000) for Canada, and Kuttner and Robinson (2010) for Australia, for example.

Reasons for the apparent attenuation of the relationship between inflation and resource utilization are often linked to the rise in globalization and an associated increase in the cross-border movement of goods, services, technology, labor, and capital since the 1990s.[2] The resulting greater openness of national economies implies that a greater share of an increase in domestic demand is satisfied through imports, rather than domestic production. In turn, this implies that changes in the domestic output gap will have a smaller effect on domestic marginal costs, thereby reducing the responsiveness of domestic inflation to fluctuations in domestic economic slack, while increasing the sensitivity of domestic inflation to foreign economic slack. Increased international trade also gives rise to a common component for inputs such as commodities, thus implying that local costs—and hence prices—become less sensitive to domestic economic conditions. Increased openness of labor markets is another factor that attenuates the link between inflation and fluctuations in economic activity at the local level.[3]

Although prominent in recent policy discussions, the evidence in favor of a weakening in the relationship between inflation and economic activity due to increased global economic integration is mixed. Ball (2006) and Ihrig and others (2010) argue that there is little evidence to suggest that increased international trade and other globalization factors have attenuated the relationship between inflation and economic slack in the United States. Borio and Filardo (2007), Auer and others (2017), and Zhang (2017), on the other hand, present evidence that globalization has indeed led to a decline in the sensitivity of inflation to domestic factors, arguing that the integration of China and other lower-cost producers in world production networks has increased competition, thereby inducing downward pressure on wages and import prices in the U.S. and other industrial countries.

2. Another hypothesis posits that the observed flattening of the Phillips curve over the past couple of decades is due to a lower frequency of price adjustment at the firm level, reflecting the significantly lower average inflation rate that has prevailed over that period (see Ball and others, 1988). Relatedly, some economists have hypothesized that firms and households have started to pay less attention to macroeconomic conditions when setting wages and prices because of a prolonged period of low and stable inflation—the so-called rational inattention hypothesis (see Sims, 2003; Pfajfar and Roberts, 2018).

3. See Bernanke (2007) for an overview of the various channels through which ongoing global economic integration can affect inflation dynamics.

Similarly, Forbes (2018) shows that global factors have played a more prominent role in determining U.S. inflation outcomes since the 1990s; these global factors, however, are primarily linked to the food and energy component of consumer prices and play a diminished, rather than an increased, role in explaining movements in the core measures of U.S. consumer price inflation.

In this chapter, we re-examine this "globalization" hypothesis by using both U.S. aggregate data on measures of inflation and economic slack and a rich panel data set containing producer prices, wages, output, and employment at a narrowly defined industry level. Industries in our data set are defined at the 6-digit North American Industry Classification System (NAICS) level, and the data on prices and output serve as the basis for the construction of the U.S. producer price index (PPI) published by the Bureau of Labor Statistics and the industrial production index (IPI) published by the Federal Reserve Board. We also measure international trade exposure at the industry level—albeit at a somewhat coarser level of aggregation (i.e., 4-digit NAICS)—by using information on exports, imports, and value-added output. Linking these trade exposures to industry-level prices, wages, employment, and production allows us to directly determine the extent to which the response of inflation to fluctuations in output differs systematically across industries that are more or less exposed to international trade.

We begin our analysis by examining the time-series relationships between inflation and fluctuations in economic activity. Specifically, we consider the extent to which the relationship between inflation and economic activity has evolved over time. We address this question by estimating the sensitivity of both producer and consumer price inflation to economic slack using 15-year rolling-window regressions, starting in the early 1960s. This evidence shows that this relationship has indeed weakened substantially over the past 30 years or so. Importantly, our findings are robust to using both headline inflation measures, as well as core measures of inflation that remove the direct influence of swings in the volatile food and energy prices. They are also robust to measuring economic slack using alternative concepts such as the output gap or the unemployment gap.

We next consider the responsiveness of inflation to economic activity at the industry level. In this analysis, we exploit the cross-sectional dimension of our data and can directly control for the common aggregate component driving both inflation and output. We again find that fluctuations in output are an important determinant of inflation—indeed, the estimated response of industry-level inflation to variation in industry-level output is very similar to that obtained from aggregate time-series data over comparable sample periods.

Using both the industry-level and aggregate time-series data, we then examine the extent to which an increase in trade exposure has altered the response of inflation to fluctuations in economic activity. Here again our findings are consistent across both aggregate and industry-level data. In the time-series dimension, the rising exposure of the U.S. economy to international trade can indeed explain a significant fraction of the overall decline in responsiveness of aggregate inflation to economic slack. This result is confirmed by our cross-sectional evidence, which shows that increased trade exposure significantly dampens the response of inflation to fluctuations in output across industries.

The results discussed above, however, do not directly determine the causal impact of fluctuations in economic activity on inflation. While demand shocks typically move inflation and output in the same direction, supply shocks have the opposite effect. Thus any attenuation in the observed response of inflation to output may be due to changes in the mixture of demand and supply shocks that the U.S. economy has experienced over our sample period. To address this issue, we examine the effect of identified aggregate shocks on industry-level outcomes. In this exercise, we are explicitly interested in the extent to which the intensity of trade exposure at the industry level alters the responsiveness of inflation to such aggregate shocks.

Given the high dimensionality of our industry-level data, we consider the dynamic effects of identified aggregate demand shocks using a Factor-Augmented Vector Autoregression (FAVAR) model, whereby the information contained in the large panel of industries is summarized by a small subset of common factors. By using this framework, we study how shocks to broad financial conditions—a specific form of aggregate demand shocks—affect the dynamics of price and wage inflation, output, and employment at the industry level. We focus on disturbances to the financial intermediation process because we view them as readily identified from economic and financial time-series data; moreover, there exists a large body of empirical evidence indicating that financial shocks account for a sizable fraction of the variability in output and inflation over the past 30 years.[4]

4. See Gilchrist and others, 2009; Gilchrist and Zakrajšek, 2012; Jermann and Quadrini, 2012; and Peersman and Wagner, 2014.

Using the FAVAR approach, we first document that an unanticipated tightening in broad financial conditions causes a significant decline in price and wage inflation, as well as in output and employment growth across all industries. Thus financial shocks deliver the positive comovement between inflation and output that is typically associated with shocks to aggregate demand. We then examine the extent to which responses of inflation and output to financial shocks differ across industries based on their trade exposure. Our results indicate that industries with a high trade exposure exhibit a substantially smaller response of inflation to movements in output induced by the unanticipated changes in financial conditions, relative to industries with a low trade exposure.

These differential dynamics occur despite the fact that the effect of such shocks on economic activity is virtually identical across these two industrial groupings. Translated into the movements of inflation—relative to output—our results imply that, in response to such shocks, inflation is about three times more responsive to changes in output for industries with a low trade exposure, compared with industries with a high trade exposure. These findings further support the argument that external trade exposure attenuates the link between inflation and fluctuations in economic activity and that increased international trade is indeed a likely reason behind the reduced responsiveness of aggregate inflation to economic slack that has been observed in the data since the early 1990s.

The remainder of the chapter is organized as follows. Section 1 considers the aggregate time-series relationship between inflation and economic activity and documents its evolution over time. Section 2 explores the relationship between inflation and economic activity using industry-level data and documents the extent to which differences in trade exposure across industries affect this relationship. Section 3 provides the FAVAR analysis, which shows how industry-level variables respond to financial shocks, as well as the extent to which these responses differ across industries depending on their exposure to international trade. Section 4 offers a brief conclusion.

1. Aggregate Phillips Curve

In this section, we establish some stylized facts about the relationship between inflation and economic slack by using aggregate time-series data, which serve as a useful benchmark for the subsequent industry-level analysis. While the vast literature on this topic has focused on consumer price inflation, we analyze inflation dynamics at both the producer and consumer levels; the focus on the former is especially important because movements in producer prices directly capture the price response of production units to changes in the underlying economic conditions.[5] In terms of data used in this analysis, the solid line in panel A of figure 1 shows the behavior of prices received by U.S. producers for their output, measured by the four-quarter percent change in the PPI for final demand, while the solid line in panel B shows the four-quarter percent change in the consumer price index (CPI), a measure of prices paid by urban consumers for a market basket of consumer goods and services. The slashed lines in each panel show the corresponding core inflation, which strips out items belonging to the food and energy categories from each headline price index.[6]

Clearly evident in the data are several distinct inflation regimes. First, the 1970s, a period of high and volatile inflation that was early on influenced importantly by the OPEC-induced increases in oil prices (Hamilton, 1983) and later by the Federal Reserve's overly optimistic view of the natural rate of unemployment (Orphanides and Williams, 2013). The early 1980s, in contrast, were marked by a gradual step-down in inflation reflecting the tightening of monetary policy under Chairman Volcker, who was determined to fight inflation and reverse the rise in inflation expectations (Lindsey and others, 2005). Since the mid-1980s, inflation—at both the producer and consumer levels—has stabilized in a narrow range around two percent, a pattern consistent with the well-anchored inflation expectations engendered by credible monetary policy, aimed at achieving the so-called dual mandate stipulated by the Full Employment and Balanced Growth Act of 1978.[7]

5. It is worth noting that the frequency of price changes in the narrow-item categories that are both in the consumer and producer micro-level price data sets collected by the Bureau of Labor Statistics are highly correlated (see Nakamura and Steinsson, 2008).

6. Each quarterly price index is constructed as a simple average of the monthly (seasonally adjusted) index values, and four-quarter percent changes are computed as 100 times the four-quarter log-difference of the specified series. In addition, while we use the CPI to measure inflation at the consumer level, all the results reported below are robust to using the personal consumption expenditure (PCE) price index to track the change in prices of goods and services purchased by the U.S. consumers throughout the economy.

7. More commonly known as the Humprey-Hawkins Act, the Full Employment and Balanced Growth Act established price stability and full employment as national economic policy objectives.

Figure 1. Producer and Consumer Price Inflation

A. Producer price inflation

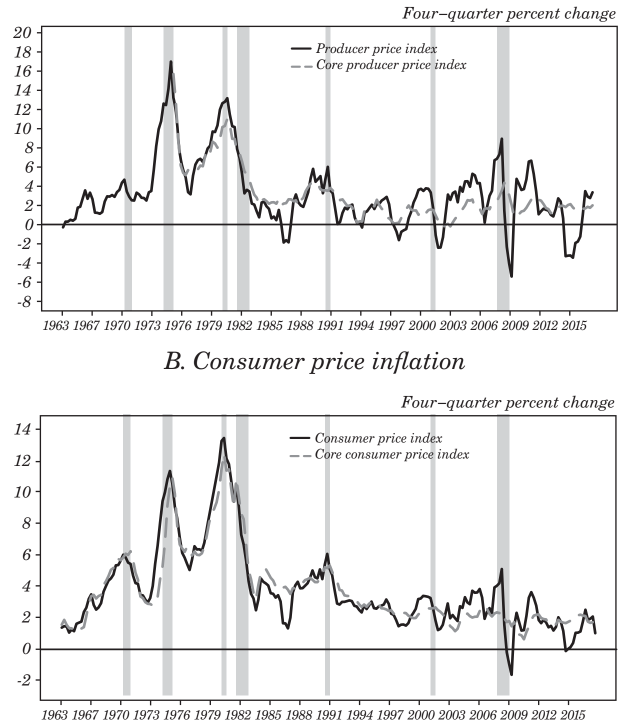

Source: Authors' calculations using data from the Bureau of Labor Statistics.
Note: All price indices are seasonally adjusted. The shaded vertical bars denote the NBER-dated recessions.

A striking way to illustrate how inflation is unresponsive to fluctuations in economic activity—in other words, how flat the Phillips curve is—is to focus on economic downturns. To that end, figure 2 examines the relationship between inflation and economic activity during the past five recessions, downturns in which supply-side disturbances—which cause inflation and economic activity to move in opposite directions—were arguably not the dominant factor. The first three panels of the figure depict the behavior of detrended prices two years before and after each NBER-dated cyclical peak since the early 1980s; the bottom right panel, by contrast, shows the corresponding dynamics of detrended real GDP, a simple measure of economic slack.

Figure 2. Inflation and Output in Recessions

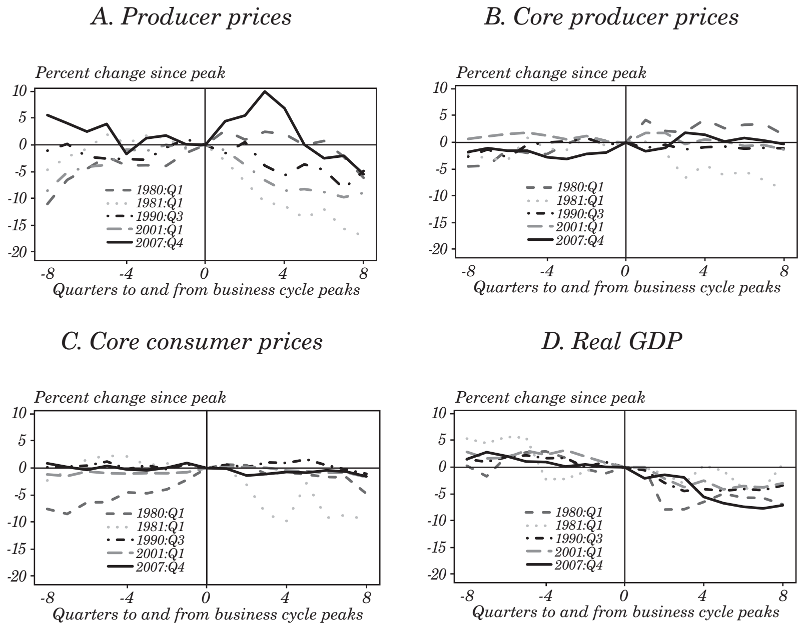

Source: Authors' calculations using data from the Bureau of Economic Analysis, the Bureau of Labor Statistics, and the Federal Reserve Board.
Note: The panels depict the behavior of various price measures and real GDP eight quarters before and eight quarters after the specified NBER-dated cyclical peak. All series are plotted as deviations from their respective stochastic trends, estimated using the Hamilton (2018) filter.

As shown in the top two panels, with the exception of the 2001 recession, producer prices—especially those that exclude the volatile food and energy components—showed virtually no deceleration during the past five economic downturns, relative to their trends. And even during the bursting of the tech bubble in 2001, the decline in both the headline and core PPI is due entirely to the plunge in producer prices in the immediate aftermath of the September 11 terrorist attacks—in October 2001, the Bureau of Labor Statistics reported that the PPI dropped almost 20 percent at an annual rate.[8] As shown in the bottom left panel, the resilience of inflation in response to the emergence of

8. It is also worth noting that the sharp increase in commodity prices prompted by the First Gulf War confounds the behavior of PPI inflation during the 1990 recession to some extent.

substantial economic slack is also evident at the consumer level. At the same time, as shown in the bottom right panel, real GDP declined markedly—relative to its trend—during these five episodes.

1.1 Baseline Estimates

To investigate more formally how the relationship between inflation and fluctuations in economic activity may have changed over time, we begin by estimating a standard Phillips curve specification, which expresses inflation as a linear function of expected inflation and a measure of economic slack. Specifically, letting lower-case variables denote variables in logarithms and defining $\Delta_h x_{t+h} = \frac{400}{h}(x_{t+h} - x_t)$, we estimate the following Phillips curve specification:

$$\Delta_{h+1} p_{t+h} = \mu + \lambda \text{gap}_t + \sum_{s=1}^{4} \phi_s \Delta p_{t-s} + \epsilon_{t+h}, \tag{1}$$

where p_t denotes the logarithm of a price index (i.e., PPI or CPI) and gap_t is a measure of economic slack, a degree of resource over- or under-utilization. Thus equation (1) posits a relationship between (annualized) inflation from quarter $t-1$ to quarter $t+h$ and a measure of economic slack in quarter t, while the lags of inflation $\Delta p_{t-s}, s = 1,\ldots,4$, are a proxy for expected inflation.[9,10] In this canonical formulation, the error term ϵ_{t+h} encompasses cost-push shocks—shock to commodity prices, for example—which are assumed to be uncorrelated with the contemporaneous economic slack.[11]

Figure 3 shows two measures of economic slack used in our analysis: the output gap and the unemployment gap. The output gap, denoted by $[y_t - y_t^*]$, is defined as (100 times) the logarithm of the ratio of real GDP to its estimate of potential, while the unemployment

9. See Gordon, 1982; Stock and Watson, 2009.

10. An alternative approach to using lagged values of inflation to capture expected inflation would be to use survey measures of expected inflation. However, as documented by Mankiw and others (2004), such survey measures do not appear to be consistent with either rational expectations or adaptive expectations used in specification (1).

11. It is worth noting that the presence of very low frequency variation in both the producer and consumer inflation rates (see figure 1) has the potential to confound the relationship between inflation and fluctuations in economic slack at the business cycle frequency, which is the primary interest of our analysis. To ensure that our baseline time-series results are not unduly affected by this low frequency variation, appendix A contains a robustness analysis in which all inflation series are expressed as deviations from their respective local means. As evidenced by those results, the main conclusions of this section are robust to this transformation of the data.

gap, denoted by $[U_t - U_t^*]$ corresponds to the unemployment rate less its estimate of the natural rate. The estimates of both the potential real GDP and the natural rate of unemployment are taken from the FRB/US model, a large-scale estimated general equilibrium model of the U.S. economy that has been in use at the Federal Reserve Board since 1996. While the definition of these two slack measures naturally produces series of the opposite sign, they paint a very similar picture of cyclical resource utilization over the last 50 years or so. One exception to this pattern has occurred during the past several years, a period in which the unemployment rate has moved below its natural rate, whereas the real GDP has yet to return to its potential.[12]

Table 1 present estimates of the coefficient λ for producer price inflation at horizons of one and four quarters (i.e., $h = 1, 4$), with panel A showing estimates of λ for headline PPI inflation and panel B showing estimates of λ for core PPI inflation; the corresponding estimates of λ for consumer price inflation—both headline and core—are shown in table 2.

Figure 3. Economic Slack

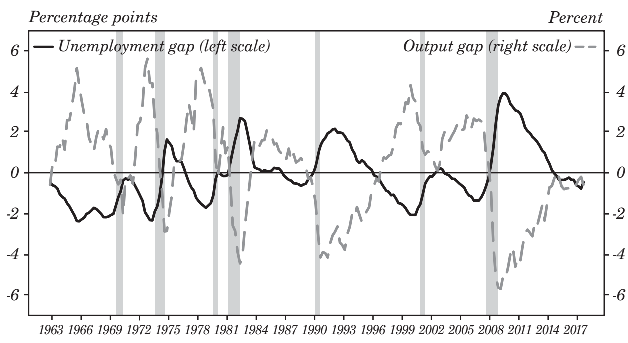

Source: Bureau of Economic Analysis, Bureau of Labor Statistics, and Federal Reserve Board.
Note: The output gap is defined as (100 times) the log-ratio of real GDP to its estimate of potential; the unemployment gap is defined as the civilian unemployment rate less its estimate of the natural rate. The shaded vertical bars denote the NBER-dated recessions.

12. Movements in the output gap can be interpreted as capturing fluctuations in real marginal cost, which microfounded models emphasize as a key determinant of inflation dynamics (see Roberts, 1995; Galí and Gertler, 2000; Galí and others, 2001; Sbordone, 2002; and Galí and others, 2007).

Table 1. Phillips Curve – Producer Price Inflation

Explanatory variables	*h* = 1 (*1*)	(*2*)	*h* = 4 (*3*)	(*4*)
A. Producer prices				
$[y_t - y_t^*]$	0.356**	-	0.414***	-
	(0.144)		(0.153)	
$[U_t - U_t^*]$	-	-0.396*	-	-0.469*
		(0.238)		(0.257)
Sum: inflation lags[a]	0.578***	0.600***	0.470***	0.495***
	(0.113)	(0.113)	(0.093)	(0.100)
sup W^b	15.185***	11.345**	33.370***	23.284***
	[81.Q2]	[91.Q4]	[80.Q3]	[93.Q4]
q_{LL}^c	-6.230	-5.375	-5.297	-4.325
Adj. R^2	0.360	0.333	0.392	0.343
B. Core producer prices				
$[y_t - y_t^*]$	0.186***	-	0.223***	-
	(0.056)		(0.067)	
$[U_t - U_t^*]$	-	-0.243**	-	-0.273**
		(0.105)		(0.131)
Sum: inflation lags[a]	0.776***	0.797***	0.730***	0.755***
	(0.071)	(0.076)	(0.071)	(0.081)
sup W^b	21.278***	18.00***	70.033***	39.261***
	[81.Q4]	[93.Q4]	[81.Q4]	[82.Q3]
q_{LL}^c	-9.554**	-7.550*	-6.304	-5.737
Adj. R^2	0.743	0.725	0.760	0.727

Source: Authors' calculations.
Note: Sample: 1962:Q2 to 2017:Q4 for headline PPI (panel A), and 1974:Q1 to 2017:Q4 for core PPI (panel B). The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the annualized log-difference in the specified PPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap and $[U_t - U_t^*]$ = unemployment gap. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.
[a] Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.
[b] The Andrews (1993) sup-Wald statistic of the null hypothesis that there is no structural break in the coefficient on economic slack; the estimated break dates are reported in brackets below.
[c] The Elliott and Müller (2006) q_{LL} statistic of the null hypothesis that the coefficient on economic slack is constant over time.

Table 2. Phillips Curve – Consumer Price Inflation

Explanatory variables	*h* = 1 (*1*)	*h* = 1 (*2*)	*h* = 4 (*3*)	*h* = 4 (*4*)
A. Consumer prices				
$[y_t - y_t^*]$	0.258***	-	0.318***	-
	(0.075)		(0.084)	
$[U_t - U_t^*]$	-	-0.321***	-	-0.380***
		(0.120)		(0.128)
Sum: inflation lags[a]	0.779***	0.795***	0.690***	0.709***
	(0.066)	(0.070)	(0.068)	(0.077)
sup W^b	34.118***	28.008***	70.231***	44.548***
	[83.Q1]	[83.Q2]	[83.Q1]	[83.Q1]
$q_{LL}{}^c$	-8.199*	-6.892	-6.347	-4.986
Adj. R^2	0.657	0.635	0.676	0.632
B. Core consumer prices				
$[y_t - y_t^*]$	0.176***	-	0.265***	-
	(0.044)		(0.060)	
$[U_t - U_t^*]$	-	-0.263***	-	-0.364***
		(0.079)		(0.107)
Sum: inflation lags[a]	0.868***	0.875***	0.787***	0.797***
	(0.056)	(0.060)	(0.065)	(0.074)
sup W^b	38.828***	38.420***	112.255***	66.416***
	[83.Q1]	[83.Q1]	[83.Q1]	[83.Q1]
$q_{LL}{}^c$	-8.639***	-6.259	-7.278*	-6.056
Adj. R^2	0.802	0.794	0.778	0.750

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4. The dependent variable in each Phillips curve specification is $\Delta_{h+1}p_{t+h}$, the annualized log-difference in the specified CPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap, and $[U_t - U_t^*]$ = unemployment gap. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

[a] Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.

[b] The Andrews (1993) sup-Wald statistic of the null hypothesis that there is no structural break in the coefficient on economic slack; the estimated break dates are reported in brackets below.

[c] The Elliott and Müller (2006) q_{LL} statistic of the null hypothesis that the coefficient on economic slack is constant over time.

According to the entries in panel A of table 1, fluctuations in economic slack have a significant effect on the subsequent behavior of producer prices. A decrease in resource utilization of one percentage point in quarter *t*—that is, a decline in the output gap or a rise in the unemployment gap of that magnitude—is estimated to reduce annualized headline producer price inflation over the next several quarters about 40 basis points. The corresponding estimates for core PPI inflation shown in panel B are about one-half as large as those reported in panel A, though the estimates are significant in both economic and statistical terms. As shown in table 2, economic slack is also a significant determinant of consumer price inflation. In that case, a decrease in resource utilization of one percentage point is estimated to shave off about 25 basis points from annualized CPI inflation over the subsequent few quarters.

As a first pass on the question of whether the relationship between economic slack and subsequent inflation may have changed over the past 50 years or so, we report results of two statistical tests. The first is the well-known Andrews (1993) test of a structural break—at an unknown date— in the coefficient λ. The second is the Elliott and Müller (2006) test of stability of the coefficient λ, which encompasses diverse forms of parameter instability—from relatively rare (including a single break) to frequent small breaks, persistent temporal parameter variation, and breaks occurring with a regular pattern.[13] This statistical analysis, however, yields a mixed picture. Turning first to producer prices (table 1), the Andrews (1993) test provides strong evidence of a structural break in λ, with the point estimate of a break date generally falling in the early 1980s, a result consistent with that of Roberts (2006). The evidence of parameter instability from the Elliott and Müller (2006) test, in contrast, is considerably weaker. A similar picture emerges when we look at consumer prices (table 2). Here again, the Andrews (1993) test strongly suggest a structural break in λ that occurred in the early 1980s, whereas the results from the Elliott and Müller (2006) test are far less conclusive.

All told, the results reported in tables 1 and 2 clearly indicate an important role—in both economic and statistical terms—for economic

13. In both tests, the null hypothesis is that the coefficient λ is stable over the sample period. The alternative in the Andrews (1993) test is that $\lambda = \lambda_1$ for $t = 1,2,\ldots,\tau - 1$ and $\lambda = \lambda_2$ for $t = \tau, \tau + 1, \ldots T$, where τ is the unknown (single) break date. The alternative in the Elliott and Müller (2006) test is $\lambda = \lambda_t$, where the time variation in the parameter λ_t is unspecified and can take on a variety of forms.

slack as a determinant of cyclical inflation dynamics. Nevertheless, empirical Phillips curves of the type given by equation (1) predicted a significantly greater downward pressure on inflation—if not outright deflation—during the Great Recession than was actually realized. Economists have advanced a number of hypotheses to explain this case of "missing deflation." A prominent hypothesis that received a lot of attention in policy circles argues that the Federal Reserve's credibility has led businesses and households to discount inflation outcomes that fall outside the narrow range bracketing the Federal Open Market Committee's inflation target of two percent; this anchoring of agents' expectations has—through the standard expectational effects—prevented actual inflation from falling significantly below that level.[14]

Another frequently cited hypothesis posits that the relevant measure of economic slack in empirical Phillips curves is not the overall unemployment rate gap, but rather the short-term unemployment rate.[15] Compared with the former, this latter indicator of slack increased notably less during the Great Recession and has also returned more quickly to its pre-recession levels, thus providing substantially less deflationary impetus. And although it has proven difficult to identify structural changes in the economy that could account for the diminished sensitivity of inflation to the level of unemployment, a number of economists have singled out the apparent flattening of the Phillips curve as an important reason for the fact that the U.S. economy did not experience a Fisherian debt-deflation spiral during the 2008–2009 global financial crisis.[16, 17]

1.2 Time-Varying Estimates

In light of the above discussion and the results reported in tables 1 and 2, it seems clear that a further investigation in the time-varying nature of the relationship between inflation and economic slack is warranted. As a simple and relatively straightforward way to

14. See Bernanke, 2010; Yellen, 2013.

15. Underlying this argument is the idea that workers who have been unemployed for a relatively short time are the relevant margin for wage adjustment. The longer-term unemployed, by contrast, do not put much downward pressure on wages because these potential workers are disconnected from the labor market (see Stock and Watson, 2010b; Gordon, 2013; Krueger and others, 2014).

16. See Ball and Mazumder, 2011; Simon and others, 2013.

17. Gilchrist and Zakrajšek (2016); Gilchrist and others (2017), in contrast, emphasize how the interaction of financial distress and customer markets attenuated deflationary pressures during the Great Recession.

consider the possibility of time variation in the coefficient λ—as well as in other parameters of the standard Phillips curve—we re-estimate specification (1) by using a 15-year rolling window. We then plot the time-varying coefficient on the specified measure of economic slack, along with its 95-percent confidence interval. To conserve space, we focus on the Phillips curve specifications for inflation at the horizon of four quarters (i.e., $h = 4$). The resulting time-varying estimates of the coefficient λ, for both the headline and core PPI inflation, are shown in figure 4, with panel A showing the time-varying sensitivity to the output gap and panel B showing the time-varying sensitivity to the unemployment gap; the comparable estimates for CPI inflation are shown in figure 5.[18]

The left chart in panel A of figure 4 shows the evolution of the response of headline PPI inflation to the output gap. In the early part of the sample, the estimates of λ are greater than one and significantly different from zero, according to the 95-percent confidence intervals. Starting in the mid-1980s, however, these estimated sensitivities begin to decline steadily before stabilizing in the late 1990s. From then onward, the estimates of λ fluctuate in a fairly narrow range between zero and 0.5, though for most of this latter sample period, one would not reject the hypothesis that the coefficient on the output gap is statistically different from zero.

The left chart in panel B shows the evolution of the response of headline PPI inflation to the unemployment gap. We observe roughly the same general pattern in this case. The estimates of λ start out negative and large in economic terms, as well statistically different from zero, according to the conventional significance levels. Once the late 1980s enter the sample period, however, the estimates begin to converge rapidly to zero. The estimated response of headline PPI inflation to the unemployment gap then remains around zero for the remainder of the sample period.

18. The convention is that the data point labeled “1994:Q4,” for example, represents an estimate based on the 1980:Q1–1994:Q4 estimation window. For both the headline producer and consumer price inflation, as well as for the core consumer price inflation, our sample period—allowing for lags—starts in 1962:Q2, so that the rolling-window estimates begin in 1978:Q1 and run through 2017:Q4, the end of our sample period. Core producer prices, by contrast, start in 1974:Q1, which implies that the first rolling-window estimates—again allowing for lags—become available in 1990:Q1.

Figure 4. Time-Varying Coefficient on Economic Slack
(Phillips Curve – Producer price inflation)

A. Economic slack: output gap

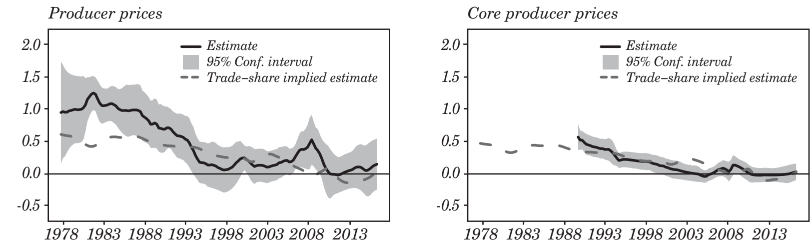

B. Economic slack: unemployment gap

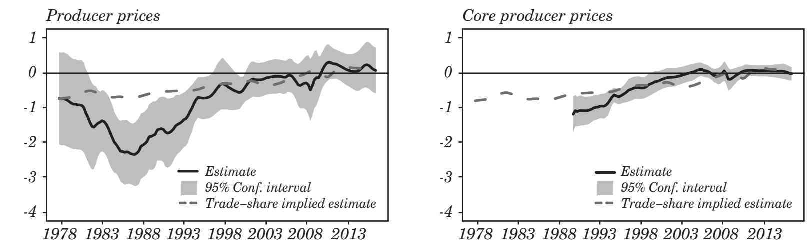

Source: Authors' calculations.
Note: The dependent variable in each Phillips curve specification is $\Delta_5 p_{t+4}$, the annualized log-difference in the specified PPI from date $t-1$ to date $t+4$. The solid line in each panel depicts the time-varying coefficient on the specified measure of economic slack estimated using a 60-quarter moving window; the dashed lines depict the corresponding time-varying coefficients implied by specifications (3) and (4) in table 3 (see notes to the table and the text for details).

The corresponding right charts of figure 4 trace out the estimated sensitivities of core PPI inflation to the output gap (panel A) and the unemployment gap (panel B). Although the sample begins later in this instance, the rolling-window estimates of the coefficient λ in the Phillips curve for core PPI inflation are much more precisely estimated than their counterparts for headline inflation. The estimates of λ for the output gap begin at about 0.5 for the sample that extends from the mid-1970s to the end of the 1980s and then decline monotonically to zero as the sample period moves forward; in fact, the estimate of λ based on the last 15 years of available data implies a sensitivity of core PPI inflation to the output gap that is economically and statistically indistinguishable from zero. The time-series pattern of coefficients on the unemployment gap is very similar: The estimates of λ start out negative, large in absolute value, and are precisely estimated and then converge to zero by the end of the 1990s.

Figure 5. Time-Varying Coefficient on Economic Slack
(Phillips Curve – Consumer price inflation)

A. Economic slack: output gap

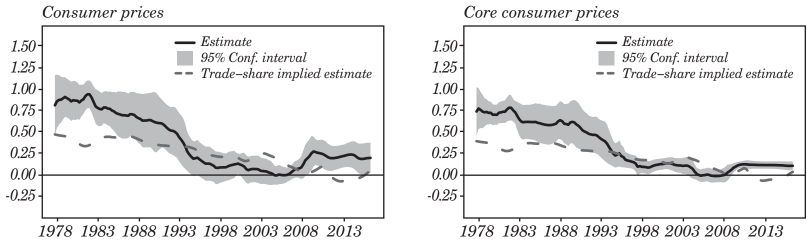

B. Economic slack: unemployment gap

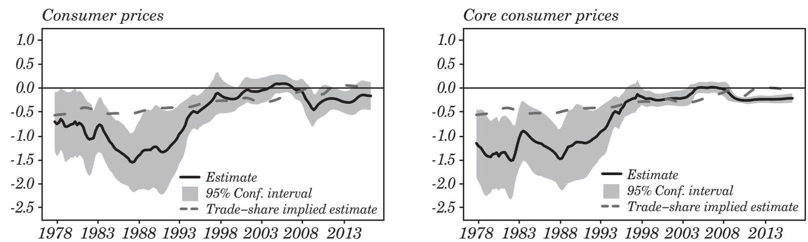

Source: Authors' calculations.
Note: The dependent variable in each Phillips curve specification is $\Delta_5 \, p_{t+4}$, the annualized log-difference in the specified CPI from date $t-1$ to date $t+4$. The solid line in each panel depicts the time-varying coefficients on the specified measure of economic slack estimated using a 60-quarter moving window; the dashed lines depict the corresponding time-varying coefficients implied by specifications (3) and (4) in table 4 (see notes to the table and the text for details).

Figure 5 shows the time-varying coefficient estimates on economic slack for both the headline and core measures of CPI inflation. As before, panel A shows coefficient estimates on the output gap, while the corresponding estimates for the unemployment gap are shown in panel B. Consistent with the full-sample estimates of λ reported in panel A of table 2, the time-varying coefficient estimates of the response of inflation to the output gap for headline CPI inflation are very similar to those for core inflation, both in terms of their magnitude and their evolution over time. They also show a pattern similar to that shown in figure 4: The estimates of λ are positive, economically and statistically significant in the early part of the sample, and then begin to decline sharply once the 1990s enter the estimation window. In contrast to the estimated response coefficients for PPI inflation shown in panel A of figure 4, the sensitivity of CPI inflation to the output gap is estimated to have increased notably at the end of our sample period. That said, these late-sample estimates of λ, though statistically different from zero, are only about one-fourth of those estimated during the early part of our sample period. According to panel B of figure 5, these patterns are robust to using the unemployment gap, rather than the output gap, as a measure of economic slack.

While there are a variety of phenomena that may help explain the declining sensitivity of aggregate inflation to fluctuations in economic activity, we are specifically interested in the extent to which increased globalization and trade may have contributed to the flattening of the Phillips curve. The notion that increased trade may help account for such changes is consistent with the rising trade intensity in the United States—defined as the sum of exports and imports relative to GDP—shown in figure 6. According to this metric, the trade intensity of the U.S. economy has risen by nearly a factor of three over the past 50 years or so.

To test the hypothesis that increased trade intensity of the U.S. economy may have contributed to the observed decline of the sensitivity of inflation to economic slack, we estimate the following variant of our baseline Phillips curve specification:

$$\Delta_{h+1}\mathrm{p}_{t+h} = \mu + \lambda_1 \mathrm{gap}_t + \lambda_2\left[\mathrm{gap}_t \times \mathrm{TrdShr}_{t-1}\right] + \sum_{s=1}^{4}\phi_s \Delta p_{t-s} + \epsilon_{t+h}, \qquad (2)$$

where TrdShr_t denotes an eight-quarter trailing moving average of the U.S. trade share shown in figure 6.[19] The resulting coefficient estimates of λ_1 and λ_2 for PPI inflation are reported in table 3, while those for CPI inflation are reported in table 4.

19. The Phillips curve specification (2) is similar to that used by Ball (2006), except that it does not include the “smoothed” trade share, TrdShr_{t-1}, as a separate explanatory variable; the inclusion of this term, however, had no material effect on any of the results reported below. Note also that appendix A contains results from the estimation, which controls for the slow-moving changes in the average inflation rate over our sample period; again, those results are qualitatively and quantitatively similar to those reported in the main text.

Figure 6. U.S. Trade Share

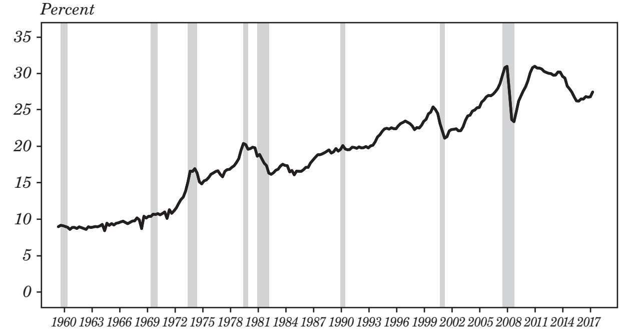

Source: Bureau of Economic Analysis.
The trade share is defined as the sum of the nominal value of U.S. imports and exports, expressed as a percent of nominal GDP. The shaded vertical bars denote the NBER-dated recessions.

According to the entries reported in panel A of table 3, the coefficient on the interaction term between the output gap and the trailing moving average of the U.S. trade share is negative— though not statistically different from zero—at the one-quarter horizon (column 1) and negative and marginally significant at the four-quarter horizon (column 3). Similarly, the interaction effect between the unemployment gap and trade share is positive and imprecisely estimated for $h=1$, whereas the coefficient on the interaction term for $h=4$ is positive and statistically different from zero at the 10-percent significance level. On balance, therefore, the evidence based on headline PPI inflation does not seem to support strongly the hypothesis that increased trade exposure of the U.S. economy can explain the decline in the sensitivity of inflation to fluctuations in economic activity.

As shown in panel B, however, the corresponding estimates for core PPI inflation paint a very different picture. The coefficients on the interaction terms between the output gap and trade share are negative and quite precisely estimated at both the one- and four-quarter horizons (columns 1 and 3). And similarly, the coefficients on the interaction terms between the unemployment gap and trade share are negative and statistically different from zero for both $h=1$ and $h=4$ (columns 2 and 4). Moreover, these estimates are economically meaningful. At the four-quarter horizon, they imply that when the trade share was at the 5th percentile of its distribution, the sensitivity of core PPI inflation to the output gap was 0.766 (std. error = 0.186) and −1.262 (std. error = 0.416) when the unemployment gap is used to gauge the degree of resource utilization in the economy; the corresponding estimates at the 95th percentile, in contrast, are −0.058 (std. error = 0.122) and 0.092 (std. error = 0.188), respectively. In other words, these results indicate a robust relationship between the rising trade share and the diminished sensitivity of core PPI inflation to fluctuations in economic activity.

In table 4, we report the estimates of coefficients λ_1 and λ_2 for Phillips curve specifications involving headline (panel A) and core (panel B) CPI inflation. These results again imply an economically large and statistically significant reduction in the responsiveness of inflation to economic slack as the trade share rises. This is true for both the headline and core measures of CPI inflation and holds at both the one- and four-quarter horizons. Moreover, the strong attenuation of the response of CPI inflation to economic slack is robust to using either the output or the unemployment gap as a gauge of cyclical resource utilization.

To summarize how the increasing exposure of the U.S. economy to international trade over the past 50 years affected the responsiveness of inflation to fluctuations in economic activity, we calculate the time-series evolution of the response coefficients associated with economic slack, as implied by the estimates of coefficients λ_1 and λ_2 reported in tables 3 and 4 and the trajectory of the U.S. trade share shown in figure 6. We then plot these estimates, as dashed lines in figures 4 and 5, next to their corresponding time-varying estimates based on the 15-year rolling window. The comparison of solid and dashed lines in the panels of these two figures shows that this specific parametrization of the time-varying slope of the Phillips curve—a simple interaction between the trade share and economic slack—can account for about one-half of the decline in the responsiveness of PPI and CPI inflation to economic slack observed over the past 50 years. In addition, this interaction effect captures remarkably well the attenuation in the response of core PPI inflation to changes in economic slack that we observe during the latter part of the sample period.

Table 3. Phillips Curve and the Trade Share – Producer Price Inflation

Explanatory variables	h = 1		h = 4	
	(1)	(2)	(3)	(4)
A. Producer prices				
$[y_t - y_t^*]$	0.968**	-	1.459**	-
	(0.449)		(0.584)	
$[y_t - y_t^*] \times \text{TrdShr}_{t-1}$	-0.030	-	-0.052*	-
	(0.023)		(0.07)	
$[U_t - U_t^*]$	-	-1.349**	-	-1.759**
		(0.674)		(0.822)
$[U_t - U_t^*] \times \text{TrdShr}_{t-1}$	-	0.047	-	0.063*
		(0.033)		(0.038)
Sum: inflation lags[a]	0.583***	0.604***	0.479***	0.500***
	(0.111)	(0.108)	(0.086)	(0.093)
Adj. R^2	0.369	0.342	0.440	0.371
B. Core producer prices				
$[y_t - y_t^*]$	0.903*	-	1.125***	-
	(0.309)		(0.305)	
$[y_t - y_t^*] \times \text{TrdShr}_{t-1}$	-0.031**	-	-0.040***	-
	(0.014)		(0.014)	
$[U_t - U_t^*]$	-	-1.591***	-	-1.851***
		(0.596)		(0.648)
$[U_t - U_t^*] \times \text{TrdShr}_{t-1}$	-	0.056**	-	0.065**
		(0.024)		(0.026)
Sum: inflation lags[a]	0.751***	0.790***	0.698***	0.747***
	(0.067)	(0.067)	(0.058)	(0.068)
Adj. R^2	0.762	0.742	0.794	0.754

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4 for headline PPI (panel A), and 1974:Q1 to 2017:Q4 for core PPI (panel B). The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the annualized log-difference in the specified PPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap; $[U_t - U_t^*]$ = unemployment gap, and TrdShr_{t-1} = eight-quarter (trailing) moving average of the trade share. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

[a] Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.

Table 4. Phillips Curve and the Trade Share – Consumer Price Inflation

Explanatory variables	$h = 1$		$h = 4$	
	(1)	*(2)*	*(3)*	*(4)*
A. Consumer prices				
$[y_t - y_t^*]$	0.728***	-	1.093***	-
	(0.231)		(0.311)	
$[y_t - y_t^*] \times \mathrm{TrdShr}_{t-1}$	-0.023**	-	-0.038***	-
	(0.011)		(0.013)	
$[U_t - U_t^*]$	-	-0.997***	-	-1.282***
		(0.353)		(0.425)
$[U_t - U_t^*] \times \mathrm{TrdShr}_{t-1}$	-	0.033**	-	0.044**
		(0.016)		(0.018)
Sum: inflation lags[a]	0.801***	0.812***	0.727***	0.734***
	(0.063)	(0.067)	(0.063)	(0.075)
Adj. R^2	0.670	0.645	0.721	0.656
B. Core consumer prices				
$[y_t - y_t^*]$	0.526***	-	0.918***	-
	(0.133)		(0.200)	
$[y_t - y_t^*] \times \mathrm{TrdShr}_{t-1}$	-0.017***	-	-0.032***	-
	(0.005)		(0.008)	
$[U_t - U_t^*]$	-	-0.802***	-	-1.197***
		(0.237)		(0.340)
$[U_t - U_t^*] \times \mathrm{TrdShr}_{t-1}$	-	0.026***	-	0.040***
		(0.009)		(0.013)
Sum: inflation lags[a]	0.891***	0.896***	0.831***	0.851***
	(0.053)	(0.058)	(0.061)	(0.072)
Adj. R^2	0.811	0.803	0.815	0.874

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4. The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the annualized log-difference in the specified CPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap, $[U_t - U_t^*]$ = unemployment gap, and TrdShr_{t-1} = eight-quarter (trailing) moving average of the trade share. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

[a] Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.

2. Industry-Level Phillips Curve

The combination of a rising trade share with the concomitant decline in the responsiveness of aggregate inflation to fluctuations in economic activity provides suggestive evidence that the observed flattening of the Phillips curve is at least partly due to increased trade intensity of the U.S. economy. The variation used to estimate this effect, however, relies solely on the secular increase in the U.S. trade share over the past 50 years or so and moreover does not fully explain the substantial reduction in the estimated slope of the aggregate Phillips curve. To provide a more thorough analysis of this phenomenon, we now turn to industry-level data, where we can exploit variation in trade shares across industries to test whether a differential trade exposure influences the sensitivity of inflation to economic slack.

2.1 Data Sources and Methods

To construct the panel data set used in this analysis, we utilize the most detailed (i.e., 6-digit NAICS) industry-level PPIs published by the Bureau of Labor Statistics, which we merge with the corresponding industry-level data on industrial production—a measure of output—constructed by the Federal Reserve.[20] The resulting data set covers all 6-digit NAICS industries—excluding those in the Utilities sector (i.e., 2-digit NAICS code 22)—that are used to produce both the producer price and industrial production indices for the U.S. economy. The industry-level price and production data are available at the monthly frequency, and we convert them to quarterly frequency by simply averaging the values of each index over the three months of each quarter.

The industry-level price and production data are available starting in the early 1970s. However, the data are not available for every industry from the beginning—that is, the panel is unbalanced—and there is an especially large expansion in the number of industries covered that occurred in the mid-1980s. To capture this broad array of industries, we thus begin our sample in 1984:Q1. All told, our

20. IPIs are not available for the full set of 6-digit NAICS industries. At such a fine level of disaggregation, there are in some cases an insufficient number of production units to construct a meaningful estimate of the index. In those instances, the staff at the Federal Reserve Board aggregates the underlying data across several of such closely related industries. In our matching algorithm, we assigned such industrial production data to all the 6-digit industries in the index.

unbalanced panel includes price and production data for 319 industries at the 6-digit NAICS level, covering the period from 1984:Q1 to 2017:Q4. We complement these industry-level data on output and prices with the corresponding data on wages and employment from the Quarterly Census of Employment and Wages (QCEW), a data collection program that publishes a quarterly count of employment, total wages, and average weekly wages per employee, reported by companies covering more than 95 percent of U.S. jobs. The QCEW data, however, are available only starting in 1990:Q1. We thus also consider a more in-depth analysis by using a balanced panel of 185 industries for which all of these variables are available over the 1990:Q1–2017:Q4 period.[21]

To measure trade exposure at the industry level, we rely on the annual (nominal) import and export data, which are made available by the Center for International Data at the University of California Davis and cover the period from 1972 to 2006.[22] The data provided are disaggregated by country (source for imports and destination for exports) and Schedule B number. These data were first aggregated to the total annual imports and exports at the industry level by using the 5-digit Standard Industrial Classification (SIC) codes. The annual (nominal) imports and exports for the 2007–2017 period were obtained from the U.S. Census Bureau's USA Trade Online database and are available at the 10-digit Harmonized System Code (HTS) level. By using various crosswalks, all of these data had to be first mapped to industries at the 6-digit NAICS level. At such a fine level of disaggregation, however, there are numerous missing industry/year observations. Accordingly, we aggregated trade date to the 4-digit NAICS level. The resulting panel data set was then merged with the annual 4-digit NAICS data on (nominal) value-added output provided by the Bureau of Economic Analysis; these data were then used to calculate trade exposure—the sum of imports and exports relative to output—for each 4-digit NAICS industry.

21. The industry-level data exhibit significant seasonal fluctuations. Accordingly, we filtered all industry-level variables by using the Census Bureau's X12 seasonal adjustment procedure—thus all of our growth rates (i.e., log differences) are constructed by using seasonally adjusted level series. To ensure that our results were not influenced by a small number of extreme observations, all quarterly growth rates were winsorized at the 0.5th and 99.5th percentiles.

22. These data were assembled by Robert Feenstra through the project funded by a grant from the National Science Foundation to the NBER; see http://cid.econ.ucdavis.edu/usix.html for further details.

Figure 7. Industry-Specific Producer Prices and Industrial Production

A. Producer prices

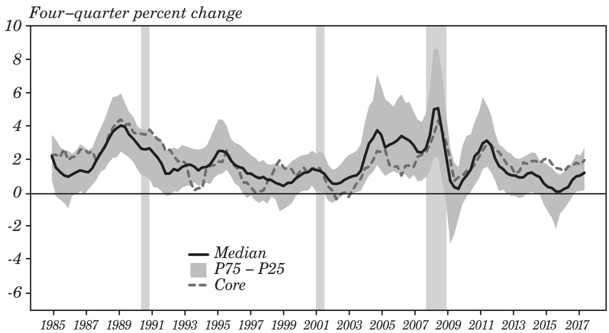

B. Industrial production

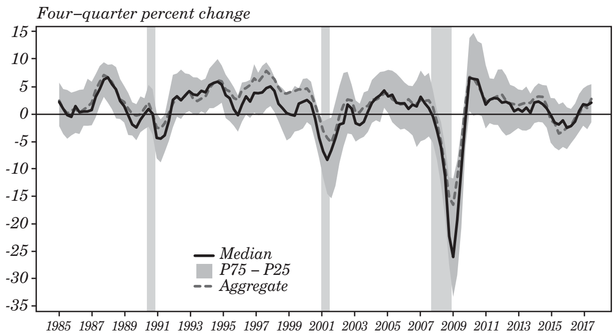

Source: Authors' calculations using data from the Bureau of Labor Statistics and the Federal Reserve Board.
Note: All industry-specific producer price and industrial production indices are seasonally adjusted. The solid lines depict the cross-sectional medians of the specified series, while the shaded bands depict the corresponding interquartile (P75–P25) ranges. For comparison purposes, the dashed line in panel A shows the four-quarter log-difference of the published core PPI, while the dashed line in panel B shows the four-quarter log-difference of the published IPI. The shaded vertical bars denote the NBER-dated recessions.

The solid line in panel A of figure 7 shows the time-series evolution of the (unweighted) cross-sectional median of the four-quarter percent change in PPI inflation across 319 industries in our unbalanced panel, while the shaded band depicts the corresponding (unweighted) interquartile range. The dashed line, in contrast, shows the four-quarter percent change based on the published core PPI. In panel B, the solid line and the shaded band depict the same moments of the four-quarter percent change in industrial production across the same set of industries, while the dashed line shows the corresponding growth rate of total industrial production. As evidenced by the shaded bands, the inflation rates and output growth vary significantly across industries. At the same time, the time-series fluctuations in the two medians closely match dynamics of their corresponding aggregates, an indication that our industry-level data are representative of the economy as a whole.

2.2 Baseline Estimates

To analyze the relationship between producer prices and economic activity at the industry level, we reformulate our baseline Phillips curve specification given by equation (1) above to accommodate the cross-sectional aspect of the industry-level data. Specifically, we estimate the following panel-data version of the Phillips curve:

$$\Delta_{h+1} p_{i,t+h} = \lambda \mathrm{gap}_{it} + \sum_{s=1}^{4} \phi_s \Delta p_{i,t-s} + \mu_i + \eta_t + \epsilon_{i,t+h}, \tag{3}$$

where $p_{i,t}$ denotes the logarithm of the PPI for industry i in quarter t and gap_{it} is a measure of economic slack (or activity) in that industry. This specification also allows for an industry-specific intercept μ_i that is estimated using industry fixed effects and a full set of time dummies—denoted by η_t, $t = 1,2,\ldots,T$—that capture variation in common factors across industries. To measure the extent of resource utilization within each industry, we compute the "industrial production" gaps for each industry—denoted by $[q_{it} - \tilde{q}_{it}]$—as (100 times) the log-deviation of IPI (q_{it}) from its stochastic trend ($\tilde{q}_{it}$), where the latter is estimated by using the Hamilton (2018) filter. As an alternative, we also consider a simple four-quarter log-difference of industrial production, denoted by $\Delta_4 q_{it}$.

Columns (1) and (2) of table 5 report estimates of the Phillips curve at the four-quarter horizon (i.e., $h = 4$) for the full sample of industries from 1984:Q1 to 2017:Q4. Columns (3) and (4), on the other hand, provide comparable estimates for a subsample based on the 1998:Q1–2017:Q4 period, which corresponds to the time period in which the slope of the aggregate Phillips curve for PPI inflation is estimated to have stabilized near zero (figure 4).[23] According to

23. Because our panel data set is unbalanced, the coefficient estimates are not strictly comparable across these two periods.

columns (1) and (2), fluctuations in economic activity—measured either as deviations of industrial output from its trend or as four-quarter growth in output—are important determinants of producer price inflation at the industry level. Although precisely estimated, the economic magnitudes of these coefficients are fairly small: An increase in the industrial production gap of 10 percentage points in quarter t—an increase of a bit less than one standard deviation—is estimated to boost annualized PPI inflation from quarter $t-1$ to $t+4$ a mere 15 basis points; the same-sized increase in the four-quarter growth of industrial output leads to a rise in PPI inflation of about a quarter of a percentage point over the same horizon.

While small in economic terms, these estimates are nonetheless broadly consistent with those based on the aggregate time-series data. For example, the coefficient on the output gap in the aggregate Phillips curve for core producer prices estimated over the 1984:Q1–2017:Q4 period is 0.063 (std. error = 0.054), while the corresponding coefficient estimate based on the 1998:Q1–2017:Q4 sample is –0.021 (std. error = 0.056).

Table 5. Industry-Level Phillips Curve

Explanatory variables	*Sample: 1984:Q1-2017:Q4* (1)	(2)	*Sample: 1998:Q1-2017:Q4* (3)	(4)
$[q_{it} - \tilde{q}_{it}]$	0.014**	-	0.020***	-
	(0.006)		(0.007)	
$\Delta_4 q_{it}$	-	0.027***	-	0.030***
		(0.008)		(0.008)
Sum: inflation lags[a]	-0.057*	-0.054*	-0.082**	-0.079**
	(0.031)	(0.030)	(0.037)	(0.037)
Adj. R^2	0.220	0.222	0.246	0.246
Panel dimensions				
No. of industries	319	319	319	319
Avg. T_i (quarters)	95.6	95.8	60.4	60.5
No. of observations	30,512	30,566	19,266	19,287

Source: Authors' calculations.

Note: The dependent variable in each Phillips curve specification is $\Delta_5 p_{i,t+4}$, the annualized log-difference in industry-specific PPI from date $t-1$ to date $t+4$. Explanatory variables: $[q_{it} - \tilde{q}_{it}]$ = industry-specific industrial production gap, and $\Delta_4 q_{it}$ = log-difference in industry-specific IPI from date $t-4$ to date t. All specifications include industry and time fixed effects and lags 1,...,4 of Δp_{it} (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are clustered across industries and time, according to Cameron and others (2011): * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

[a] Sum of coefficients on $\Delta p_{i,t-s}$, $s = 1,\ldots,4$.

In other words, the slope of the aggregate Phillips curve for core PPI inflation is statistically indistinguishable from zero over this period. It is also worth noting that the estimates of coefficients on economic activity reported in table 5 are remarkably stable across the two sample periods. Thus, the industry-level estimates of the response of PPI inflation to fluctuations in industrial output do not show the same kind of attenuation pattern that we estimate by using the aggregate time-series data.

2.3 The Role of the Trade Share

With these results in hand, we now turn to the question of whether differences in external trade exposure across industries influence the sensitivity of PPI inflation to economic slack. A straightforward way to test this hypothesis would be to estimate our baseline industry-level Phillips curve given in equation (3) on a sample of "low" trade intensity industries and compare the results with those based on a sample of "high" trade intensity industries. However, to make a statement of whether differences in trade exposure across industries matter in the aggregate, we must specify some kind of a weighting scheme.[24] Unfortunately, the value of shipments, which would provide an economically most sensible weighting scheme for the industry-specific inflation rates, is not available at the 6-digit NAICS level. As an alternative, we rely on the QCEW employment data and aggregate the industry-specific PPI inflation rates using the industry-specific *average* employment shares as weights. Because the employment data are available only starting in 1990:Q1, we restrict the analysis to the balanced panel of 185 industries, which ensures that our aggregation scheme is not affected by changes in the composition of industries over time.

To gauge the reasonableness of our aggregation scheme, the solid line in figure 8 shows the time-series evolution of a weighted cross-sectional average of four-quarter PPI inflation rates across the 185 industries in our balanced panel, while the dashed and dashed-dotted lines show the corresponding behavior of the headline and core producer price inflation, respectively. As can be seen from the figure,

24. Note that in the above regression analysis, each industry received an equal weight. As such, the results in table 5 may not provide an accurate picture of the aggregate relationship between inflation and economic slack that is central to our analysis.

our employment-weighted aggregate inflation broadly tracks a mix of the headline and core PPI inflation. It is clearly more cyclical than the core inflation and somewhat less cyclical than the headline inflation. Importantly, this aggregation exercise gives us confidence that an employment-weighted version of the 6-digit industry data captures the cyclical variation that we see in other time-series aggregates and hence provides a meaningful laboratory from which one can infer aggregate phenomena from the industry-level estimates.

We use the balanced panel—with the associated average employment shares—to examine the extent to which the responsiveness of inflation to fluctuations in economic activity differs with the degree of trade intensity across industries. As noted above, we split our sample of 185 industries into two groups, based on whether their average trade share is above or below 5 percent. This cutoff corresponds to the median of the industry-specific average trade shares, weighted by the industry-specific average employment shares, and implies that the low and high trade intensity industry groups account for about one-half each of total employment in our balanced panel.

Figure 8. Industry vs. Aggregate Producer Price Inflation

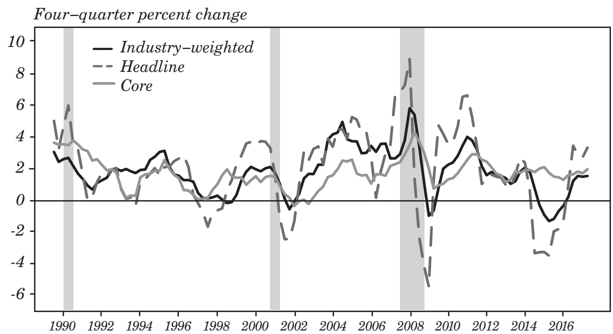

Source: Authors' calculations using data from the Bureau of Labor Statistics.
Note: The solid line depicts a cross-sectional weighted average of producer price inflation across 185 industries in the balanced panel, with weights equal to the corresponding average industry-specific employment shares. The dashed grey line depicts the headline (core) producer price inflation. The shaded vertical bars denote the NBER-dated recessions.

Table 6. Industry-Level Phillips Curve and the Trade Share
(weighted vs. unweighted estimates)

Explanatory variables	*Industry category*		
	All	*Low trade shr.*	*High trade shr.*
A. Weighted estimates			
$[q_{it} - \tilde{q}_{it}]$	0.015	0.029***	0.006
	(0.010)	(0.011)	(0.011)
Sum: inflation lags[d]	−0.060	−0.159***	0.044
	(0.041)	(0.043)	(0.043)
Adj. R^2	0.243	0.228	0.306
B. Unweighted estimates			
$[q_{it} - \tilde{q}_{it}]$	0.025***	0.035***	0.014**
	(0.007)	(0.013)	(0.006)
Sum: inflation lags[d]	−0.060	−0.091**	0.004
	(0.036)	(0.042)	(0.045)
Adj. R^2	0.198	0.198	0.227

Source: Authors' calculations.
Note: Sample: a balanced panel of 185 industries from 1990:Q1 to 2017:Q4 (Obs. = 19,239). The dependent variable in each Phillips curve specification is $\Delta_5 p_{i,t+4}$, the annualized log-difference in industry-specific PPI from date $t-1$ to date $t+4$. Explanatory variables: $[q_{it} - \tilde{q}_{it}]$ = industry-specific industrial production gap. All specifications include industry and time fixed effects and lags 1,...,4 of Δp_{it} (not reported). In panel A, the specifications are estimated by WLS—using average industry employment shares as weights—while in panel B, they are estimated by OLS. Asymptotic standard errors reported in parentheses are clustered across industries and time, according to Cameron and others (2011): * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.
[a] Sum of coefficients on $\Delta p_{i,t-s}$, $s = 1,\dots,4$.

Table 6 reports the results of this exercise for inflation at the four-quarter horizon (i.e., $h = 4$) and using the industrial production gap, $[q_{it} - \tilde{q}_{it}]$, to measure slack at the industry level. In the first column of panel A, we report the weighted least squares (WLS) estimates of the coefficient on the industrial production gap for all industries, while in the second and third column, we report the corresponding WLS estimates for low and high trade share industry groupings, respectively; for comparison purposes, panel B contains the corresponding ordinary least squares (OLS) estimates, which weight all industries equally.

The WLS estimate of the coefficient on economic slack for all industries is a bit smaller than its corresponding OLS estimate—0.015 vs. 0.025—and also less precisely estimated. More importantly, the WLS estimates of coefficients on economic slack show a clear difference across the two industry groupings: In low trade intensity industries, the coefficient on economic slack is positive and statistically highly significant, whereas in high trade intensity industries, the coefficient on economic slack is essentially zero, in both economic and statistical terms. These results provide further support for the argument that globalization and increased international trade may be responsible, at least in part, for the observed attenuation in the response of inflation to fluctuations in economic activity. However, swings in producer prices at the industry level are far more likely to reflect a confluence of demand shocks, which push prices and output in the same direction, and supply shocks, which push them in opposite directions. Thus one should be cautious in providing a structural interpretation to the coefficient estimates reported in table 6.

3. Trade Share and the Effects of Aggregate Shocks

In this section, we employ an alternative approach to investigate the role that international trade may play in determining domestic inflation outcomes. Specifically, we identify aggregate shocks that simultaneously influence inflation and output dynamics and trace out their effects on industry-level outcomes. We then examine the extent to which the industry-level responses of prices, wages, output, and employment to such aggregate shocks differ across industries with a differential exposure to international trade and thus to global factors.

3.1 Econometric Methodology

As in the previous section, we focus on a balanced panel of 185 industries for which all variables are available over the 1990:Q1–2017:Q4 sample period. Given the high dimensionality—in both the cross-sectional and time-series dimensions—of our industry-level data, we use the FAVAR methodology proposed by Bernanke and Boivin (2003) and Bernanke and others (2005) to identify aggregate shocks and trace out their effect on price and wage inflation and the growth of output and employment at the industry level. To identify an aggregate shock of interest, we study the response of industry-level variables to a sudden deterioration in broad domestic financial conditions. An adverse shock to financial conditions may be interpreted as a reduction in aggregate demand and such shocks have featured prominently in recent discussions regarding the source of business cycle fluctuations over the time period under our consideration.[25]

25. See Stock and Watson, 2012.

Our estimation and identification procedure broadly follows the empirical methodology outlined in Gilchrist and others (2009). In particular, we combine the industry-level data on price and wage inflation and on the growth of output and employment in an $(n_1 \times 1)$-dimensional vector X_{1t}.[26] We then consider a set of macro-level variables that summarize domestic financial conditions—these series are combined in an $(n_2 \times 1)$-dimensional vector X_{2t}. This data-rich environment can be succinctly represented by an $(n \times 1)$-dimensional vector $X_t = [X'_{1t}, X'_{2t}]'$, where $n = n_1 + n_2$ and $t = 1,2,\ldots,T$. We assume that X_t has a (linear) factor structure, whereby $X_{it} = \lambda_i' F_t + v_{it}$, $i = 1,\ldots,n$, where F_t is a $(k \times 1)$-dimensional vector of common latent factors (with $k \ll n$), λ_i is the corresponding vector of factor loadings, and v_{it} is an idiosyncratic random disturbance that is assumed to be uncorrelated across i and t.

When analyzing the dynamic effects of aggregate financial shocks, we assume that a subset of these common factors—denoted by a $(k_2 \times 1)$-dimensional vector F_{2t}—are factors that are specific to the aggregate financial variables contained in the vector X_{2t}. These factors do not contemporaneously influence the industry-level variables in the vector X_{1t}, but they do affect contemporaneously the variables in the vector X_{2t}. The rest of the factors—denoted by a $(k_1 \times 1)$-dimensional vector F_{1t}, where $k = k_1 + k_2$—are assumed to span the information contained in the entire data vector X_t. The relationship between the observed variables and the unobserved factors is assumed to be linear and is given by the following system of measurement equations:

$$\begin{bmatrix} X_{1t} \\ X_{2t} \end{bmatrix} = \begin{bmatrix} \Lambda_{11} & \Lambda_{12} \\ \Lambda_{21} & \Lambda_{22} \end{bmatrix} \begin{bmatrix} F_{1t} \\ F_{2t} \end{bmatrix} + \begin{bmatrix} v_{1t} \\ v_{2t} \end{bmatrix}, \tag{4}$$

where

$$\Lambda = \begin{bmatrix} \Lambda_{11} & \Lambda_{12} \\ \Lambda_{21} & \Lambda_{22} \end{bmatrix}$$

is an $(n \times k)$ matrix of factor loadings.

The latent factors are assumed to follow a vector autoregressive process of the form:

$$\begin{bmatrix} F_{1t} \\ F_{2t} \end{bmatrix} = \Phi(L) \begin{bmatrix} F_{1,t-1} \\ F_{2,t-1} \end{bmatrix} + \begin{bmatrix} \in_{1t} \\ \in_{2t} \end{bmatrix}, \tag{5}$$

26. Note that $n_1 = 4 \times 185 = 740$; that is, four series for each of the 185 industries. Wage inflation is measured as the log-difference in the average weekly earnings.

where $\Phi(L)$ is a matrix lag-polynomial of finite order p. As it is standard in these models, we assume that $E[v_{it} \in_{st}] = 0$, for all $i = 1,2,\ldots,n$ and $s = 1,2,\ldots,k$; and $E[\in_{it} \in_{jt}] = 0$, for all $i \neq j$. In this form, our model constitutes a static representation of a dynamic factor model;[27] it is static in the sense that factors enter only contemporaneously in the system of measurement equations (4).

To identify the aggregate factors F_{2t}, we impose the following restrictions on the system of measurement equations. First, we assume that the matrix $\Lambda_{12} = 0$. This restriction on the factor loading matrix Λ implies that, once we have conditioned on the factors F_{1t}, the remaining variation in the aggregate block X_{2t} has a systematic component that is reflected in its own factor structure. Although the aggregate factors F_{2t} have no contemporaneous effect on the vector X_{1t}, they affect the factors F_{1t} and, by extension, the variables in the industry block X_{1t} with a lag through the autoregressive dynamics of equation (5). The second identifying assumption is that the contemporaneous innovations associated with the factors F_{1t} and F_{2t} are orthogonal, an assumption that separates the residual information content in the aggregate block from the factors summarizing the state of the economy, as measured by the full set of industry-specific information contained in the vector X_{1t}.[28]

In implementing this identification strategy, we let the vector X_{2t} include a broad array of domestic financial indicators. Specifically, when considering how financial shocks affect industry-level outcomes, the vector X_{2t} consists of the following five financial indicators: the GZ corporate bond credit spread and the associated excess bond premium;[29] the Moody's Baa-Aaa corporate bond credit spread; the term spread measured as the difference in yields on the ten- and two-year U.S. Treasury coupon securities; and the option-implied volatility on the S&P 500 stock price index, the VIX. The GZ and Baa-Aaa credit spreads and the excess bond premium are widely used indicators of financial strains obtained from the corporate bond market. The VIX, on the other hand, is a measure of risk appetite in equity markets, while

27. See Stock and Watson, 2010a.

28. We can estimate the FAVAR model given by equations (4) and (5) by using a Gaussian maximum likelihood method and a Kalman filter to construct the likelihood function. However, in the presence of identifying assumptions with large n, this method is computationally demanding. We, therefore, follow the four-step procedure outlined in Gilchrist and others (2009), as it is straightforward to implement and directly imposes the necessary identification restrictions.

29. See Gilchrist and Zakrajšek, 2012.

the term spread primarily reflects investors' appetite for duration risk. These five indicators provide a broad summary of domestic financial conditions that are entirely reliant on market prices and, therefore, should capture changes in broad financial conditions in a timely manner. In the FAVAR specifications, we allow for four common factors in the industry-level block X_{1t}—that is, $k_1 = 4$— and for one factor in the aggregate block X_{2t}, that is $k_2 = 1$.[30]

3.2 The Impact of Financial Shocks

In this section, we present impulse responses of variables in the industry block X_{1t} to the identified aggregate financial shock. We begin by reporting these baseline results for all industries. Next, we examine how international trade exposure influences industry-level inflation dynamics by again dividing our sample of industries into those with a low external trade exposure and those with a high trade exposure.

Figure 9 plots the distribution of industry-level outcomes in response to an adverse financial shock of one standard deviation in quarter zero. Though not shown, this shock causes a broad-based tightening of domestic financial conditions, implying an increase in the excess bond premium of about 30 basis points upon impact.[31] The solid line in each panel shows the median industry response of the specified variable to such a shock, while the dark shaded bands denote the range of responses between the 75th and 25th percentiles (the $P75$–$P25$ range) and the light shaded bands denote the range of responses between the 95th and 5th percentiles (the $P95$–$P5$ range). Recall that the factor F_{2t} is, by assumption, contemporaneously orthogonal to the variables in the industry block and thus aggregate shocks have no effect on industry-level outcomes upon impact.

30. These choices were based on the information criteria proposed by Bai and Ng (2002); however, all of the results reported in this paper are robust to allowing a greater number of factors in either block.

31. Over the 1990:Q1–2017:Q4 period, the standard deviation of the excess bond premium is about 50 basis points. As a point of comparison, the excess bond premium shot up more than 300 basis points following the collapse of Lehman Brothers in September 2008.

Figure 9. Implications of an Adverse Financial Shock
(all industries)

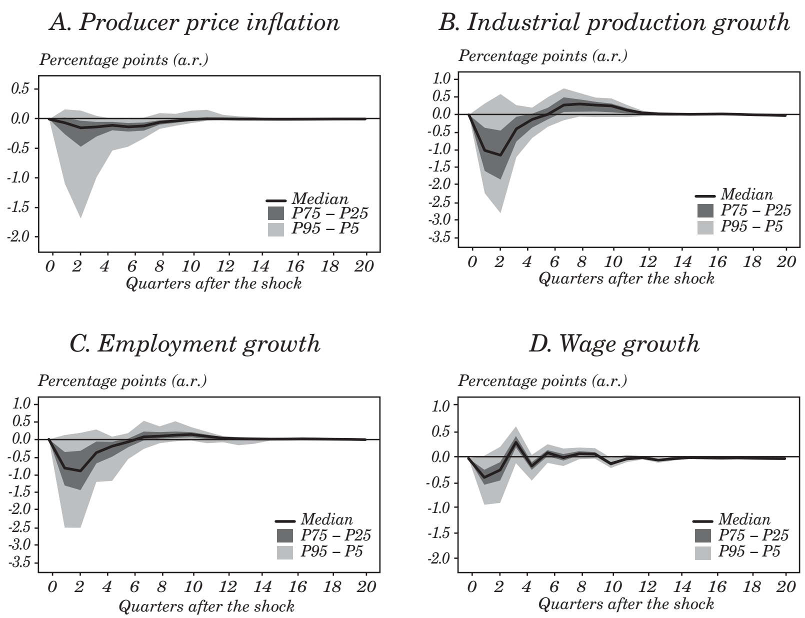

Source: Authors' calculations.
Note: The solid line in each panel depicts the median response of the specified variable to an adverse financial shock of one standard deviation across 185 industries; the shaded bands depict the corresponding *P75–P25* and *P95–P5* ranges. See the text for details.

The identified financial shock is clearly contractionary—it induces a substantial decline in the growth of industrial production and employment for the median industry. A couple of quarters after its impact, this shock is cutting 1.2 percentage points from the annualized growth of output and 0.8 percentage points from the annualized growth of employment at the median. It also causes a significant step-down in both price and wage inflation: For the median industry, annualized price inflation is lowered 0.1 percentage points, whereas the reduction in annualized wage inflation is on the order of 0.4 percentage points. Notably, the reduction in the rate of growth of economic activity, prices, and wages occurs relatively quickly, peaking a mere two quarters after the shock. Economic growth remains depressed for several more quarters before recovering slowly and returns to its long-run level only after about eight quarters. Judging by the shaded regions, the estimated range of industry responses implies that these effects are broad based. The combination of steep declines in the growth of output, employment, prices, and wages implies that the deterioration in broad domestic financial conditions delivers a response that is consistent with a reduction in aggregate demand within a New Keynesian framework.

With these baseline results in hand, we now analyze the extent to which differential trade exposure across industries changes the results reported above. As before, we sort our sample of industries based on their average trade exposure over the 1990:Q1–2017:Q4 period and group them into a low and high trade exposure categories. We then separately estimate our two FAVAR specifications for each of the two groupings, an approach that ensures that we do not artificially constrain the factor structure to be the same across industries with a differential trade exposure. As a reminder, recall that each category of industries accounts, on average, for about 50 percent of total employment in our sample.

Unlike our baseline exercise, this exercise is focused on the implications of the common financial shock for aggregate outcomes. Specifically, for each industry-level endogenous variable, we compute a weighted-average response across industries, where weights are equal to the industry-specific average employment shares within each group of industries (i.e., low vs. high trade exposure industry categories). In addition, we report the aggregate responses for all industries by weighting the industry-specific responses from figure 9 with their corresponding average employment shares; these results are shown in figure 10.[32]

As shown by the solid lines in figure 10, the aggregate responses of producer price and wage inflation and the growth of output and employment to an adverse financial shock follow closely the contours of the corresponding median industry-level responses shown in figure 9: Price and wage inflation, along with output and employment growth, all fall sharply, with peak responses occurring one to two quarters after the impact of the shock. Moreover, these aggregate responses remain persistently below their respective long-run values for six to eight quarters after the shock.

32. Figures B.1–B.2 in appendix B show the industry-level responses for the low and high trade share industry categories when the economy is perturbed by an aggregate financial shock.

Figure 10. Implications of an Adverse Financial Shock
(low vs. high trade share industries)

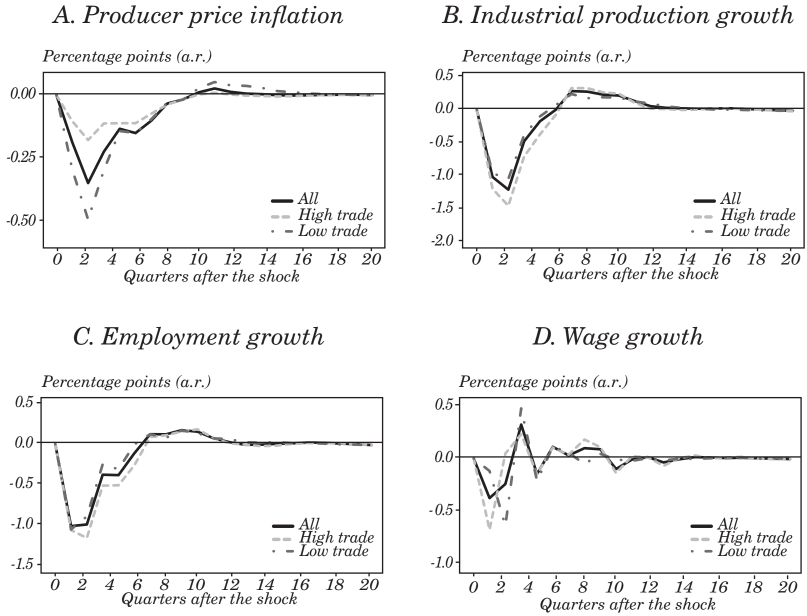

Source: Authors' calculations.
Note: The solid line in each panel depicts the employment-weighted-average response of the specified variable to an adverse financial shock of one standard deviation across 185 industries; the dashed (dashed-dotted) lines depict the corresponding employment-weighted-average responses for a subset of industries with a high (low) average trade share. See the text for details.

Note that the (absolute) magnitude of responses for the aggregates—as defined by the employment-weighted averages of industry-level responses—are somewhat larger than their corresponding unweighted median responses across industries. In particular, the annualized output and employment growth both fall by more than one percentage point, while the annualized producer price inflation declines about 30 basis points. The estimated decline in the growth of output in response to a financial shock is consistent with other studies that find that such disturbances lead to a significant contraction in economic activity.[33] That said, the estimated drop in producer price inflation is both larger and

33. See Gilchrist and others, 2009; Gilchrist and Zakrajšek, 2012; and Boivin and others, 2018.

occurs more quickly than the one estimated by VARs that use aggregate time-series data. Overall, these results indicate that producer price inflation is fairly sensitive to fluctuations in economic activity induced by changes in broad financial conditions—producer price inflation declines roughly 25 basis points when a tightening of financial conditions induces a one percentage point decline in the growth of industrial output.

Figure 10 also displays the aggregate responses to an adverse financial shock for high and low trade industries. As shown in the upper left panel, the dynamics of inflation differ markedly across industries with a differential trade exposure. Notably, the peak decline in producer price inflation of 0.5 percentage points for industries with low trade exposure is more than three times as large as that for industries with high trade exposure. Although the unanticipated tightening of financial conditions causes a somewhat greater contraction in economic activity among high trade industries, the responses of output and employment growth are broadly similar—in terms of both timing and their magnitudes—across the two industry groupings. Wage inflation also behaves in a similar manner across these two industry groupings, though in high trade industries, the deceleration in wages occurs more quickly.

A useful way to highlight the difference in inflation dynamics between low and high trade industries is to compute the cumulative responses of price inflation and output growth. The ratio of the resulting price response to the output response then provides an estimate of the decline in prices relative to output that occurs at different horizons in response to an adverse financial shock. As shown in figure 11, in low trade intensity industries, producer prices are estimated to decline about 0.3 percent for every one percent decline in output at very short horizons and about one percent for the same-sized reduction in output at the two-year horizon. In high trade intensity industries, by contrast, producer prices are estimated to decline about 0.1 percent for a one percent reduction in output at very short horizons and about 0.3 percent at the two-year horizon. In sum, these findings imply that the inflation-output tradeoff is—at every horizon—three times larger in low trade intensity industries than in their high trade intensity counterparts.

Figure 11. Inflation-Output Tradeoff

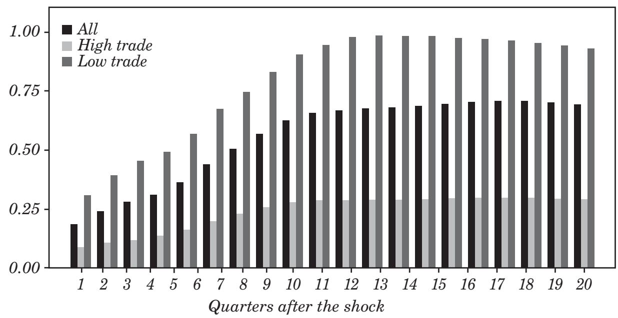

Source: Authors' calculations.
Note: The bars in the figure depict the estimated sensitivity of producer prices to fluctuations in output induced by aggregate financial shocks. See the text for details.

In summary, our FAVAR analysis implies that producer price inflation is three to four times more responsive to aggregate demand shocks in low trade intensity industries than their high trade intensity counterparts. Responses of wages, output, and employment, by contrast, are strikingly similar across the two industry groupings. These results are consistent with the notion that the Phillips curve is indeed much flatter in industries that are more exposed to international trade and are thus broadly consistent with our earlier findings, which show that the estimated flattening of the aggregate Phillips curves coincides to a substantial degree with the increased exposure of the U.S. economy to international trade.

4. Conclusions

In this chapter, we examine the extent to which the response of inflation to fluctuations in economic activity has weakened over time. Furthermore, we analyze the role of globalization and rising trade shares behind these structural shifts. Our evidence points to a significant flattening of the Phillips curve that occurred in the 1990s. Although there is some evidence of a recent rise in the responsiveness of CPI inflation to changes in economic slack, it remains the case that both PPI and CPI inflation are substantially less responsive to fluctuations in economic activity today, relative to estimates that rely on the pre-1990 data. To a significant degree, this reduced responsiveness of inflation to economic slack coincides with a rising U.S. trade share and a concomitant increase in global economic integration.

Industry-level data provide further evidence in favor of the notion that trade intensity attenuates the response of inflation to fluctuations in economic activity. Industry-level estimates of the Phillips curve imply a substantially lower sensitivity of PPI inflation to output in industries with a high trade share, relative to those with a low trade share. We confirm these results by examining the response of industry-level PPI inflation and output to identified aggregate financial shocks. This evidence implies that the inflation-output tradeoff is about three times larger for low trade intensity industries than for their high trade intensity counterparts. In this sense, increased international trade and globalization do indeed appear to help explain the observed flattening of the aggregate Phillips curve over the past several decades.

References

Andrews, D.W.K. 1993. "Tests for Parameter Instability and Structural Change With Unknown Change Point." *Econometrica* 61: 821–56.

Atkeson, A. and L.E. Ohanian. 2001. "Are Phillips Curves Useful for Forecasting Inflation?" *Federal Reserve Bank of Minneapolis Quarterly Review* (Winter): 2–11.

Auer, R., C. Borio, and A. Filardo. 2017. "The Globalisation of Inflation: The Growing Importance of Global Value Chains." Working Paper No. 602, Bank for International Settlements.

Bai, J. and S. Ng. 2002. "Determining the Number of Factors in Approximate Factor Models." *Econometrica* 70: 191–221.

Ball, L. 2006. "Has Globalization Changed Inflation?" NBER Working Paper No. 12687.

Ball, L.M. and S. Mazumder. 2011. "Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity* 42: 337–405.

———. 2018. "A Phillips Curve with Anchored Expectations and Short-Term Unemployment." *Journal of Money, Credit and Banking*. Available at https://doi:10.1111/jmcb. 12502.

Ball, L.M., N.G. Mankiw, and D.H. Romer. 1988. "The New Keynesian Economics and the Output-Inflation Tradeoff." *Brookings Papers on Economic Activity* 19: 1–82.

Beaudry, P. and M. Doyle 2000. "What Happened to the Phillips Curve in the 1990s in Canada?" In *Proceedings of the Bank of Canada Conference on Price Stability and the Long-Run Target for Monetary Policy*, Bank of Canada.

Berganza, J.C., P. Del Río López, and F. Borrallo. 2016. "Determinants and Implications of Low Global Inflation Rates." Banco de España Occasional Paper No. 1608. Available at *http://dx.doi.org/10.2139/ssrn.2891619*.

Bernanke, B.S. 2007. "Globalization and Monetary Policy." Speech delivered at the Fourth Economic Summit, Stanford Institute for Economic Policy Research, Stanford, CA. Available at *https://www.federalreserve. gov/newsevents/speech/bernanke20070302a.htm*.

———. 2010. "The Economic Outlook and Monetary Policy." Speech delivered at the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming, 27 August. Available at *http://www.federalreserve. gov/newsevents/speech/bernanke20100827a.htm*.

Bernanke, B.S. and J. Boivin. 2003. "Monetary Policy in a Data-Rich Environment." *Journal of Monetary Economics* 50: 525–46.

Bernanke, B.S., J. Boivin, and P. Eliasz. 2005. "Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach." *Quarterly Journal of Economics* 120: 387–422.

Blanchard, O.J. 2016. "The Phillips Curve: Back to the '60s?" *American Economic Review* 106: 31–4.

———. 2018. "Should We Reject the Natural Rate Hypothesis?" *Journal of Economic Perspectives* 32: 97–120.

Boivin, J., M.P. Giannoni, and D. Stevanović. 2018. "Dynamic Effects of Credit Shocks in a Data-Rich Environment." *Journal of Business and Economic Statistics* 25(1): 52–60. Available at *https://doi.org/ 10.1080/07350015.2018.1497507*.

Borio, C. and A. Filardo. 2007. "Globalisation and Inflation: New Cross-Country Evidence on the Global Determinants of Domestic Inflation." Working Paper No. 227, Bank for International Settlements.

Cameron, A.C., J.B. Gelbach, and D.L. Miller. 2011. "Robust Inference with Multi-Way Clustering." *Journal of Business and Economic Statistics* 29(2): 238–49.

Elliott, G. and U.K. Müller. 2006. "Efficient Tests for General Persistent Time Variation in Regression Coefficients." *Review of Economic Studies* 73: 907–40.

Forbes, K.J. 2018. "Has Globalization Changed the Inflation Process?" Working Paper, MIT Sloan School of Management.

Friedrich, C. 2016. "Global Inflation Dynamics in the Post-Crisis Period: What Explains the Twin Puzzle?" *Economic Letters* 142: 31–4.

Galí, J. and M. Gertler. 2000. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44: 195–222.

Galí, J., M. Gertler, and D. López-Salido. 2001. "European Inflation Dynamics." *European Economic Review* 45: 1237–70.

———. 2007. "Markups, Gaps, and the Welfare Costs of Business Fluctuations." *Review of Economics and Statistics* 89: 44–59.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek. 2017. "Inflation Dynamics During the Financial Crisis." *American Economic Review* 107; 785–823.

Gilchrist, S., V. Yankov, and E. Zakrajšek. 2009. "Credit Market Shocks and Economic Fluctuations: Evidence From Corporate Bond and Stock Markets." *Journal of Monetary Economics* 56: 471–93.

Gilchrist, S. and E. Zakrajšek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102: 1692–720.

————. 2016. "Customer Markets and Financial Frictions: Implications for Inflation Dynamics." Federal Reserve Bank of Kansas City, *Proceedings – Economic Policy Symposium – Jackson Hole*.

Gordon, R.J. 1982. "Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment." In *Workers, Jobs, and Inflation*, edited by M. Baily. Washington D.C.: The Brookings Institution.

————. 2013. "The Phillips Curve is Alive and Well: Inflation and the NAIRU during the Slow Recovery." NBER Working Paper No. 19390.

Hamilton, J.D. 1983. "Oil and the Macroeconomy Since World War II." *Journal of Political Economy* 91: 228–48.

————. 2018. "Why You Should Never Use the Hodrick-Prescott Filter." *Review of Economics and Statistics* 100: 831–43.

Ihrig, J.E., S.B. Kamin, D. Linder, and J.R. Marquez. 2010. "Some Simple Tests of the Globalization and Inflation Hypothesis." *International Finance* 13; 343–75.

Jermann, U.J. and V. Quadrini. 2012. "Macroeconomic Effects of Financial Shocks." *American Economic Review* 102: 238–71.

Krueger, A.B., J. Cramer, and D. Cho. 2014. "Are the Long-Term Unemployed on the Margins of the Labor Market?" *Brookings Papers on Economic Activity* 48: 229–98.

Kuttner, K.N. and T. Robinson. 2010. "Understanding the Flattening of the Phillips Curve," *North American Journal of Economics and Finance* 21: 110–25.

Lindsey, D.E., A. Orphanides, and R.H. Rasche. 2005. "The Reform of October 1979: How It Happened and Why." *Review, Federal Reserve Bank of St. Louis* 87: 187–235.

Mankiw, N.G., R. Reis, and J. Wolfers. 200.: "Disagreement about Inflation Expectations." In *NBER Macroeconomics Annual*, edited by M. Gertler and K. Rogoff, Cambridge MA: MIT Press.

Mavroeidis, S., M. Plagborg-Møller, and J.H. Stock. 2001. "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve." *Journal of Economic Literature* 52: 124–88.

Miles, D., U. Panizza, R. Reis, and A. Ubide. 2017. "And Yet It Moves: Inflation and the Great Recession." Geneva Reports on the World Economy 19. The International Center for Monetary and Banking Studies and the Centre for Economic Policy Research.

Nakamura, E. and J. Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123: 1415–64.

Newey, W.K. and K.D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation-Consistent Covariance Matrix." *Econometrica* 55: 703–08.

Orphanides, A. and J.C. Williams. 2013. "Monetary Policy Mistakes and the Evolution of Inflation Expectations." In *The Great Inflation: The Rebirth of Modern Central Banking*, edited by M. D. Bordo and A. Orphanides, Chicago, IL: University of Chicago Press.

Peersman, G. and W. Wagner. 2014. "Shocks to Bank Lending, Risk-Taking, Securitization, and their Role for U.S. Business Cycle Fluctuations." CEPR Discussion Paper No. 10547.

Pfajfar, D. and J.M. Roberts. 2018. "The Role of Expectations in Changed Inflation Dynamics." Finance and Economics Discussion Series Paper No. 2018-062, Federal Reserve Board.

Roberts, J.M. 1995. "New Keynesian Economics and the Phillips Curve." *Journal of Money, Credit, and Banking* 25: 975–84.

————. 2006. "Monetary Policy and Inflation Dynamics," *International Journal of Central Banking* 2: 193–230.

Sbordone, A.M. 2002. "Prices and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49: 265–92.

Simon, J., T. Matheson, and D. Sandri. 2013. "The Dog that Didn't Bark: Has Inflation Been Muzzled or Was It Just Sleeping?" In *World Economic Outlook: Hopes, Realities, and Risks*, Washington D.C.: International Monetary Fund.

Sims, C.A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50: 665–90.

Stock, J.H. and M.W. Watson. 2009. "Phillips Curve Inflation Forecasts." In *Understanding Inflation and the Implications for Monetary Policy*, edited by J. Fuhrer, Y. Kodrycki, J. Little, and G. Olivei, Cambridge, MA: MIT Press.

————. 2010a. "Dynamic Factor Models." In *Oxford Handbook of Economic Forecasting*, edited by M.P. Clements and D.F. Henry, Oxford, UK: Oxford University Press.

———. 2010b. "Modeling Inflation after the Crisis." Federal Reserve Bank of Kansas City, *Proceedings – Economic Policy Symposium – Jackson Hole*.

———. 2012. "Disentangling the Channels of the 2007–09 Recession." *Brookings Papers on Economic Activity* 44: 81–135.

———. 2018. "Slack and Cyclically Sensitive Inflation." Working Paper, Harvard University.

Yellen, J.L. 2013. "Panel Discussion on 'Monetary Policy: Many Targets, Many Instruments. Where Do We Stand?'" Speech delivered at the "Rethinking Macro Policy II," a conference sponsored by the International Monetary Fund, Washington, D.C., 16 April. Available at *http://www.federalreserve.gov/newsevents/speech/yellen10130416a.htm*.

Zhang, C. 2017. "The Great Globalization and Changing Inflation Dynamics." *International Journal of Central Banking* 13: 191–226.

APPENDIX

A. Controlling for Trend Inflation

As discussed in the main text, both producer and consumer price inflation exhibit significant low frequency variation over our sample period (figure 1). To ensure that this low frequency variation does not affect our baseline time-series estimates of the aggregate Phillips curves, this appendix reports a set of results in which all inflation series were "detrended" to eliminate very low frequency variation. Specifically, following Stock and Watson (2012), we calculated the deviations of each quarterly inflation series from a local mean, where the latter is estimated using a bi-weight kernel with a bandwidth of 100 quarters. As noted by Stock and Watson (2012), these local mean estimates are roughly the same as those computed using a centered moving-average window of ±30 quarters. This approach of eliminating low frequency variation in inflation rates has the desirable feature that it makes no assumption about reversion to the local mean.

As shown in figures A.1 and A.2, the values of these local means change substantially over our sample period. Tables A.1 and A.2 contain estimates of the baseline Phillips curve specifications for producer and consumer price inflation, respectively, which use the detrended inflation data; these estimates are directly comparable with those reported in tables 1 and 2 of the main text, which use the untransformed inflation series. Tables A.3 and A.4, in contrast, use the detrended inflation series to examine the role the trade share in influencing the slope of the Phillips curve, and the results in those tables are directly comparable to those reported in tables 3 and 4 of the main text.

Figure A1. Producer Price Inflation

A. Headline

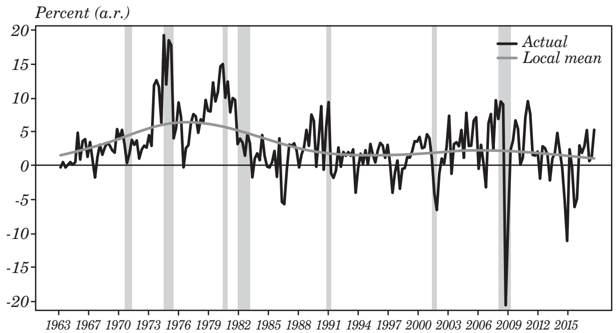

B. Core

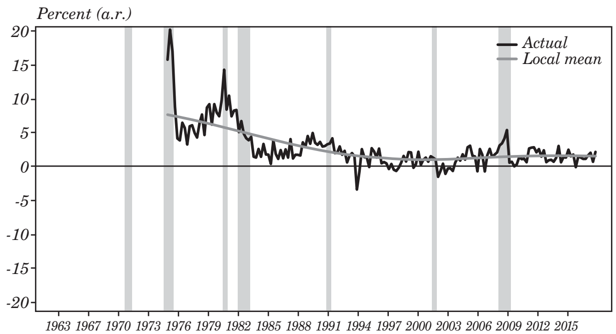

Source: Authors' calculations using data from the Bureau of Labor Statistics.
Note: Panel A depicts the annualized quarterly log-difference of headline PPI inflation and its estimated local mean, while panel B depicts the corresponding series for core PPI inflation (see the text for details). The shaded vertical bars denote the NBER-dated recessions. All price indices are seasonally adjusted.

Figure A2. Consumer Price Inflation

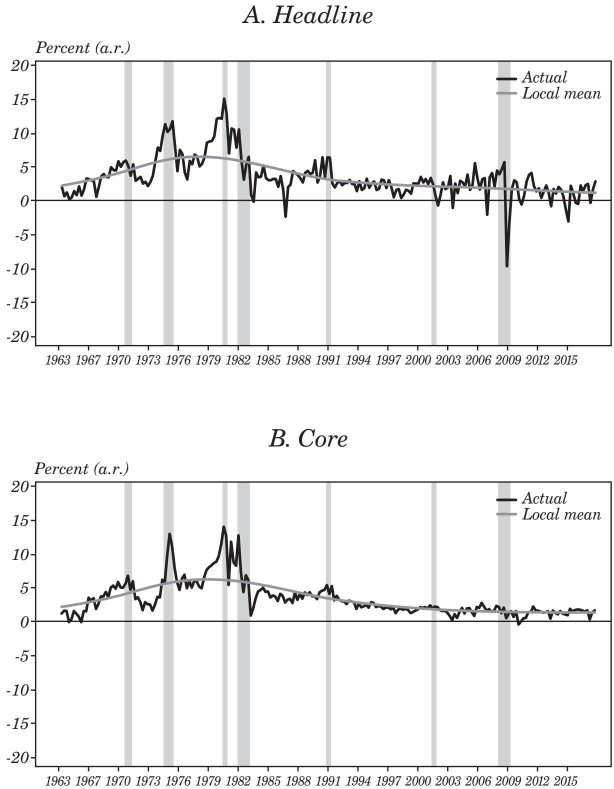

Source: Authors' calculations using data from the Bureau of Labor Statistics.
Note: Panel A depicts the annualized quarterly log-difference of headline CPI inflation and its estimated local mean, while panel B depicts the corresponding series for core CPI inflation (see the text for details). The shaded vertical bars denote the NBER-dated recessions. All price indices are seasonally adjusted.

Table A1. Phillips Curve – Detrended Producer Price Inflation

Explanatory variables	h = 1		h = 4	
	(1)	*(2)*	*(3)*	*(4)*
A. Producer prices				
$[y_t - y_t^*]$	0.263**	-	0.294**	-
	(0.129)	-	(0.135)	-
$[U_t - U_t^*]$	-	-0.306	-	-0.331
	-	(0.224)	-	(0.232)
Sum: inflation lags[a]	0.371**	0.369**	0.212*	0.211*
	(0.147)	(0.147)	(0.112)	(0.112)
sup W[b]	16.168***	115.353***	32.027***	34.338***
	[09:Q1]	[09:Q1]	[84:Q1]	[08:Q1]
q_{LL}[c]	-4.536	-4.975	-3.793	-3.178
Adj. R^2	0.161	0.141	0.160	0.115
B. Core producer prices				
$[y_t - y_t^*]$	0.136**	-	0.163**	-
	(0.052)	-	(0.062)	-
$[U_t - U_t^*]$	-	-0.240**	-	-0.271**
	-	(0.099)	-	(0.122)
Sum: inflation lags[a]	0.565***	0.538***	0.434***	0.403***
	(0.117)	(0.114)	(0.101)	(0.096)
sup W[b]	17.562***	25.350***	60.378***	78.601***
	[84:Q1]	[84:Q1]	[81:Q4]	[82:Q2]
q_{LL}[c]	-7.776*	-7.716*	-5.355	-5.395
Adj. R^2	0.378	0.373	0.409	0.389

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4 for headline PPI (panel A), and 1974:Q1 to 2017:Q4 for core PPI (panel B). The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the detrended annualized log-difference in the specified PPI from date t−1 to date t+h. Explanatory variables: $[y_t - y_t^*]$ = output gap, and $[U_t - U_t^*]$ = unemployment gap. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

a. Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.

b. The Andrews (1993) sup-Wald statistic of the null hypothesis that there is no structural break in the coefficient on economic slack; the estimated break dates are reported in brackets below.

c. The Elliott and Müller (2006) q_{LL} statistic of the null hypothesis that the coefficient on economic slack is constant over time.

Table A2. Phillips Curve – Detrended Consumer Price Inflation

Explanatory variables	h = 1		h = 4	
	(1)	(2)	(3)	(4)
A. Consumer prices				
$[y_t - y_t^*]$	0.196***	-	0.232***	-
	(0.072)	-	(0.075)	-
$[U_t - U_t^*]$	-	-0.264**	-	-0.290**
	-	(0.114)	-	(0.113)
Sum: inflation lags[a]	0.571***	0.550***	0.392***	0.369***
	(0.114)	(0.118)	(0.101)	(0.108)
sup W[b]	22.185***	17.394***	52.617***	30.743***
	[83:Q2]	[83:Q2]	[83:Q1]	[91:Q3]
q_{LL}[c]	-6.173	-6.729	-4.367	-3.846
Adj. R^2	0.339	0.316	0.346	0.290
B. Core consumer prices				
$[y_t - y_t^*]$	0.124***	-	0.179***	-
	(0.045)	-	(0.053)	-
$[U_t - U_t^*]$	-	-0.200***	-	-0.257***
	-	(0.071)	-	(0.085)
Sum: inflation lags[a]	0.677***	0.652***	0.479***	0.441***
	(0.100)	(0.100)	(0.116)	(0.122)
sup W[b]	20.048***	19.449***	65.059***	38.216***
	[83:Q2]	[83:Q2]	[83:Q1]	[83:Q1]
q_{LL}[c]	-6.486	-6.384	-6.068	-5.870
Adj. R^2	0.499	0.495	0.427	0.394

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4. The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the detrended annualized log-difference in the specified CPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap, and $[U_t - U_t^*]$ = unemployment gap. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

a. Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.

b. The Andrews (1993) sup-Wald statistic of the null hypothesis that there is no structural break in the coefficient on economic slack; the estimated break dates are reported in brackets below.

c. The Elliott and Müller (2006) q_{LL} statistic of the null hypothesis that the coefficient on economic slack is constant over time.

Table A3. Phillips Curve and the Trade Share – Detrended Producer Price Inflation

Explanatory variables	*h* = 1		*h* = 4	
	(1)	(2)	(3)	(4)
A. Producer prices				
$[y_t - y_t^*]$	0.702*	-	1.112**	-
	(0.397)	-	(0.494)	-
$[y_t - y_t^*] \times \text{TrdShr}_{t-1}$	-0.022	-	-0.040*	-
	(0.021)	-	(0.023)	-
$[U_t - U_t^*]$	-	-1.033	-	-1.310*
	-	(0.627)	-	(0.739)
$[U_t - U_t^*] \times \text{TrdShr}_{t-1}$	-	0.037	-	0.048
	-	(0.030)	-	(0.033)
Sum: inflation lags[a]	0.382**	0.374**	0.231**	0.217**
	(0.146)	(0.143)	(0.105)	(0.106)
Adj. R^2	0.166	0.147	0.210	0.142
B. Core producer prices				
$[y_t - y_t^*]$	0.636**	-	0.808***	-
	(0.270)	-	(0.281)	-
$[y_t - y_t^*] \times \text{TrdShr}_{t-1}$	-0.022*	-	-0.028**	-
	(0.011)	-	(0.018)	-
$[U_t - U_t^*]$	-	-1.908***	-	-2.239***
	-	(0.529)	-	(0.516)
$[U_t - U_t^*] \times \text{TrdShr}_{t-1}$	-	0.069***	-	0.081***
	-	(0.021)	-	(0.021)
Sum: inflation lags[a]	0.562***	0.490***	0.431***	0.346***
	(0.113)	(0.110)	(0.096)	(0.083)
Adj. R^2	0.402	0.443	0.466	0.519

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4 for headline PPI (panel A), and 1974:Q1 to 2017:Q4 for core PPI (panel B). The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the detrended annualized log-difference in the specified PPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap, $[U_t - U_t^*]$ = unemployment gap, TrdShr_{t-1} = eight-quarter (trailing) moving average of the trade share. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

a. Sum of coefficients on Δp_{t-s}, $s = 1,\ldots,4$.

Table A4. Phillips Curve and the Trade Share – Detrended Consumer Price Inflation

Explanatory variables	*h* = 1		*h* = 4	
	(*1*)	(*2*)	(*3*)	(*4*)
A. Consumer prices				
$[y_t - y_t^*]$	0.505**	-	0.772***	-
	(0.227)	-	(0.275)	-
$[y_t - y_t^*] \times \text{TrdShr}_{t-1}$	-0.015	-	-0.028**	-
	(0.011)		(0.012)	
$[U_t - U_t^*]$	-	-0.717**	-	-0.867**
	-	(0.320)	-	(0.355)
$[U_t - U_t^*] \times \text{TrdShr}_{t-1}$	-	0.022	-	0.028*
	-	(0.015)	-	(0.015)
Sum: inflation lags[a]	0.602***	0.566***	0.447***	0.390***
	(0.117)	(0.116)	(0.173)	(0.105)
Adj. R^2	0.348	0.324	0.398	0.313
B. Core consumer prices				
$[y_t - y_t^*]$	0.315**	-	0.576***	-
	(0.146)	-	(0.186)	-
$[y_t - y_t^*] \times \text{TrdShr}_{t-1}$	-0.009*	-	-0.019***	-
	(0.006)	-	(0.007)	-
$[U_t - U_t^*]$	-	-0.503**	-	-0.698***
	-	(0.201)	-	(0.242)
$[U_t - U_t^*] \times \text{TrdShr}_{t-1}$	-	0.015*	-	0.021**
	-	(0.008)	-	(0.009)
Sum: inflation lags[a]	0.711***	0.623***	0.549***	0.471***
	(0.102)	(0.099)	(0.112)	(0.118)
Adj. R^2	0.504	0.501	0.465	0.414

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4. The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the annualized log-difference in the specified CPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap, $[U_t - U_t^*]$ = unemployment gap, and TrdShr_{t-1} = eight-quarter (trailing) moving average of the trade share. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

a. Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.

B. Supplementary FAVAR Results

Figure B1. Implications of an Adverse Financial Shock
(industries with a low trade share)

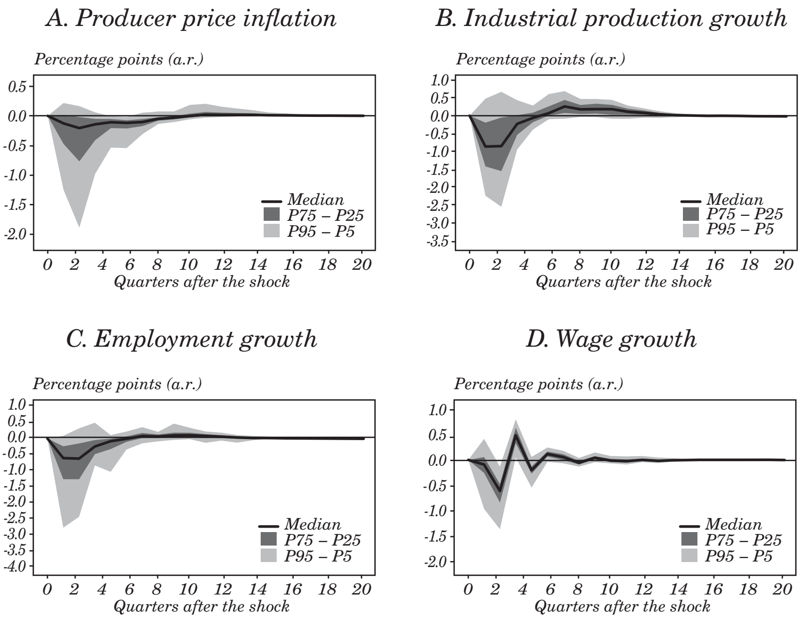

Source: Authors' calculations.
Note: The solid line in each panel depicts the median response of the specified variable to an adverse financial shock of one standard deviation across a subset of industries with a low average trade share; the shaded bands depict the corresponding *P75 – P25* and *P95 – P5* ranges. See the main text for details.

Figure B2. Implications of an Adverse Financial Shock
(industries with a high trade share)

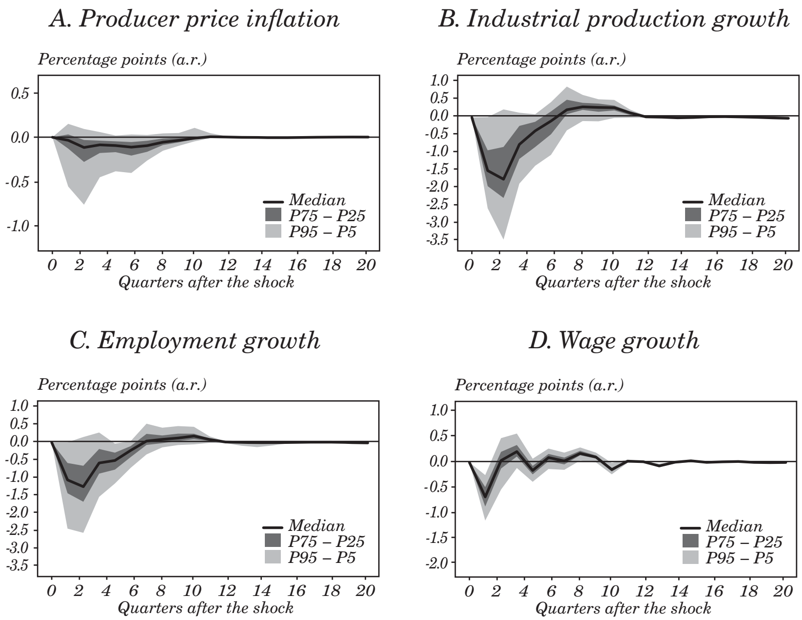

Source: Authors' calculations.
Note: The solid line in each panel depicts the median response of the specified variable to an adverse financial shock of one standard deviation across a subset of industries with a high average trade share; the shaded bands depict the corresponding *P75 – P25* and *P95 – P5* ranges. See the main text for detail.

The Supply-Side Origins of U.S. Inflation

Bart Hobijn
Arizona State University

In recent years, we have not seen much of a negative correlation between inflation, the time series plotted in figure 1, and measures of resource slack, based on real GDP plotted in figure 2. This flattening of the Phillips curve in many countries across the world has startled monetary policymakers. In fact, it has some former policymakers ask whether the Phillips curve is dead.[1] It is often interpreted as the disappearance of a short-run output-inflation tradeoff that central banks can exploit for stabilization purposes.[2]

In this paper I argue that this is too pessimistic an assessment. What the flattening of the Phillips curve really indicates is that recent economic fluctuations were not mainly driven by movements in aggregate demand (AD) but, instead, by joint movements in aggregate demand and aggregate supply (AS). It is these movements in aggregate supply that are at the root of the "supply-side origins of inflation" that I refer to in the title.

In the first part of this paper, I illustrate that, once one is willing to drop the assumption in a textbook aggregate demand-aggregate supply (AD-AS) framework that business-cycle fluctuations are mainly the result of movements in aggregate demand, it is not hard to imagine how joint inward shifts in both aggregate demand and aggregate supply can result in economic downturns without much of a, if any, decline in inflation. I discuss how a broad range of recent papers and explanations

Paper prepared for the XXII Annual Conference of the Central Bank of Chile, "Changing Inflation Dynamics, Evolving Monetary Policy." I would like to thank Dennis Bonam for discussions, suggestions, and help with the simulations of the NK model. Ricardo Ruiz has provided excellent research assistance.

1. See Blinder (2018).

2. The potential for such an output-inflation tradeoff was first emphasized in Samuelson and Solow (1960)'s reinterpretation of Phillips (1958).

Changing Inflation Dynamics, Evolving Monetary Policy edited by Gonzalo Castex, Jordi Galí, and Diego Saravia, Santiago, Chile.

can be interpreted as shifts in the short-run aggregate supply (SRAS) curve that is the backbone of the upward-sloping Phillips curve.[3]

Figure 1. Inflation Rates in U.S. and Chile: 1996-2018

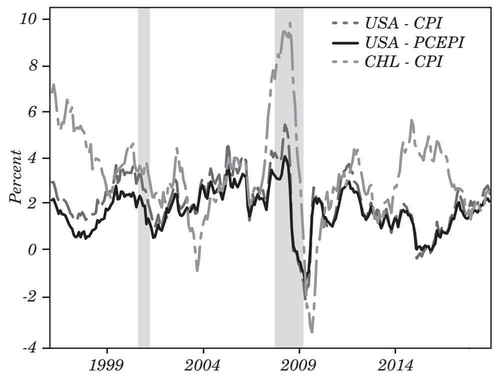

Sources: Bureau of Labor Statistics (BLS), Bureau of Economic Analysis (BEA), and Organization of Economic Cooperation and Development(OECD).
Note: 12-month inflation rates. Shading shows U.S. recessions.

Figure 2. Log Real GDP in U.S. and Chile: 1996-2018

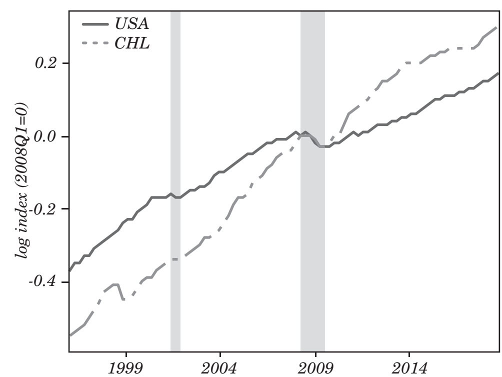

Sources: BEA and OECD.
Note: Log index 2008Q1=0. Shading shows U.S. recessions.

3. This includes Ravenna and Walsh (2006), Gilchrist and others (2017), Daly and Hobijn (2014), and Carlstrom and others (2017) among many.

Looking at the flattening of the Phillips curve through this joint AD-AS shift lens reveals some important insights. First of all, it implies that the flattening of the Phillips curve is not indicative of the absence of a transmission of monetary policy to the real economy. Instead, it suggests this transmission works through both the AD and the SRAS curves. Secondly, as a consequence of this first insight, this means that monetary policymakers have to think beyond the common focus on keeping "the growth of aggregate demand stable in order to prevent fluctuations in real output and inflation."[4] Finally, thinking beyond this common focus involves identifying and quantifying the supply-side effects of monetary policy and their impact on output and, most importantly for the second part of this paper, inflation.

In order to study the supply-side effects of monetary policy and their impact on inflation, we need to be able to measure how important supply-side factors, like factor costs, technology, and markups, are for inflation. One way would be to use a New Keynesian (NK) dynamic stochastic general equilibrium (DSGE) model.

But it is exactly that type of model that has not been particularly satisfactory in furthering our understanding of recent inflation dynamics. This is the reason I explore a different approach in this paper. Namely, to apply growth-accounting techniques that are generally used for the medium- to long-run analysis of the supply side of the economy for decomposing the sources of inflation.

In the second part of the paper I present the results obtained with this approach. I use dual growth-accounting methods to quantify the supply-side factors that underlie inflation in the headline personal consumption expenditures (PCE) price index[5] in the U.S. from 1999 to 2015.

The value chain of the PCE goods and services, whose price changes are captured in personal consumption expenditures price index (PCEPI) inflation, has not changed a lot from 1999 to 2015. The relative contributions of domestic industries to the cost of these goods has remained approximately constant over time. What has changed is the importance of imports and where they flow into the supply chain.

4. See Taylor (1997).
5. This is the price index that the Federal Reserve explicitly targets.

Since 1998 the share of the cost of PCE traceable to imports has increased from 7.6 percent on the dollar to 10.6 percent. This share peaked in 2008. Imports increasingly flow into the U.S. supply chain at more advanced stages of production. In terms of the production factors that contribute to these costs, the share of labor has declined steadily. This largely reflects the decline in the factor requirement of unskilled labor over time.

Import-price fluctuations played an outsized role in the dynamics of PCEPI inflation in the U.S. Even though imports only account for a tenth of the cost of PCE spending, import-price movements account for 45 percent of the variance in inflation. The contributions of changes in the costs of capital and total factor productivity (TFP) growth to inflation largely offset each other. This is possibly due to movements in markups that the dual growth-accounting method I use does not explicitly take into account. Labor compensation, even though it makes up half of the cost of PCE spending, accounts for less than a fifth of inflation fluctuations.

The data requirements for the dual growth-accounting methods I use are steep and the relevant data is released with a substantial delay. However, the contributions of import-price inflation, measured TFP growth, and, to a lesser extent, labor, can be reasonably approximated by using simple rules of thumb that can be implemented almost in real time.

The results in this paper show how the application of growth-accounting methods, normally used to analyze long-run growth and productivity trends, to short-run movements in inflation uncovers useful facts about the supply-side origins of inflation. These growth-accounting methods are based on neoclassical assumptions and do not, yet, allow for disentangling markups. Neither are they applicable in many countries other than the U.S. due to a lack of data. These are two areas that central banks possibly can contribute to with their research and resources.

1. Beyond Demand-Driven Inflation Fluctuations

To understand what I mean by the "supply-side origins" of inflation, it is useful to start with the textbook explanation of the AD-AS model. Though such a textbook-type exposition definitely does not do justice to the numerous academic studies that employ the three-equation NK model and variations and extensions thereof, it does capture the main intuition of many of the core principles that leading macroeconomists agreed on in 1997.[6,7]

The textbook explanation is illustrated in panel (i) of figure 3. The diagram in this panel can be understood in terms of the core principles laid out in 1997. First, the short-run aggregate supply curve in the panel captures that "there is a short-run tradeoff between inflation and unemployment".[8] Second, the shifts in the aggregate demand curve reflect the commonly-held belief that most fluctuations of output around its long-run trend "...are predominantly driven by aggregate demand impulses".[9] The latter is the equivalent of an identifying assumption in an instrumental variables (IV) regression.

In its purest form, plotted here, this implies that business-cycle fluctuations only shift the AD curve and are orthogonal to shifts in the SRAS curve. As a result, business-cycle fluctuations (to the extent they are not dampened by stabilization policies) result in shifts of the AD curve along the (fixed) SRAS curve. Thus, under this identifying assumption, business-cycle fluctuations allow for the identification of the slope of the SRAS curve, i.e., the sacrifice ratio.

If AD fluctuations are the (main) driver of business cycles, then the focus of stabilization policies should be to "...keep the growth of aggregate demand stable in order to prevent fluctuations in real output and inflation."[10] Though not easy to implement in practice, this is a remarkably simple conceptual description of optimal stabilization policies, including monetary policy.

The problem is that, in recent years, the empirical Phillips curve that such AD fluctuations imply is not in, or hard to extract from, the data.[11] The reason I emphasized the IV interpretation of the identifying assumptions underlying the Phillips curve above is that it provides us with a way to think through why we are not retrieving a positive correlation between output and inflation from the data.

6. See Blanchard (1997), Blinder (1997), Eichenbaum (1997), Solow (1997), and Taylor (1997).

7. The version of the AD-AS model that I plot here has the inflation rate on the vertical axis, rather than the price level. This is to bring the exposition more in line in with NK models.

8. See Taylor (1997).

9. See Solow (1997).

10. Taylor (1997).

11. It is important to realize that the Phillips curve implied by panel (i) of figure 3 is a simplification. Most empirical Phillips curve relationships include long lags. Moreover, even historically, the empirical Phillips curve worked well and was relatively stable only in the United States. (Blinder, 1997).

Within this textbook framework, there are three reasons why we could observe a flat Phillips curve. The first two maintain that business-cycle fluctuations are mainly driven by demand shocks. In that case, the SRAS curve can have flattened. Thus, firms' price-setting decisions depend less on the current level of economic activity. In a conventional three-equation NK model this could, for example, happen if there is an increase in nominal rigidities (especially price rigidities). Empirical studies using micro-price data do not reveal such an increase.

Another possibility would be that the AD curve has flattened. For example, a country where the central bank is hawkish on inflation will have a flatter AD curve than a country with a more dovish central bank. Of course, in this very simple stylized framework, a flat AD curve means that demand shocks do not affect the level of real activity, i.e., output, in the economy. Thus, in this simple diagram output fluctuations cannot be demand driven when the AD curve is flat. Though this is an artifact of the simple framework I use here, it does bring me to the third possible reason that the empirical Phillips curve has not been stable in recent years.

This third reason is what is plotted in panel (ii) of figure 3. It is that economic fluctuations in recent years have been driven by positively correlated demand and supply shocks of similar magnitude. That is, the sources of recent economic fluctuations violate the IV identifying restriction that allows us to recover the sacrifice ratio. That is, declines in demand, like during the Great Recession and its aftermath, were accompanied by shifts in the SRAS curve. As a result, the downward pressures on inflation from the AD shifts are offset by the upward pressures on inflation resulting from the shift in the SRAS curve. Panel (ii) of figure 3 illustrates the case in which the correlated shocks fully offset each other in terms of inflation.

The textbook AD-AS framework that I use to illustrate my point in figure 3 might seem rather simplistic. However, the main insight translates directly to a standard three-equation NK model. In fact, figure 4 plots the NK Phillips curve, i.e., the relationship between the percent deviation of output and inflation from their steady-state values in two cases.

The case in the left panel is the one that satisfies the conventional assumption that short-run economic fluctuations are due to demand shocks. Demand shocks in the context of this model reflect fluctuations in the representative household's discount factor.[12] As you can see, the NK model in that case results in a conventional Phillips curve that reflects a positive short-run output-inflation tradeoff.

12. The log-linearized version of the model is described in appendix A.

Figure 3. Slope of Phillips Curve Depends on Relative Demand and Supply Shocks

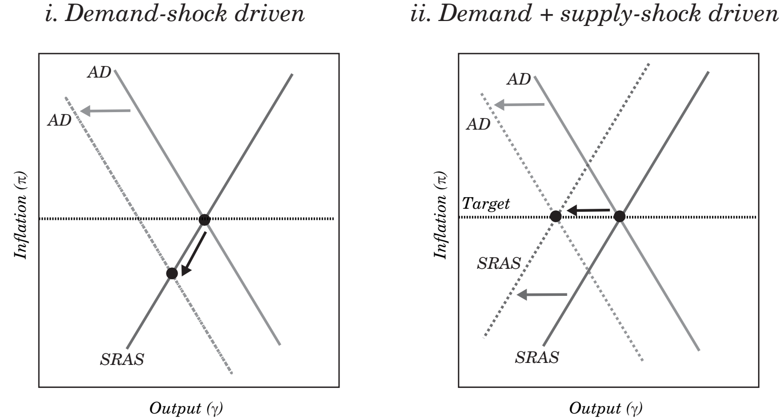

Source: Author's calculations.

Figure 4. Phillips Curve in NK-model with Uncorrelated and Correlated Shock

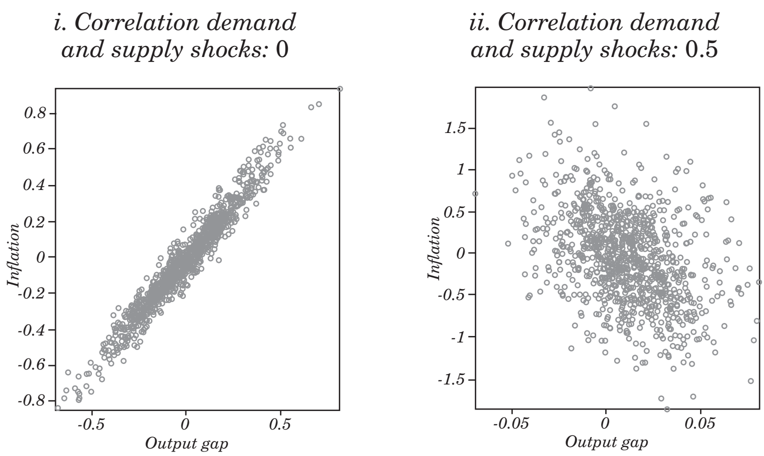

Source: Author's calculations.
Note: Plotted are comovement of inflation and output (percent deviations from steady state) in demand-shock-driven fluctuations under different assumptions about the correlation between demand and supply shocks.

The panel on the right in figure 3 shows the NK Phillips curve from the same model, but now for the case in which the demand (discount-factor) shocks are positively correlated with the supply shocks in the model. These supply shocks affect the marginal cost of production and shift firms' price-setting decisions. The panel plots the relationship between the percent deviation of output and inflation from their steady-state values when this correlation is 0.5. Even at this low correlation, the sign of the equilibrium reduced-form regression coefficient of inflation on output in the NK model changes from positive, i.e., the sacrifice ratio plotted in the left panel, to negative.

Thus, the importance of the correlation between demand and supply shocks for the empirical identification of the Phillips curve is not a moot point. It is relevant in the class of models most commonly used for monetary policy analysis by central banks.

Note that this observation that supply shocks might be important for shaping the recent relationship between output and inflation does not necessarily render monetary policy ineffective. Instead, it should make us think beyond (recent) monetary policy measures only affecting aggregate demand, as in the textbook AD-AS model as well as the conventional NK model.

In fact, there is a large number of research papers that, though not explicitly put in this context, already do so. For example, Ravenna and Walsh (2006) explicitly focus on the cost channel of monetary policy, where the interest rate that the central bank sets directly affects the marginal cost of production through the cost of financing working capital needed in production. Daly and Hobijn (2014) discuss how the equilibrium impact of downward nominal wage rigidities can be interpreted as a supply shock in that they affect the relationship between marginal cost and resource slack and thus firms' price-setting decisions and, in the simple AD-AS framework, the SRAS curve. The result is a flattening of the (wage) Phillips curve in their model. Gilchrist and others (2017) show how firms' liquidity levels affected their price-setting decisions, and thus the SRAS curve, during the financial crisis. Finally, Carlstrom and others (2017) show how quantitative easing also can have an effect on the supply side of the economy and potentially offset a negative supply shock.

The distinction of demand and supply shocks itself is largely a product of the AD-AS model being the workhorse model for the analysis of stabilization policies, where demand shocks affect preferences and supply shocks affect technology. This is in line with Ramey (2016), who defines "…shocks we seek to estimate as the empirical counterparts to the shocks we discuss in our theories, such as shocks to technology, monetary policy, and fiscal policy." However, Ramey (2016) also points out that shocks "…(1) should be exogenous with respect to the other current and lagged endogenous variables in the model; (2) they should be uncorrelated with other exogenous shocks; otherwise, we cannot identify the unique causal effects of one exogenous shock relative to another."

In this sense "correlated demand and supply shocks" is an oxymoron. The oxymoronic observation that we have "correlated demand and supply shocks" poses challenges at three different levels.

At a theoretical level, it means that the common source that drives both of these shocks needs to be modeled. Since this common source moves both the AD and SRAS curves, the specific distinction between these two curves in the AD-AS, as well as NK, framework might not necessarily be the most useful in this case. As I discussed above, however, there are already many papers that are up to this challenge and introduce mechanisms that result in joint shifts of the AD and SRAS curves.

At a policy level, it is important that we realize that such mechanisms might invalidate our narrative of monetary policy offsetting demand shocks and managing fluctuations in aggregate demand along a relatively fixed SRAS curve. This means that the Fed's dual mandate of "price stability and maximum employment" does not necessarily involve a positive output-inflation tradeoff inherent in the existence of a Phillips curve.

Moreover, it also means that it is important for policymakers to clearly communicate the mechanisms through which monetary policy measures are transmitted to the supply side of the economy. The reason I cited the four papers with such mechanisms above is that all four of them provide clear insights into how monetary policy decisions affect the supply side of the economy: through affecting the cost of working capital of firms, greasing the wheels of the labor market, alleviating financial constraints, and quantitative easing.

Finally, at a measurement level, it is important to improve our understanding of and to account for the supply-side factors that drive the inflation rate that the central bank targets, i.e. PCEPI inflation in the United States. In the rest of this paper, I address this third challenge.

2. Measuring the Supply-Side Origins of Inflation

One approach is to study these supply-side factors that drive inflation in the context of a model. A model is useful because it allows for counterfactual analyses and is very explicit about the general equilibrium effects at play. In the simple three-equation NK model that I used in the previous section, the supply-side factors that determine current inflation are: (i) Expected future inflation, (ii) the degree of nominal (price) rigidities, and (iii) all things that affect the marginal cost of production. Of course, most of these models imply paths of demand and supply shocks that are correlated and thus do not have a structural interpretation.

Another approach, which is the one I am taking here, is to use an accounting framework to measure these supply-side factors. The type of accounting exercise, using dual growth-accounting techniques, that I perform here explicitly takes the scope of the costs of PCE into account and traces these costs along the domestic value-added chain as well as the costs of imports to account for the production factors that contribute to the value added that makes up personal consumption expenditures.

For example, wages make up the bulk of the (marginal) cost of production in the economy. Thus, using the right measure of wages is important.[13] The problem is that the wage measures most often used by economists are not constructed to measure the cost of production of *consumption goods*, but instead to cover all value added in the economy. This is also true for other measures of factors that capture marginal costs. The growth-accounting exercise that I perform is meant to construct the factor costs relevant for the production of PCE.

Of course, I am not the first to use growth-accounting techniques to account for supply-side factors in the economy. Long-run trend forecasts, like that for potential output in table 1–2 in Congressional Budget Office (2018) and the table on page 24 in Federal Reserve Board of Governors (2012), are mostly derived by using growth-accounting methods.

What distinguishes my accounting exercise from those that focus on trend growth is the following. First, the scope of my analysis is

13. For example, to deal with this, Justiniano and others (2013) use measurement equations for compensation per hour and average hourly earnings in the empirical state-space model that they estimate based on their DSGE model.

different. Because the Federal Reserve, just like most other central banks, focuses on consumer price inflation, and in particular the PCEPI, I focus on personal consumption expenditures rather than GDP. Second, I perform a *dual* growth-accounting exercise. Using this dual approach allows me to focus on the *price* of consumption goods rather than on the quantity. Finally, I consider the short-run rather than the long-run in that I decompose the annual percent change in the PCEPI.

The data requirements for the accounting exercise I perform here are steep. However, for the U.S. the data needed are part of the integrated Bureau of Labor Statistics/Bureau of Economic Analysis (BLS/BEA) industry-level production account and the BEA's annual input-output accounts. The combined annual data that I use cover 1998–2015.

2.1 The PCE Value Chain Has Been Relatively Stable

The first step in disentangling the supply-side factors that drive PCE inflation is to identify the sectors in the U.S. economy as well as the types of imports that account for the value added embodied in the final goods and services that households (and non-profits) buy. The PCE value chain uncovered in this step has been relatively stable over the 18 years covered in the data. This result, and how it is derived, is best understood in the context of figure 5.

Figure 5. Tracing Sources of Costs of PCE

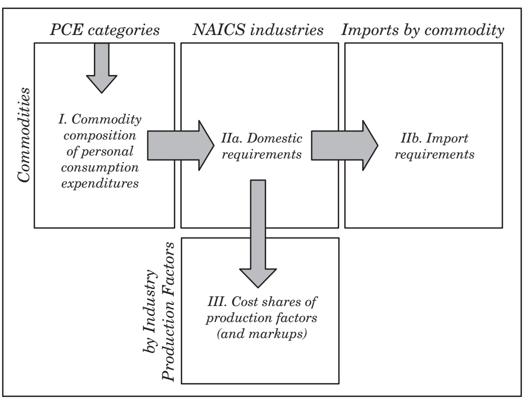

Sources: BEA and OECD.

Panel I of the figure shows how the cost of consumer spending on different categories of goods and services is tracked to the commodities that make up these goods and services. For example, when one buys a bottle of milk at the supermarket, then part of this spending is classified as a retail sales commodity, i.e., the markup the supermarket charges, and part of it as a food manufacturing commodity, i.e., the supermarket's cost of the bottle of milk.[14]

Panels IIa and IIb show how we can trace the cost of the retail sales and food manufacturing commodities of this bottle of milk up the domestic supply chain. For example, part of the retail sales cost of the bottle of milk reflects the intermediate goods and services the supermarket buys, like its electricity bill which, in turn, reflects the cost of utilities. Part of the cost of the bottle of milk reflects the cost of intermediate goods and services bought by the dairy producer. Some of these intermediate goods and services, like the glass bottle and the milk, are themselves commodities produced in the United States. These domestically produced intermediate inputs can be traced further up the domestic value chain in terms of panel IIa of the figure. Other intermediate inputs of the dairy producer, like the plastic cap that seals the bottle, are imported from abroad. These imported intermediates cannot be traced further along the domestic value chain and are accounted for as separate supply-side factors.[15]

The part of the cost of the supermarket that sells the bottle of milk that is not due to the cost of intermediate goods and services is the value added that the supermarket contributes to the cost of the bottle of milk sold to consumers. Similarly, the part of the producer price of the bottle of milk that is not due to the intermediate goods and services the dairy producer buys is the value that dairy producer adds. At the end, the cost of the bottle of milk for consumers reflects both value added by domestic industries at different stages along the value chain as well as the cost of imported intermediates at different stages along the value chain.

14. Because the bottle of milk is simply resold by the supermarket and not transformed in the process of production it is not counted as an intermediate input of the supermarket.

15. The imports that are counted in the value-added chain are imports that are directly sold to final demand, consumers in the case of the analysis in this paper, and imports used as intermediate inputs. Imported capital goods that are used in production are accounted for as part of the factor cost of capital.

Tracing the cost of PCE up the domestic value chain to figure out the value added required in each industry as well as the imports required to produce the goods and services bought by consumers, as illustrated in figure 5, can be done by using input-output analysis. This yields what is known as total requirements for the production of the final goods and services that make up PCE. The math involved in this calculation is explained in subsection A.2 of appendix A.[16]

The results of tracing these total domestic and foreign requirements per dollar of PCE by subperiod, as well as the average over the whole period, are reported in table 1. As an example, the 15.8 in the row "Trade and transportation" for 1999 means that 15.8 cents per dollar of PCE spending in 1999 was produced as value added in the retail and wholesale trade and transportation industries.

Two things stand out from this table. The composition of the domestic requirements in part (a) of the table does not vary much over the subperiods reported. This suggests that the domestic part of the PCE value chain is relatively stable over time.[17] Most notable are the declines in the importance of manufacturing and of trade and transportation during the sample period, and the rise of the importance of education and health. Also note the low total requirement for government production for PCE.

The biggest change is the increased importance of imports for PCE spending from 1998 to the Great Recession in 2008, reported in the "Total imports" row in part (b) in table 1. Over that period, the import requirements for PCE spending increase from 7.9 cents on the dollar to 11.8 cents. Since the Great Recession, this has declined to 9.8 cents on the dollar in 2015. A lot of this decline has to do with energy imports.

Overall, though, the composition of the industries and the imports that account for the production of the value added that makes up the cost of PCE spending has been relatively stable over the 18 years in the sample. The relative stability of this composition does not necessarily mean the value chain itself has been stable. For example, the length of the value chain might have changed because of vertical specialization, as in Yi (2003).

16. See also ten Raa (2006) for an exposition of input-output analysis. The calculation of the total requirements for PCE here generalizes those applied in Hobijn (2008), Hale and Hobijn (2011), and Hale and others (2012).

17. Part of this might reflect that input-output data are collected relatively infrequently. This might result in these data understating the actual higher frequency fluctuations in these shares.

Table 1. Domestic and foreign Requirements per Dollar of Personal Consumption Expenditures (PCE) by Year

A. 1998–2007

	1998	*1999*	*2000*	*2001*	*2002*	*2003*	*2004*	*2005*	*2006*	*2007*
					Domestic requirements					
Agriculture, forestry, fishing, and hunting	1.3	1.1	1.1	1.1	1.0	1.2	1.3	1.1	1.1	1.1
Mining and utilities	2.7	2.7	2.7	2.6	2.5	2.6	2.8	2.8	3.0	3.1
Construction	0.5	0.5	0.6	0.6	0.7	0.7	0.7	0.8	0.8	0.9
Manufacturing	9.9	9.9	9.5	9.4	9.2	9.1	8.8	8.5	8.4	8.1
Trade and transportation	16.3	15.8	15.6	15.3	15.1	15.1	14.9	14.9	14.8	14.5
Information	4.1	4.2	3.9	3.9	4.1	4.0	4.2	4.0	3.8	4.1
Finance, insurance, and non-housing real estate	11.8	11.8	11.9	12.4	12.2	12.0	11.5	11.9	11.7	11.6
Housing	13.7	13.6	13.5	13.8	14.0	13.7	13.7	13.7	13.8	13.7
Professional and business services	8.6	8.7	8.9	8.8	8.9	8.8	8.8	8.8	8.8	9.1
Education and health	9.8	9.7	9.5	9.8	10.2	10.4	10.4	10.3	10.4	10.4
Arts, entertainment, and food services	5.1	5.3	5.3	5.1	5.3	5.2	5.3	5.2	5.2	5.2
Other services	5.0	4.9	4.8	4.5	4.5	4.4	4.3	4.2	4.2	4.1
Government	3.5	3.4	3.3	3.3	3.3	3.3	3.2	3.1	3.1	3.0
Other	0.5	0.6	0.6	0.6	0.6	0.4	0.4	0.4	0.4	0.4
Total value added	92.1	91.7	90.7	90.9	90.9	90.6	89.9	89.4	89.0	88.9
					Import requirements					
Materials	6.3	6.4	6.8	6.6	6.7	6.8	7.1	7.2	7.2	7.2
Energy	0.7	0.9	1.5	1.3	1.1	1.4	1.7	2.3	2.4	2.5
Services	0.5	0.5	0.5	0.6	0.7	0.7	0.7	0.7	0.9	1.0
Other	0.5	0.6	0.6	0.6	0.6	0.4	0.4	0.4	0.4	0.4
Total imports	7.9	8.5	9.4	9.1	9.0	9.3	10.0	10.5	10.9	11.0

Table 1. (continued)

B. 2008–2015 and 1998-2015 Average

	2008	*2009*	*2010*	*2011*	*2012*	*2013*	*2014*	*2015*	*Average*
					Domestic requirements				
Agriculture, forestry, fishing, and hunting	1.1	1.1	1.2	1.3	1.3	1.4	1.3	1.1	1.2
Mining and utilities	3.4	3.2	3.3	3.4	3.2	3.4	3.3	2.8	3.0
Construction	0.8	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.8
Manufacturing	7.8	8.4	8.1	8.0	7.9	7.9	7.8	7.9	8.6
Trade and transportation	14.1	14.3	14.1	13.9	14.0	14.2	14.0	14.3	14.7
Information	4.1	4.0	4.1	3.9	3.8	4.0	3.8	4.0	4.0
Finance, insurance, and non-housing real estate	10.2	10.7	10.9	10.8	11.4	11.2	11.5	12.0	11.5
Housing	14.1	14.8	14.4	14.2	14.0	13.9	14.0	14.1	13.9
Professional and business services	9.5	9.0	9.0	9.1	9.2	9.2	9.3	9.5	9.0
Education and health	10.9	11.7	11.6	11.5	11.5	11.4	11.3	11.4	10.7
Arts, entertainment, and food services	5.1	5.1	5.1	5.0	5.2	5.3	5.4	5.5	5.2
Other services	4.0	4.0	3.9	3.8	3.8	3.8	3.9	3.9	4.2
Government	3.0	3.2	3.2	3.1	3.1	3.1	3.1	3.2	3.2
Other	0.4	0.4	0.3	0.3	0.2	0.3	0.3	0.4	0.4
Total value added	88.2	90.4	89.7	88.8	89.2	89.6	89.6	90.5	90.0
					Import requirements				
Materials	7.1	6.1	6.6	7.2	7.1	7.1	7.3	7.2	6.9
Energy	3.2	1.8	2.3	2.7	2.4	2.1	1.9	1.1	1.8
Services	1.1	1.2	1.1	1.1	1.1	1.1	1.1	1.0	0.9
Other	0.4	0.4	0.3	0.3	0.2	0.3	0.3	0.4	0.4
Total imports	11.8	9.5	10.4	11.3	10.9	10.6	10.6	9.8	10.0

Source: Author's calculations.

Note: Reported are cents of domestically produced and imported value added required for production of a dollar of PCE. Each column contains the average over period covered in column. Totals do not add up to 100 due to rounding.

There is little evidence for that in the data, though. The length of the domestic value chain has not changed much between 1998 and 2015. The main change has been where imports flow into the value chain. This can be seen from figure 6.

Panel (a) shows the cents of domestic requirements in a dollar of PCE by how many stages of transformation they go through before they are sold to final demand for both 1998 and 2015. This distribution can be used to gauge the length of the domestic supply chain. As can be seen from the figure, little has changed over the 18 years in the sample. What has changed is displayed in the right panel, i.e. panel (b). It shows how the import requirements, in cents on the dollar of PCE, are distributed along the number of transformation steps they take before they reach consumers. As can be seen from the figure, imports in 2015 flowed into the U.S. closer to final demand than in 1998. That is, imports in the U.S. take fewer steps along the supply chain now than 20 years ago.

The reason that it is important to look at the length of the supply chain is that several studies emphasize how the distortions due to nominal rigidities can be amplified along the supply chain in the economy.[18] The evidence here suggests that such amplification has not increased over the past two decades due to a lengthening of the value chain. This is because, just like the composition of total requirements in PCE, the length of the PCE value chain has been relatively constant over time.

Figure 6. Length of Value Chain for Requirements of a Dollar of PCE (1998 and 2015)

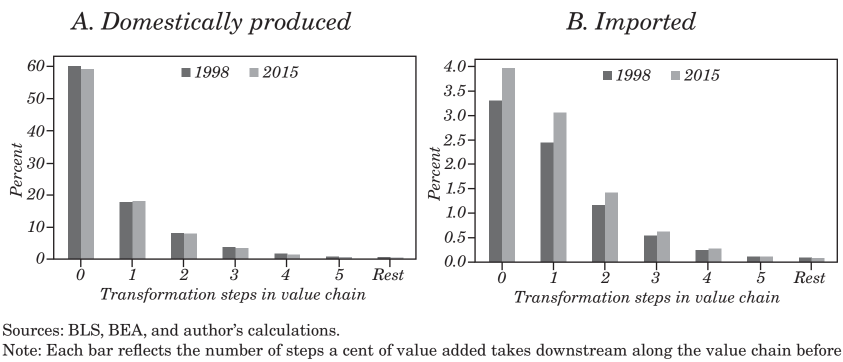

Sources: BLS, BEA, and author's calculations.
Note: Each bar reflects the number of steps a cent of value added takes downstream along the value chain before it is sold to final demand in terms of PCE.

18. See Huang and Liu (2001), Nakamura and Steinsson (2010), Pasten and others (2017).

2.2 Factor Requirements Reflect Decline of Labor Share

The next step in disentangling the supply-side factors that drive PCEPI inflation is to split up the industry value-added requirements into parts, due to different types of labor and capital used as factors of production. In terms of figure 5, this is reflected by the arrow from panel IIa to panel III. The results of this calculation are the total factor requirements that measure the cents on the dollar of PCE spending that can be traced to payments to different types of labor and capital.

These factor requirements are reported in table 2. Labor is split up into workers with and without a college education. The types of capital that are distinguished in the data are three: R&D, intangibles, information and communication technology (ICT) and a residual category. This would probably not be the classification of capital goods that a macroeconomist interested in inflation would choose, but it is the result of these data having been constructed for the analysis of long-run productivity trends.

On the labor side, the factor requirement of college-educated labor has steadily increased over the 18 years in the sample, from 22.8 cents on the dollar in 1998 to 25.9 in 2015. This increase is more than offset, however, by the decline in the factor requirement of non-college-educated labor that fell from 29.3 in 1998 to 22 in 2015. The net result is decline in the factor requirement of labor in the production of PCE goods and services, i.e., the *Labor-Total* row in the tables, from 52.1 in 1998 to 47.9 in 2015.

To compare this with more oft-cited measures of the labor share, one needs to consider this as a fraction of the domestic value-added requirement reported in the bottom row of the tables. This implies that the labor share of the domestically produced value added sold to consumers has declined from 56.5 to 52.9 percent. This means two things: First of all, the labor share in the domestic production of PCE goods and services is lower than in the nonfarm business sector. Second, the labor share in the domestic production of PCE goods and services has declined less than that in the nonfarm business sector.[19]

19. See Elsby and others (2013) for discussion of the time path of the latter labor share.

Table 2. Domestic Requirements per Dollar of Personal Consumption Expenditures (PCE) by Production Factor and Year

A. 1998–2007

	1998	*1999*	*2000*	*2001*	*2002*	*2003*	*2004*	*2005*	*2006*	*2007*
					Labor					
College	22.8	22.4	23.1	23.7	23.9	23.6	23.8	23.6	23.9	23.7
No college	29.3	28.9	27.9	27.8	26.7	26.5	25.6	25.0	24.5	24.4
Total labor	52.1	51.3	51.0	51.5	50.7	50.1	49.4	48.6	48.4	48.1
					Capital					
R&D	1.6	1.6	1.5	1.5	1.7	1.7	1.7	1.5	1.6	1.5
IT	2.7	2.8	2.8	2.8	2.7	2.6	2.4	2.3	2.2	2.1
Software	1.6	1.7	2.0	2.0	2.3	2.5	2.5	2.5	2.5	2.6
Other	33.6	33.7	33.0	32.5	32.9	33.1	33.2	33.8	33.8	33.8
Total capital	39.5	39.8	39.3	38.9	39.6	39.8	39.9	40.2	40.0	40.1
Total value added	92.1	91.7	90.7	90.9	90.9	90.6	89.9	89.4	89.0	88.9

Table 2. (continued)

B. 2008–2015 and 1998-2015 Average

	2008	*2009*	*2010*	*2011*	*2012*	*2013*	*2014*	*2015*	*Average*
	Labor								
College	23.6	24.3	23.9	24.2	24.9	24.9	24.8	25.9	24.0
No college	23.6	23.7	22.9	22.3	22.3	22.2	22.1	22.0	24.9
Total labor	47.2	48.0	46.8	46.5	47.2	47.0	46.9	47.9	48.8
					Capital				
R&D	1.5	1.7	1.7	1.7	1.6	1.7	1.8	1.9	1.6
IT	2.2	2.0	1.9	1.7	1.7	1.7	1.6	1.6	2.2
Software	2.5	2.6	2.6	2.5	2.5	2.5	2.6	2.8	2.4
Other	34.2	35.5	36.0	35.7	35.5	36.0	36.0	35.7	34.3
Total capital	40.3	41.8	42.1	41.6	41.3	41.9	42.0	41.9	40.6
Total value added	88.2	90.4	89.7	88.8	89.2	89.6	89.6	90.5	90.0

Source: Author’s calculations.

Note: Reported are cents of domestically produced and imported value added required for production of a dollar of PCE. Each column contains the average over period covered in column. Totals do not add up to 100 due to rounding.

There is little evidence that the labor share of low-skilled workers in the domestic production of PCE goods and services has declined because of capital-labor substitution between low-skilled workers and ICT capital. If this was the case, then the decline in the factor requirement of non-college-educated workers should be mostly offset by an increase in the factor requirements of ICT capital and software. However, we have only seen a small increase in the factor requirements of these two types of capital.

Instead, two other mechanisms seem to be putting downward pressure on the PCE factor requirement of labor. They can be best seen when looking at figure 7. The figure plots the time series of factor requirements per dollar of PCE for labor and capital as well as the import requirements. As can be seen in the figure by comparing the line for "Labor" with the other two, the decline in the factor requirement of labor over the 18 years of the sample can be split into two episodes. In the first, from 1998 to 2008, when the labor requirement declined by 4 percentage points, it was offset by an increase in the import requirement. This is consistent with the cross-industry evidence from Elsby and others (2013) that declines in labor shares occurred in industries with more import competition, i.e., that there was import substitution of unskilled labor. During the second episode, from 2008 to 2015, the factor requirement of labor did not decline much, but that of unskilled labor did, and it was offset by an increase in that of skilled labor. The decline in the factor requirement of unskilled labor coincided with an increase in the factor requirement of other non-ICT and non-R&D capital. There are several potential explanations that are consistent with such a shift in factor requirements. Capital/non-skilled-labor substitution in response to low interest rates would be one of them.[20]

What is most striking from table 2 as well as figure 7 is that there are no obvious cyclical fluctuations in the factor requirements and that what is most important is the longer-run trends. An important caveat is the question whether the pattern in the eight years post-2008 is partly reflective of the prolonged low-interest rate regime the economy was in or a continuation of longer-run shifts in factor usage in the production of consumer goods and services.

20. Rognlie (2015) points out the importance of the increase in the factor share of housing and structures for the trend in the U.S. labor share in longer-run data.

Figure 7. Factor Requirements of a Dollar of PCE

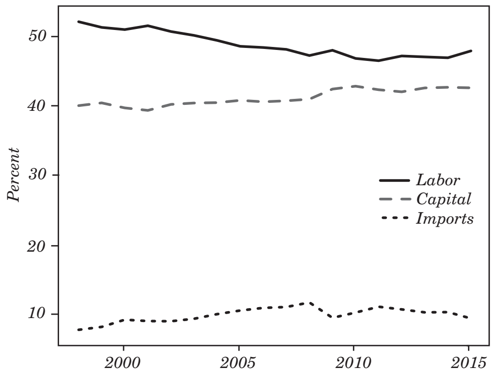

Sources: BLS, BEA, and author's calculations.
Note: Shares of value added embodied in PCE traced to capital, labor, and imports.

2.3 Bulk of Inflation Fluctuations Related to Import Prices

The final step in disentangling the supply-side factors that drive PCEPI inflation is to calculate the importance of changes in the costs of production factors and import prices for PCEPI *inflation*. This translation can be done by using the realization that PCEPI inflation is approximately a weighted average of the percent changes in factor costs and import prices. The weights in this average correspond to the requirements reported in the previous two subsections. The formal mathematical derivation of this result is in subsection A.3 of appendix A.

I present the results obtained in this final step in three parts. First, I look at how much industries and imports contribute to PCE inflation. That is, I calculate the PCEPI inflation contributions based on the domestic and import requirements from panels IIa and IIb from figure 5. I then split up the contributions of domestically produced value added into those of different types of labor and capital. That is, I calculate the inflation contributions based on the factor requirements from panel III of figure 5. Finally, I take a more aggregate perspective and look at the PCEPI contributions of labor, capital, and imports over time.

How much each industry contributes to PCEPI inflation, as well as the inferred residual contribution of imports, is reported in table 3.[21] The top row of the table is the time series of annual PCEPI inflation that is decomposed.

21. The contribution of imports cannot be split up by type of imports because there are no import-price data by NAICS category before 2005.

Table 3. Contributions to Annual Personal Consumption Expenditures Price Index (PCEPI) Inflation from Value-Added Deflators by Industry and Imports

A. 1999–2007

	1999	*2000*	*2001*	*2002*	*2003*	*2004*	*2005*	*2006*	*2007*
PCEPI	1.48	2.45	1.90	1.31	1.92	2.46	2.80	2.68	2.51
	Domestically produced value added								
Agriculture, forestry, fishing, and hunting	–0.16	–0.09	0.08	–0.09	0.12	0.19	–0.19	–0.05	0.28
Mining and utilities	–0.05	0.29	0.32	–0.17	0.42	0.23	0.43	0.22	0.11
Construction	0.03	0.04	0.04	0.03	0.03	0.05	0.08	0.07	0.05
Manufacturing	0.17	0.04	0.14	–0.04	–0.03	–0.04	0.35	0.13	0.02
Trade and transportation	0.12	0.31	0.05	0.02	–0.07	0.28	0.32	0.44	0.35
Information	–0.01	–0.04	–0.00	–0.02	–0.00	–0.04	–0.11	–0.07	–0.06
Finance, insurance, and non-housing real estate	–0.20	0.20	–0.01	0.24	0.35	0.32	0.26	0.19	0.20
Housing	0.42	0.40	0.51	0.54	0.30	0.28	0.29	0.37	0.28
Professional and business services	0.36	0.43	0.18	0.15	0.16	0.45	0.30	0.35	0.45
Education and health	0.31	0.33	0.47	0.35	0.33	0.32	0.32	0.27	0.42
Arts, entertainment, and food services	0.21	0.17	0.21	0.19	0.08	0.11	0.18	0.18	0.21
Other services	0.20	0.23	0.26	0.20	0.13	0.14	0.17	0.17	0.16
Government	0.09	0.11	0.18	0.11	0.23	0.30	0.31	0.21	0.01
Total domestically produced	1.50	2.40	2.43	1.51	2.04	2.59	2.71	2.48	2.48
	Imported value added								
Imports	–0.02	0.05	–0.53	–0.20	–0.12	–0.13	0.09	0.20	0.03

Table 3. (continued)

B. 2008–2015 and 1999-2015 Average

	2008	*2009*	*2010*	*2011*	*2012*	*2013*	*2014*	*2015*	*Average*
PCEPI	2.95	–0.09	1.70	2.50	1.89	1.34	1.49	0.25	1.86
					Domestic requirements				
Agriculture, forestry, fishing, and hunting	0.02	–0.27	0.15	0.33	0.03	0.03	–0.11	–0.22	0.00
Mining and utilities	0.41	–0.57	0.24	0.19	–0.19	0.11	0.09	–0.60	0.09
Construction	0.00	0.01	–0.02	0.01	0.03	0.03	0.05	0.03	0.03
Manufacturing	0.19	0.35	0.10	0.39	0.36	0.03	0.18	0.13	0.14
Trade and transportation	0.35	0.49	0.19	0.28	0.41	0.18	0.17	0.34	0.25
Information	–0.05	–0.00	–0.03	–0.02	0.00	0.03	–0.00	–0.07	–0.03
Finance, insurance, and non-housing real estate	0.07	–0.52	0.28	0.19	0.42	0.36	0.41	0.34	0.18
Housing	0.19	0.22	–0.09	0.12	0.29	0.28	0.30	0.44	0.30
Professional and business services	0.12	0.14	0.08	0.10	0.13	0.15	0.10	0.26	0.23
Education and health	0.23	0.41	0.26	0.17	0.21	0.15	0.15	0.21	0.29
Arts, entertainment, and food services	0.16	0.26	–0.00	–0.01	0.20	0.14	0.13	0.25	0.16
Other services	0.16	0.19	0.10	0.09	0.10	0.12	0.10	0.14	0.16
Government	–0.05	–0.03	0.10	0.06	0.10	0.12	0.13	–0.00	0.12
Total domestically produced	1.81	0.69	1.37	1.90	2.09	1.72	1.71	1.26	1.92
					Imported value added				
Imports	1.15	–0.78	0.33	0.60	–0.21	–0.38	–0.22	–1.01	–0.07

Sources: BLS, BEA, and author's calculations.
Note: The "Average" column contains the average over the 17 years in the sample.

Housing, education and health, and trade and transportation are, on average, the biggest contributors to headline inflation. This can be seen from the final column of table 3, labeled "Average". It lists the average percentage point contribution of each of the industries as well as imports to the 1.86 percent average annual rate of PCEPI inflation from 1999 to 2015. Together these three top contributing sectors account for 0.84 percentage points of the 1.86 percent average inflation.

However, these averages do not reflect the importance of these industries for inflation fluctuations. Three quarters of inflation *fluctuations* can be traced back to imports and to mining and utilities. This can be seen from figure 8, which decomposes the variance of annual PCEPI inflation over the 17 years in the sample into fluctuations in the contributions by industries and by imports. This result emphasizes the importance of commodity, especially oil, price fluctuations for headline PCEPI inflation.

The contributions of domestically produced value added to PCEPI inflation are divided into the parts due to different types of labor and to different types of capital in table 4. The "Average" column of the table shows that, in terms of levels, labor inputs account for two thirds of the average 1.86 percent of inflation over the 17 years for which I have data, IT capital costs reduce inflation by 0.2 percentage points, while measured TFP growth lowers inflation by 0.25 percentage points.

Figure 8. Variance Decomposition of Annual PCEPI Inflation by Industry and Imports

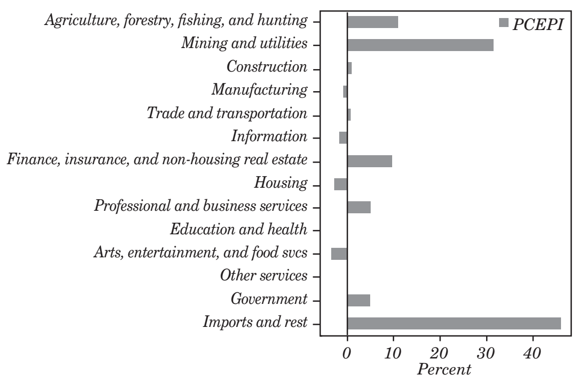

Sources: BLS, BEA, and author's calculations.
Note: Percent of variance of PCEPI due to industry and imports.
Reported is covariance between PCEPI inflation and industry and import contribution to PCEPI inflation as share of variance of PCEPI inflation.

Table 4. Contributions from Production Factors to Annual PECPI Inflation, 1999–2015

A. 1999–2007

	1999	*2000*	*2001*	*2002*	*2003*	*2004*	*2005*	*2006*	*2007*
					Labor				
College	0.35	1.81	1.54	0.26	0.41	1.37	1.15	0.88	0.08
No college	1.22	0.95	1.22	0.60	1.07	0.70	0.80	0.72	0.96
Total labor	1.57	2.76	2.76	0.85	1.48	2.07	1.96	1.60	1.04
					Capital				
Art	0.11	–0.16	0.04	0.12	–0.03	0.06	–0.01	0.01	0.05
R&D	0.01	–0.04	0.01	0.16	0.04	0.07	0.00	0.10	–0.00
IT	–0.40	–0.43	–0.43	–0.36	–0.18	–0.22	–0.11	–0.20	–0.22
Software	–0.09	0.04	–0.12	0.18	0.03	0.04	–0.10	–0.02	0.10
Other	1.34	0.60	–0.06	1.01	1.22	1.67	1.98	0.78	0.92
Total capital	0.98	0.01	–0.56	1.12	1.07	1.62	1.76	0.68	0.84
					Measured productivity growth				
TFP	–1.05	–0.37	0.23	–0.46	–0.51	–1.10	–1.00	0.20	0.59
Total value added	1.50	2.40	2.43	1.51	2.04	2.59	2.71	2.48	2.48

Table 4. (continued)

B. 2008–2015 and 1999-2015 Average

	2008	2009	2010	2011	2012	2013	2014	2015	Average
					Labor				
College	0.35	1.81	1.54	0.26	0.41	1.37	1.15	0.88	0.08
No college	1.22	0.95	1.22	0.60	1.07	0.70	0.80	0.72	0.96
Total labor	1.57	2.76	2.76	0.85	1.48	2.07	1.96	1.60	1.04
					Capital				
Art	0.01	−0.03	0.07	−0.02	−0.00	−0.00	−0.06	0.02	0.01
R&D	−0.03	−0.01	0.05	0.02	0.00	0.08	0.05	0.05	0.03
IT	−0.13	−0.34	−0.14	−0.13	−0.03	−0.04	−0.06	−0.00	−0.20
Software	−0.13	0.00	0.02	−0.01	−0.02	−0.04	0.13	0.04	0.00
Other	0.45	0.11	2.27	1.28	1.02	1.25	1.17	0.26	1.02
Total capital	0.16	−0.27	2.28	1.14	0.98	1.25	1.22	0.36	0.86
					Measured productivity growth				
TFP	1.46	0.00	−1.74	0.17	−0.15	0.13	−0.68	−0.29	−0.27
Total value added	1.81	0.69	1.37	1.90	2.09	1.72	1.71	1.26	1.92

Sources: BLS, BEA, and author's calculations.
Note: The “Average” column contains the average over the 17 years in the sample.

Just like for the industry-level analysis in table 3, the factor-level analysis in table 4 is misleading about the relative importance in terms of inflation *fluctuations*. The relative importance of changes in the cost of domestic production factors for inflation is shown in figure 9. Three things stand out from this figure. The first is the relative importance of the fluctuations in factor costs of other types of capital for inflation.

The second is the importance of fluctuations in TFP growth. In its purest form, these are the supply shocks I discussed above. In practice, of course, the *measured* contributions of capital and TFP to PCE inflation are both potentially affected by the cyclicality of markups, which the type of growth-accounting method I use here does not take into account.

Finally, most surprisingly, fluctuations in the compensation of college-educated labor are four times more important for inflation fluctuations than those of non-college-educated labor. This possibly reflects two things. First of all, that wages of non-college-educated workers are stickier, partly due to minimum wage restrictions and to them being disproportionately determined by union bargaining. Secondly, as Elsby and others (2013) show, a large part of aggregate fluctuations in compensation per hour is accounted for by sectors that pay bonuses. Thus, to some extent, the relative importance of fluctuations in the compensation of college-educated workers for inflation might be due to non-wage and salary aspects of compensation.

Figure 9. Variance Decomposition of Annual PCEPI Inflation for Production Factors

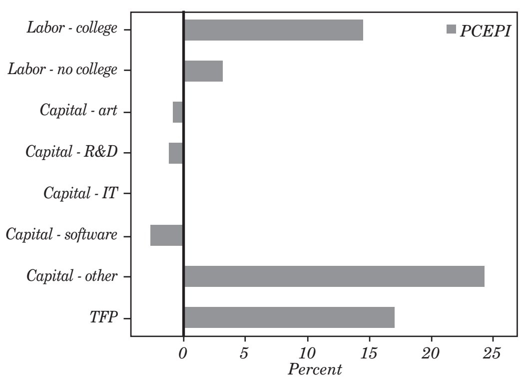

Source: BLS, BEA, and author's calculations.
Note: Percent of variance of PCEPI due to production factors.
Reported is covariance between PCEPI inflation and factor contribution to PCEPI inflation as share of variance of PCEPI inflation. Total does not add up to 100 because figure excludes contribution of imports.

Figure 10. Factor Contributions to Annual PCEPI Inflation

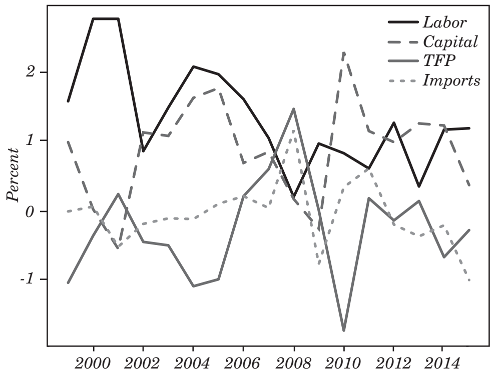

Sources: BLS, BEA, and author's calculations.
Note: Percentage point contribution of production factors, measured productivity growth (TFP), and import-price inflation to annual (yr/yr) PCEPI inflation.

Of course, most macroeconomists neither distinguish between college- and non-college-educated labor nor between different types of capital. For that reason, figure 10 plots the time series for the contributions of labor, capital, measured TFP, and imports to PCEPI inflation. The shares of inflation fluctuations that they account for are: 17.6, 19.7, 17.0, and 45.7 percent, respectively. That is, even though labor compensation accounts for the bulk of the cost if PCE spending, it only accounts for less than a fifth of inflation fluctuations. Fluctuations in the measured cost of capital and measured TFP growth tend to largely offset each other, possibly because of unaccounted movements in markups. This results in the contributions of these factors not comoving that much with headline inflation. Finally, though imports only make up a tenth of the cost of PCE spending, they play an outsized role in fluctuations in PCEPI inflation.

3. Real-Time Rule-of-Thumb Approximation

The measurement of the supply-side origins of PCEPI inflation that I presented in the previous section relies on data on U.S. input-output relationships and productivity accounts by industry, which are released with a substantial delay. In fact, the data that I use was released in November 2017 and only covers years through 2015.

Thus, in practice, the type of supply-side accounting for inflation that I do here might not be practical for the real-time analysis of inflation. It turns out, however, that several of the main results of subsection 2.3 can be approximated by using simple rules of thumb that are implementable in real time. These real-time rule-of-thumb approximations are shown in figure 11.

Figure11. Real-Time Rule-of-Thumb Approximation of Supply-Side Origins of Inflation

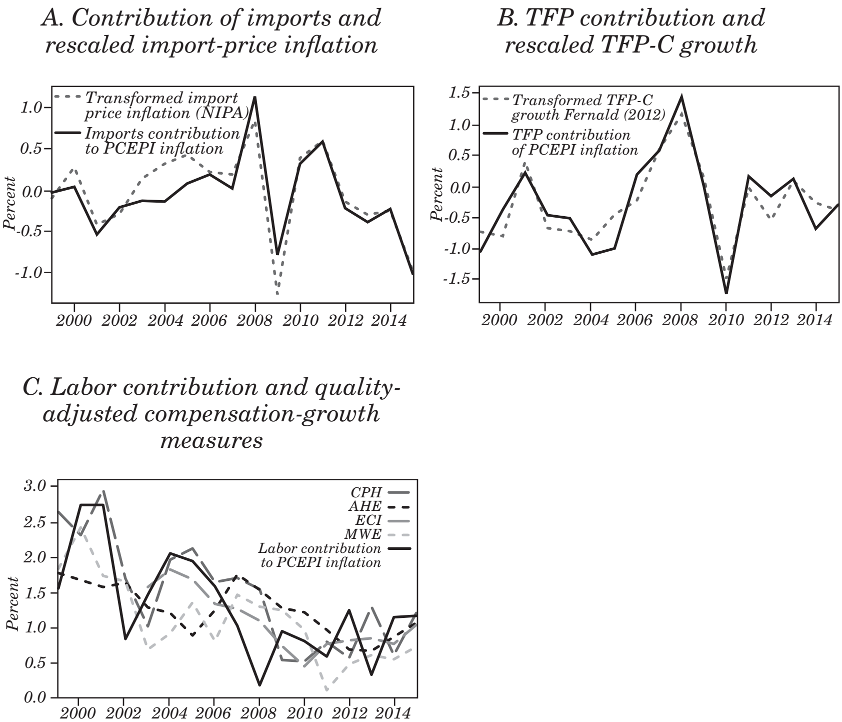

Sources: BLS, BEA, Fernald (2012), and author's calculations.
Note: Rescaled import-price inflation is $0.1\pi_t^M - 0.15$, where π_t^M is annual inflation in the implicit price deflator of imports of goods and services (NIPA, table 4.2.4, line 26). Rescaled TFP –C growth is $-0.5\Delta tfp_{c,t} - 0.25$, where $\Delta tfp_{c,t}$ is TFP-C from Fernald (2012). Quality-adjusted compensation-growth is $0.5(\Delta w_t - \Delta LQ_t)$, where Δw_t is annual growth rate of the respective compensation measure and ΔLQ_t is the growth rate of labor quality, based on Aaronson and Sullivan (2003), from Fernald (2012).

The top panel of the figure, i.e., panel (a), shows how the contribution of imports to annual PCE inflation can be closely approximated by $0.1\pi_t^M - 0.15$, where π_t^M is annual inflation in the implicit price deflator of imports of goods and services (NIPA, table 4.2.4, line 26).[22] The coefficient of 0.1 is in line with total import requirements reported in table 1. The deduction of 0.15 is a mean correction due to the rescaling of the import-price inflation rate.

As can be seen from the figure, this rule-of-thumb approximation does a very good job tracking the contribution of import-price inflation to PCE inflation. It is simple to calculate when one wants to gauge the importance of import-price inflation for PCE inflation when one does not have the input-output and productivity data that I relied on here.

The middle panel, i.e., panel (b), shows that the TFP contribution to PCE inflation lines up closely with total factor productivity growth of consumption goods from Fernald (2012)'s quarterly TFP growth data, published by the Federal Reserve Bank of San Francisco. In particular the TFP contribution to PCEPI inflation is approximately equal to $-0.5\Delta tfp_{c,t} - 0.25$, where $\Delta tfp_{c,t}$ is annual TFP-C growth from Fernald (2012). Thus, the effect of measured productivity growth on the inflation rate that the Fed targets can, in principle, be gleaned from data published with less of a delay than the data I use and at a quarterly basis. This is with the caveat that the quarterly TFP data, based on Fernald (2012), are subject to revisions. But so is PCEPI inflation, of course.

The bottom panel, i.e., panel (c), compares the labor contribution to PCEPI inflation to four measures of quality-adjusted compensation growth for the U.S. In the BLS/BEA data that I use, labor costs are calculated based on industry compensation per quality-adjusted hour measures. The quality adjustment is done by using the method explained in Jorgenson and others (2017) and is based on CPS-ASEC (Current Population Survey – Annual Social and Economic Supplement) data on self-reported sector of employment and earnings of individuals.

I compare the labor-cost contribution to PCEPI inflation with four commonly used aggregate compensation-growth measures, ΔLQ_t, for the U.S., namely: average hourly earnings (AHE), compensation per hour (CPH), employment cost index (ECI), and median usual weekly

22. Note, however, that figure 11 is not constructed with real-time data but instead with the data available in September 2018, when the results were calculated. So the rule-of-thumb approximation that is depicted is not real-time.

earnings (MWE). I adjust these compensation-growth measures for aggregate changes in labor quality by using the measure, ΔLQ_t, based on Aaronson and Sullivan (2003), from Fernald (2012).

Panel (c) shows the labor contribution to PCEPI inflation as well as 0.5 times the growth rate in quality-adjusted labor compensation based on each of these four measures. As can be seen from the figure, the labor contribution to PCEPI inflation is best approximated by rescaled quality-adjusted CPH and ECI growth. However, these two, as well as the other compensation-growth measures, overstate the contribution of labor costs to inflation in both the 2001 and 2008 recessions. That is, the contribution of labor-cost growth to headline PCE inflation is more procyclical than commonly used in macroeconomic time series of wage growth. This might partly reflect that the cost of PCE spending does not depend much on government production and thus on the wages of government workers, which tend to be less sensitive to market forces that drive business-cycle fluctuations.

A rule of thumb for the contribution of the cost of capital to PCEPI inflation is hard to find. This is because, being a user cost, this cost depends on a lot of factors: the composition of the capital stock used in producing PCE goods and services, depreciation rates, the internal rate of return of businesses, and the price of investment goods. In addition, due to the way the productivity statistics are calculated, capital is effectively the residual claimant in the factor attribution of revenue. As a consequence, changes in the measured cost of capital are also affected by movements in markups.

Still, relatively simple rule-of-thumb calculations can be used to approximate the factor contributions to PCEPI inflation for three out of the four supply-side factors I consider. These approximations can be useful when discussing the importance of these factors for inflation in real time.

4. Beyond Neoclassical Assumptions and Beyond the U.S.

I hope the dual growth-accounting exercise in the previous two sections has convinced you that it is worthwhile for central banks to explicitly account for the supply-side factors that are at the root of the inflation rates that they target. As I discussed above, the methodology that I used is not new, I just applied it with a different scope, focused on prices rather than quantities, and used it to analyze short-run fluctuations.

Because of this, my analysis in this paper is subject to the same limitations as other studies that use growth-accounting methods. Most notably, it is based on neoclassical assumptions that ignore the possible existence of markups. It is, of course, the variation in such markups due to nominal rigidities that gives rise to the monetary transmission mechanism in most theoretical NK models. Thus, to further the use of supply-side analyses of inflation, it is important to extend growth-accounting methods to also account for markups.[23] To give an example of why accounting explicitly for markups is important, in figure 8 I found that mining and utilities accounts for about a third of inflation fluctuations in the U.S. These contributions largely reflect changes in markups in the industry due to fluctuations in oil prices.

My analysis here focused solely on the U.S. I used the integrated industry-level production accounts (ILPA) for the U.S. This data has been published since 2014. Unfortunately, doing similar analyses for other countries is hard because of the lack of recent integrated growth-accounting and input-output data. The initial vintage (2014) of the World Input-Output Database (WIOD)[24], included socio-economic accounts that allowed for the type of dual growth accounting I did here. Unfortunately, the most recent vintage (2016) does not include the data on capital needed to do so. Similarly, the current version of the OECD statistical analysis (STAN)[25], that contains data on Chile, does not include the necessary input-output data to do the analysis I did here.

This lack of data, in large part, reflects a lack of funding for statistical agencies and cross-country data collection efforts. I hope the analysis in this paper shows that such funding is important in order to collect and construct the data necessary to assess how national and global value-added chains, factor costs, and, hopefully soon, markups drive the headline numbers that policymakers focus on.

It is imperative that central banks emphasize the importance of this type of data and, if necessary, contribute to the collection and construction of data that better help us understand the changing mix and dynamics of supply-side factors that contribute to fluctuations in output and inflation.

23. Hall (1988) is an older paper that addressed growth accounting with markups for aggregate data. A similar method to apply in the context of the input-output analysis used here has not yet been developed.

24. Stehrer and others, 2014.

25. OECD, 2017.

5. Conclusion

The disappearance of an empirical Phillips curve relationship in the data is indicative of recent economic fluctuations being affected by (positively) correlated demand and supply shocks. The correlation between these shocks poses a challenge on three different fronts.

Theoretically, we need models to better understand the source of these common fluctuations in demand and supply forces. There are several existing studies that provide such explanations but that do not explicitly place their results in this context. A reinterpretation of theories in this framework is useful.

In terms of policy, this disappearance of the Phillips curve does *not* mean that monetary policy has become ineffective. It is a reminder that it is important to understand and communicate the transmission of monetary policy measures to the production, rather than spending, side of the economy. It does indicate, though, that the Fed's dual mandate of "price stability and maximum employment" does not always involve a tradeoff.

The final challenge is to better measure the supply-side factors that drive inflation. In this paper, I use dual growth-accounting methods, normally applied for the analysis of long-run growth and productivity trends, to account for the supply-side factors that drive annual PCEPI inflation from 1999 to 2015.

I show that the value chain of PCE goods and services that determines the composition of the costs that drive PCEPI has been relatively constant over time. The two main trends are the increased importance of imports from 1998 to 2008 and the steady decline of the factor requirement of (unskilled) labor over time.

The relative shares of the supply-side factors in the cost of PCE goods and services, however, are not indicative of their relative importance for inflation *fluctuations*. In terms of changes in inflation over time, import-price inflation turns out to be the most important factor. Even though imports only account for a tenth of the cost of PCE, fluctuations in import prices drive 45 percent of fluctuations in inflation. The contributions of capital and measured TFP growth largely offset each other. Finally, even though labor accounts for about half of the cost of PCE goods and services, changes in compensation only drive a fifth of inflation fluctuations.

References

Aaronson, Daniel and D.G. Sullivan. 2003. "Regional Growth in Worker Quality." *Chicago Fed Letter*.

Blanchard, Olivier. 1997. "Is There a Core of Usable Macroeconomics?" *American Economic Review* 87(2): 244–46.

Blinder, Alan S. 1997. "Is There a Core of Practical Macroeconomics That We Should All Believe?" *American Economic Review* 87(2): 240 –43.

———. 2018. "Is the Phillips Curve Dead? And Other Questions for the Fed." May 3.

Carlstrom, C.T., T.S. Fuerst and M. Paustian. 2017. "Targeting Long Rates in a Model with Segmented Markets." *American Economic Journal: Macroeconomics* 9(1): 205–42.

Congressional Budget Office. 2018. *The Budget and Economic Outlook: 2018 to 2028*.

Daly, M.C. and B. Hobijn. 2014. "Downward Nominal Wage Rigidities Bend the Phillips Curve." *Journal of Money, Credit and Banking* 46(S2): 51–93.

Eichenbaum, M. 1997. "Some Thoughts on Practical Stabilization Policy." *American Economic Review* 87(2): 236 –39.

Elsby, M., B. Hobijn, and A. Sahin. 2013. "The Decline of the U.S. Labor Share." *Brookings Papers on Economic Activity* 44(2 (Fall)): 1–63.

Federal Reserve Board of Governors. 2012. *Report to the FOMC on Economic Conditions and Monetary Policy: Book A*.

Fernald, J.G. 2012. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Technical report.

Fleck, S., S. Rosenthal, M. Russell, E.H. Strassner, and L. Usher. 2014. *A Prototype BEA/BLS Industry-Level Production Account for the United States*. Chicago, IL: University of Chicago Press.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek. 2017. "Inflation Dynamics during the Financial Crisis." *American Economic Review* 107(3): 785–823.

Hale, G. and B. Hobijn. 2011. The U.S. Content of Made in China. FRBSF *Economic Letter*.

Hale, G., B. Hobijn, and R. Raina. 2012. "Commodity Prices and PCE Inflation." FRBSF *Economic Letter.*

Hall, R.E. 1988. "The Relation between Price and Marginal Cost in U.S. Industry." *Journal of Political Economy* 96(5): 921–47.

Hobijn, B. 2008. "Commodity Price Movements and PCE Inflation." *Current Issues in Economics and Finance* 14(Nov) 1–7.

Huang, K.X.D. and Z. Liu. 2001. "Production Chains and General Equilibrium Aggregate Dynamics." *Journal of Monetary Economics* 48(2): 437–62.

Jorgenson, D.W., M.S. Ho, and J.D. Samuels. 2017. *Educational Attainment and the Revival of U.S. Economic Growth*. Chicago, IL: University of Chicago Press.

Justiniano, A., G.E. Primiceri, and A. Tambalotti. 2013. "Is There a Tradeoff between Inflation and Output Stabilization?" *American Economic Journal: Macroeconomics* 5(2): 1–31.

Nakamura, E. and Steinsson, J. 2010. "Monetary Non-neutrality in a Multisector Menu Cost Model." *Quarterly Journal of Economics* 125(3): 961–1013.

OECD. 2017. STAN: OECD Structural Analysis Statistics 2016.

Pasten, E., R. Schoenle, and M. Weber. 2017. "Price Rigidities and the Granular Origins of Aggregate Fluctuations." NBER Working Papers 23750.

Phillips, A.W. 1958. "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957." *Economica* 25(100): 283–99.

Ramey, V.A. 2016. "Macroeconomic Shocks and Their Propagation." *Handbook of Macroeconomics*, vol. 2. Elsevier.

Ravenna, F. and C.E. Walsh. 2006. "Optimal Monetary Policy with the Cost Channel." *Journal of Monetary Economics* 53(2): 199–216.

Rognlie, M. 2015. "Deciphering the Fall and Rise in the Net Capital Share: Accumulation or Scarcity?" *Brookings Papers on Economic Activity* 46(1): 1–69.

Samuelson, P.A. and R.M. Solow. 1960. "Analytical Aspects of Anti-Inflation Policy." *American Economic Review* 50(2): 177–94.

Solow, R.M. 1997. "Is There a Core of Usable Macroeconomics We Should All Believe In?" *American Economic Review* 87(2): 230–32.

Stehrer, R, B. Los, H.W.A. Dietzenbacher, M. Timmer, and G.J. de Vries. 2014. "The World Input-Output Database: Content, Concepts and Applications." Techical report.

Taylor, J.B. 1997. "A Core of Practical Macroeconomics." *American Economic Review* 87(2): 233 –35.

Ten Raa, T. 2006. *The Economics of Input-Output Analysis*. Cambridge Books: Cambridge University Press.

Yi, K.-M. 2003. "Can Vertical Specialization Explain the Growth of World Trade?" *Journal of Political Economy* 111(1): 52–102.

Appendix A

Mathematical Details

A.1 Simple New Keynesian (NK) Model

The three-equation NK model that is simulated in section 1 boils down to the following log-linearized equations:

$$\hat{y}_t = \mathbb{E}_t \hat{y}_{t+1} - \frac{1}{\sigma}\left(\hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1}\right) + (1-\rho_D)\hat{z}_{D,t}, \quad (1)$$

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa\left(\varphi + \sigma\right)\hat{y}_t + \kappa(1+\varphi)\hat{z}_{S,t}, \quad (2)$$

$$\hat{R}_t = \left(\phi_\pi \hat{\pi}_t + \phi_y \hat{y}_t\right). \quad (3)$$

The table below lists these parameters and the definition of the equilibrium variables:

Variable	*Description*	*Value*
	Equilibrium variables	
$\hat{y}_t$	Output gap	-
$\hat{\pi}_t$	Inflation	-
$\hat{R}_t$	% deviation of gross from steady state	
	Shocks	
$a_{D,t}$	Demand shock	-
$z_{S,t}$	Supply shock	
	Parameters	
θ	Price stickiness	0.75
σ	Intertemporal elasticity of substitution	2
φ	Frisch elasticity of labor supply	3
β	Discount factor	0.99
φ_π	Inflation coefficient in Taylor rule	1.5
φ_y	Output-gap coefficient in Taylor rule	0.125
ρ_D	Persistence parameter of demand shocks	0.9
ρ_s	Persistence parameter of supply shocks	0

Source: Author's calculations.
Note: $.\hat{}$ denotes percentage deviation from steady state. Here $\kappa=(1-\theta)(1-\theta\beta)/\theta$.
The parameters of this model are calibrated for a quarterly frequency.

A.2 Derivations for Subsections 2.1 and 2.2

Though supply chains are often analyzed in terms of input-output analysis, I find it easier to think of them in terms of discrete-state Markov Chains. This is the interpretation that I use here. We will follow a dollar of final demand by consumers, i.e., a dollar of PCE, up the supply chain to where it either was imported or where it was created in terms of domestic value added. We denote the number of steps it has taken up the supply chain by s.

Throughout its journey up the supply chain this dollar can end up in three states. Either it can still be going up the supply chain in the form of gross output, or it has been traced to come from imports, or it has been traced to domestic value added in a particular industry. The latter two are absorbing states in that they are the origin of the value added (either foreign or domestic) that the dollar of PCE embodies.

In the following, the $(n_c \times 1)$-vector $\boldsymbol{c}_0$ represents the distribution of the dollar of PCE across the consumption categories. Because it reflects a distribution, $\iota'\boldsymbol{c}_0 = 1$, where ι is the summation operator, i.e., a vector of ones.

The $(n_j \times 1)$-vector $\boldsymbol{y}_s$ traces the fraction of the dollar of PCE that is still going up the supply chain after s steps. That is, the k^{th} element of $\boldsymbol{y}_s$ is the fraction of the dollar of PCE that was part of output of commodity k and then took s steps of transformation along the supply chain before it was sold to consumers.

The $(n_j \times 1)$-vector $\boldsymbol{m}_s$ is the fraction of the dollar of PCE, by commodity, that is imported into the U.S. and then takes s steps before it gets sold to consumers. The $(n_j \times 1)$-vector $\boldsymbol{v}_s$ is the fraction of the dollar of PCE that is produced, by industry, and goes through s transformation steps before ending up being sold to consumers. Each element in this vector corresponds to an industry.

We combine the last three vectors into a large $((2n_j + n_j \times 1)$-$)$-vector over which we define the Markov chain.

$$\boldsymbol{x}_s = \left[\boldsymbol{y}_s' \ \boldsymbol{m}_s' \ \boldsymbol{v}_s' \right]' . \tag{4}$$

The starting value $\boldsymbol{x}_0$ is determined by whether the consumption goods and services are made in the U.S.A. or imported from abroad. The $(n_j \times n_c)$-matrix $\boldsymbol{C}_y$ has the $(k,l)^{th}$ element that is the fraction of the consumption of category l that is supplied domestically. It is the part of the l^{th} element of $\boldsymbol{c}_0$ that is part of the k^{th} element of $\boldsymbol{y}_0$. Similarly, the $(n_j \times n_c)$-matrix $\boldsymbol{C}_m$ has the $(k,l)^{th}$ element that is the fraction of the consumption of category l that is imported and directly sold to final demand.

Given this definition, the starting value $\boldsymbol{x}_0$ can be written as

$$\boldsymbol{x}_0 = \begin{bmatrix} \mathbf{C}_y \\ \mathbf{C}_m \\ 0_{(n_i \times n_j)} \end{bmatrix} \boldsymbol{c}_0 . \tag{5}$$

Note that $\iota'\boldsymbol{C} = \iota'$, i.e., the column sums of $\boldsymbol{C}$ are one. The next step is to follow the dollar of PCE that is part of $\boldsymbol{y}_0$ up the U.S. domestic supply chain.

For this purpose, I define three matrices. The first, $\boldsymbol{A}_y$, is an $(n_j \times n_j)$-matrix for which the $(k,l)^{th}$ element is the domestically produced intermediate input revenue share of commodity k in gross output of commodity l. These shares are reported as part of the domestic direct requirements matrix in the BEA's annual input-output tables. The second, $\boldsymbol{A}_m$, is an $(n_j \times n_j)$-matrix for which the $(k,l)^{th}$ element is the imported intermediate input revenue share of commodity k in gross output of commodity l. These shares are derived by subtracting the domestic direct requirements matrix from the total direct requirements matrix. Finally, $\boldsymbol{A}_v$ is an $(n_j \times n_j)$-matrix for which the $(k,l)^{th}$ element is the value-added share of industry k in gross output of commodity l.[26] This matrix is derived by combining the direct requirements matrix with the make table.

$$\boldsymbol{x}_{s+1} = \begin{bmatrix} \mathbf{A}_y & 0 & 0 \\ \mathbf{A}_y & 0 & 0 \\ \mathbf{A}_v & 0 & 0 \end{bmatrix} \boldsymbol{x}_s = \mathbf{A}\boldsymbol{x}_s . \tag{6}$$

The matrix $\boldsymbol{A}$ is defined such that I drop the value of the dollar of PCE as soon as it ends in one of the absorbing states, i.e., when I have traced back the source of the value added. Moreover,

$$\boldsymbol{\iota}'\mathbf{A} = \left[\boldsymbol{\iota}'_{\boldsymbol{n}_j} \; 0'_{\boldsymbol{n}_j} \; 0'_{\boldsymbol{n}_i} \right].$$

26. More than one industry can have a non-zero share in each column of this matrix because some commodities are produced by more than one industry.

Defining the transition matrix this way means that $\boldsymbol{x}_s$ has the following two properties:

$$\lim_{s\to\infty} \boldsymbol{x}s = 0_{((2n_j+n_i)\times 1)} \text{ and } \textstyle\sum_{s=0}^{\infty} \left[0'_{\boldsymbol{n}_j} \;\; \boldsymbol{\iota}'_{\boldsymbol{n}_j+\boldsymbol{n}_i} \right] \boldsymbol{x}s = 1. \tag{7}$$

These two properties imply that the whole dollar of value added will be distributed into either imported value added or domestic value added along the supply chain that we decompose. The latter property in (7) is useful, because it means that our decomposition of a dollar of PCE can be written as

$$1 = \boldsymbol{\iota}'\boldsymbol{c}_0 = \textstyle\sum_{s=0}^{\infty} \iota'_{n_i} \boldsymbol{v}_s + \sum_{s=0}^{\infty} \iota'_{n_j} \boldsymbol{m}_s . \tag{8}$$

This allows us to trace where the value added that is sold to final demand in the form of nominal PCE originates, both domestically, by industry, and foreign, by imported commodity. For each industry, the value-added requirements in $\boldsymbol{v}_s$ can then be divided into the factor requirements of the different types of labor and capital based on data on factor shares by industry.

A.3 Derivations for Subsection 2.3

To understand the dual growth accounting that allows us to measure the supply-side factors that drive PCE inflation, we split the nominal parts of (8) into their price and quantity components. I denote the price of PCE, i.e., the PCEPI, by P_c and the quantity by C. Thus, nominal PCE is equal to P_cC.

Throughout my derivations, I use a continuous-time notation, which I will approximate with a Törnqvist index in the empirical implementation. The goal is to account for the supply-side factors that drive the growth rate of the PCEPI, which, in continuous time, is the change in the log of P_c, i.e., $\pi_c = \dot{p}_c$. Here $\dot{}$ denotes the time derivative in continuous time and $\pi_c = \ln P_c$. The growth of nominal PCE is the sum of inflation and the growth rate of the quantity, i.e., $\pi_c = \dot{c}$.

Nominal value added of industry i that ends up being sold to consumers after s steps along the supply chain is

$$V_s^C(i) = \boldsymbol{v}_s(i) P_C C, \tag{9}$$

where $\boldsymbol{v}_s(i)$ is the i^{th} element of $\boldsymbol{v}_s$. This makes up a fraction $F_s^V(i) = \dfrac{V_s^C(i)}{V(i)}$ of total value added of industry i.

Nominal value of imports of commodity j that end up being sold to consumers after s steps along the supply chain is

$$M_s^C(j) = \boldsymbol{m}_s(j) P_C C, \tag{10}$$

where $\boldsymbol{m}_s(j)$ is the j^{th} element of $\boldsymbol{m}_s$. This makes up a fraction $F_s^M(j) = \frac{M_s^C(j)}{M(j)}$ of imports of commodity j.

This allows us to write nominal PCE in terms of the origins of the value added it encompasses. That is, we obtain that

$$P_C C = \sum_{s=0}^{\infty} \sum_{i=1}^{n_i} F_s^V(i) V(i) + \sum_{s=0}^{\infty} \sum_{j=1}^{nj} F_s^M(j) M(j). \tag{11}$$

Taking the time derivative on both sides of this expression, we find that

$$\left(P^C C\right)\dot{p}^C + \left(P^C C\right)\dot{c} = \sum_{s=0}^{\infty} \sum_{i=1}^{n_i} (F_s^V(i) V(i)) \dot{f}_s^V(i) + \sum_{s=0}^{\infty} \sum_{j=1}^{nj} (F_s^M(j) M(j)) \dot{f}_s^M(j) \tag{12}$$

$$+ \sum_{s=0}^{\infty} \sum_{i=1}^{n_i} (F_s^V(i) V(i)) \dot{v}(i) + \sum_{s=0}^{\infty} \sum_{j=1}^{nj} (F_s^M(j) M(j)) \dot{m}(j). \tag{13}$$

When we define the shares of each of the components in nominal PCE as

$$\phi_s^V(i) = \frac{F_s^V(i) V(i)}{P^C C} \text{ and } \phi_s^M(j) = \frac{F_s^M(j) M(j)}{P^C C} \tag{14}$$

and divide both sides of this equation by the value of nominal PCE, we obtain that the growth rate of nominal PCE is a share-weighted average of the growth rates of the value-added components that flow to final demand in the form of consumption. That is,

$$\dot{p}^C + \dot{c} = \sum_{s=0}^{\infty} \sum_{i=1}^{n_i} \phi_s^V(i) \dot{f}_s^V(i) + \sum_{s=0}^{\infty} \sum_{j=1}^{nj} \phi_s^M(j) \dot{f}_s^M(j) \tag{15}$$

$$= \sum_{s=0}^{\infty} \sum_{i=1}^{n_i} \phi_s^V(i) \dot{v}(i) + \sum_{s=0}^{\infty} \sum_{j=1}^{nj} \phi_s^M(j) \dot{m}(j). \tag{16}$$

The next step is to split nominal value-added growth of each industry in a price and quantity component, i.e.

$$\dot{v}(i) = \dot{p}^V(i) + \dot{q}^V(i) \tag{17}$$

and

$$\dot{m}(j) = \dot{p}^M(j) + \dot{q}^M(j). \tag{18}$$

Doing so yields that

$$\dot{p}^C + \dot{c} = \sum_{s=0}^{\infty}\sum_{i=1}^{n_i}\phi_s^V(i)\,\dot{p}^V(i) + \sum_{s=0}^{\infty}\sum_{j=1}^{n_j}\phi_s^M(j)\,\dot{p}^M(j) \tag{19}$$

$$+\sum_{s=0}^{\infty}\sum_{i=1}^{n_i}\phi_s^V(i)\,\dot{f}_s^V(i) + \sum_{s=0}^{\infty}\sum_{j=1}^{n_j}\phi_s^M(j)\,\dot{f}_s^M(j) \tag{20}$$

$$+\sum_{s=0}^{\infty}\sum_{i=1}^{n_i}\phi_s^V(i)\,\dot{q}^V(i) + \sum_{s=0}^{\infty}\sum_{j=1}^{n_j}\phi_s^M(j)\,\dot{q}^M(j). \tag{21}$$

In the above equation, the bottom two lines have to do with the growth rates of quantities; this means that PCEPI inflation, i.e., $\dot{p_c}$, is equal to the top line, namely

$$\pi^C = \dot{p}^C = \textstyle\sum_{s=0}^{\infty}\sum_{i=1}^{n_i}\phi_s^V(i)\,\dot{p}^V(i) + \sum_{s=0}^{\infty}\sum_{j=1}^{n_j}\phi_s^M(j)\,\dot{p}^M(j). \tag{22}$$

That is, consumer price inflation is the weighted sum of value-added deflator inflation by industry and import-price inflation by commodity.

Implementing (18) empirically requires combining data on nominal imports with import prices, both by commodity.[27] However, because of a lack of the necessary detail in the data, I report the second term on the right-hand side of (22) as the residual that makes the above equation hold. This is why it is labeled "Imports and rest" in the tables. The fact that the implied $\dot{p}^M$ from this residual closely lines up with rescaled import-price inflation from the NIPA, as I show in the section on rules of thumb, confirms that this is a reasonable approximation.

Under neoclassical assumptions of constant returns to scale and perfect competition in both the product and factor-input markets, (17)

27. In practice, this turns out to be infeasible in U.S. data because of the lack of import prices by NAICS classified commodities before 2005.

can be rewritten further by using dual growth-accounting methods. In particular, these methods allow us to split inflation in the value-added deflator for industry up into the changes in factor costs for the industry and measured TFP growth:

$$\dot{p}^V(i) = \sum_l s_l^V(i)\dot{w}_l(i) - \sum_k s_k^V(i)\dot{u}c_k(i) - \dot{z}(i). \tag{23}$$

Here $s_l^V(i)$ is the factor share of labor of type l in value added and $\dot{w}(l)$ is quality-adjusted compensation growth for labor of type l in industry i. Similarly, $s_k^V(i)$ is the factor share of capital of type k in value added and $\dot{u}c_k(i)$ is the growth rate of the user cost of capital of type k industry i. The term $\dot{z}(i)$ is measured TFP growth in sector i. Combined with (17), this allows for decomposing π^c into parts due to labor, capital, and TFP in different industries and due to import prices.

The derivations here are in continuous time. Of course, in practice, the data are provided on an annual basis. Following Fleck and others (2014), I use a Törnqvist index to approximate these continuous-time equations in discrete time.

Inflation Globally

Òscar Jordà
Federal Reserve Bank of San Francisco
University of California, Davis

Fernanda Nechio
Federal Reserve Bank of San Francisco

The fortunes of the Phillips curve have ebbed and flowed ever since it was proposed by Phillips (1958). Although its origins are primarily as an empirical regularity, there is now a vast literature that provides more formal justification.[1] In recent times, the Great Moderation and the modern era of central banking brought about the apparent empirical demise of this core relationship. As central banks gained credibility and inflation targeting became widespread, inflation expectations became better anchored; however, the debate is far from settled.[2] Starting in the mid-1980s, several advanced economies have generally experienced the business cycle with barely a ripple in inflation. Paradoxically, a credible, inflation-targeting central bank that cares about the tradeoff spelled by the Phillips curve and sets policy to offset fluctuations in aggregate demand will make empirical estimates of the Phillips curve appear flatter than they really are. The Phillips curve is fundamental for a central bank to evaluate counterfactual policy outcomes. It is not intended to be a forecasting tool.

From the beginning, the standing of the Phillips curve in the macroeconomics literature has been fraught.[3] Its role in the standard New Keynesian model sat uneasily with ardent proponents of a more microfounded approach to macroeconomics coming from the Real Business Cycle tradition. Yet even today, New Keynesian monetary

1. See Galí (2008) and Woodford (2010) for a textbook review of the literature.
2. See Orphanides and Williams (2005), Gürkaynak and others (2005), and Gürkaynak and others (2010).
3. See Shimer (2017) for a review.

Changing Inflation Dynamics, Evolving Monetary Policy edited by Gonzalo Castex, Jordi Galí, and Diego Saravia, Santiago, Chile.

models featuring Phillips curve mechanisms remain mainstream in central-bank circles.[4]

The global financial crisis broke mold. Inflation surprised on the downside almost everywhere around the planet, even in economies that were seemingly unaffected by the crisis. The conversation quickly switched to a discussion of what policy measures should central banks implement to avoid outright deflation. With the crisis ten years behind us, it is only now that inflation appears to be returning to more normal levels, albeit slowly, once again defying well-worn tenets.

Of course, to a student of economic history, the behavior of inflation following the crisis was not entirely surprising. Jordà and others (2013) document that inflation usually runs low after financial crises, especially if they are preceded by a credit boom, as this one was. Even a simple average of inflation across advanced economies following a financial crisis, such as that displayed in figure 1, is sufficient to clearly illustrate this point. A credit boom gone bust depresses inflation because aggregate demand flags for an extended period of time.

Against this background, we set out to investigate the Phillips curve globally. More than evaluating its empirical merits, which we also do, we use the Phillips curve as a yardstick with which to think about inflation dynamics globally. The Phillips yardstick then is a useful way to assess and contrast the recent history of advanced and developing economies, and especially between those economies that experienced the crisis (most of the advanced world) versus those that seemingly escaped unscathed.

We think that understanding what happened before and after the crisis across advanced and developing economies, and across economies that suffered and escaped the global financial crisis sheds light on important questions about the dynamics of inflation that have generated a great deal of debate. The literature on the nexus between inflation and credit is not large, in part because up to recently, credit has been a silent bystander in many macroeconomic models. A useful exception is Gilchrist and others (2017). They argue that financially constrained firms have an incentive to raise prices in response to adverse financial or demand shocks in order to preserve internal liquidity. Such behavior would be consistent with an asymmetric Phillips curve that is flatter when credit is constrained, since inflation would appear to be less responsive to fluctuations in demand.

4. Once again, two textbook references are Galí (2008) and Woodford (2010).

Relating inflation and economic slack is, however, difficult. Aggregate demand factors will tend to push inflation and economic activity in the same direction, whereas supply cost-push factors operate in the opposite direction. Identification of the Phillips mechanism thus requires a source of exogenous variation. Our paper departs from the literature in this respect and builds, on one hand, on work by Galí and Gertler (1999), and Galí and others (2005), and on the other hand, on the work of di Giovanni and others (2009), and Jordà and others (2017). Whereas the former investigate the Phillips curve using alternative measures of slack in a general setting, the latter suggest that monetary policy from large economies can bleed through into smaller economies that manage their exchange rates tightly and thus have unintended effects on domestic policy. Such a mechanism will be useful to complement some of the potential instrumental variables used in Galí and Gertler (1999), and Galí and others (2005).

More specifically, the well-known trilemma of international finance[5] explains how such bleed-through can occur. Investors trying to arbitrage returns across similar assets tend to equalize their returns. If exchange-rate risk is removed (as it largely is when economies peg), such equalization neutralizes domestic monetary policy to a large degree. Large economies set monetary policy domestically but, through the trilemma mechanism, can have effects beyond their borders. Such occurrences are the type of exogenous variation that will permit a cleaner read of the Phillips curve mechanism.

Using an instrumental variable for economic slack sets our analysis apart from most of the literature. However, we take matters one step further. Our interest in examining the effects of the financial crisis on the Phillips curve requires that we recognize that during this period there have potentially been other forces at work. Such forces include development of global value chains (Auer and others, 2017), the role of China on global inflation (Eickmeier and Kühnlenz, 2018), international spillovers from commodity prices (Fernández and others, 2017), and so on. A finding that the slope of the Phillips curve shifted from 2008 onwards could be explained by any of these factors and not by the quick deceleration of credit that followed the crisis.

We avoid such confounding with a difference-in-differences (D-i-D) strategy. In addition to comparing Phillips curve estimates before and after the crisis by using instrumental variables, we also

5. See Obstfeld and Taylor (1998), Obstfeld and Taylor (2003); Obstfeld and others (2004), Obstfeld and others (2005), and Shambaugh (2004).

compare estimates for crisis-hit versus crisis-missed countries. That is, we ask first whether there was a measurable shift in the slope and persistence of the Phillips curve and whether this shift was comparable across subpopulations. Such an empirical strategy is not novel in the applied microeconomic literature, but it is somewhat uncommon in international macroeconomics.

The bird's eye view of inflation globally is surprising. By most accounts, the global financial crisis was an "advanced economies crisis". Many of the developing economies of Asia and Latin America, with less developed financial systems, did not experience similar credit booms to those experienced in the developed world, nor did they experience the rapid loss in employment that followed. Yet their inflation appears to have taken a hit during the crisis that is reminiscent of what happened in the developed world. More interestingly and with a few glaring exceptions, inflation globally continued to decline toward levels more consistent with what advanced economies have experienced in the past 20 years. That raises some interesting links to a nascent literature on the neutral rate of interest, and secular stagnation. This apparent widespread decline in inflation takes place despite record low levels of nominal interest rates and increasing evidence of a decline in the natural rate of interest.[6]

Our findings provide insights on some of the main hypotheses entertained in the literature. To a casual observer, the relative stability of inflation early on as economies were experiencing turmoil might have seemed consistent with the Gilchrist and others' (2017) argument that credit constrains forced firms to pass-through price increases when demand was flailing. However, we show that the evidence in support of this hypothesis is mixed. Instead, we document that there was a gradual movement in the Phillips curve toward a larger role for inflation expectations and a diminished role for backward-looking terms, at the same time that the Phillips curve became somewhat flatter. Such a trend is visible in countries both hit and missed by the crisis.

Had we only analyzed crisis-hit economies, we might have concluded that such changes were due to the financial crisis. Hence, pairing our difference-in-differences identification strategy with the trilemma instrument appears to have paid off in avoiding confounding from underlying trends in global inflation consistent with central banks' gaining greater credibility globally.

6. See Summers (2014), Carvalho and others (2016), and Eggertsson and others (2017).

It would likely be miscalculation to interpret the flattening of the Phillips curve as a license to stimulate the economy. With the crisis came severe distortions in labor markets that forced many individuals out of the labor force. As the economy recovered and the unemployment rate declined to more normal levels, some voices have pushed for maintaining accommodative monetary policy for a while longer to try and reclaim workers back into the labor force. Such a view rests on the observation that inflation has remained relatively quiescent as the labor market quickly improved, consistent with the Phillips curve being flat. But of course, the stability of inflation relies on the central bank adjusting the degree of accommodation as the economy improves. Deviating from this norm may have greater effects on inflation than is widely understood by casual observation.

1. International Evidence: Inflation before and after the Global Financial Crisis

Financial crises tend to depress inflation long after the crisis has past. This pattern was repeated in the aftermath of the global financial crisis and arguably represents a sort of natural experiment. That is, we can evaluate the effects of the crisis (and of a credit boom gone bust) on the dynamics of inflation (and hence its effects on the Phillips curve) by comparing crisis-hit countries against those that avoided it. Of course, this will require that we account for possible unrelated pre-existing trends and for spillovers or contagion effects. These are some of the issues that we discuss in more detail when we introduce the methods that we employ in our analysis.

Importantly, the recent financial crisis is the first to affect a large number of advanced economies in the era of modern central banking—a time where we tend to think that inflation expectations had been well-anchored in many economies—and where the consensus is that central banks largely responded to the crisis in the right direction, even if the question of whether they did with sufficient force remains open.

A simple way to begin our journey and to illustrate the main ideas is to use historical data. Financial crises are relatively rare events (at least in advanced economies) so it helps to have a long sample. Using historical data for 17 advanced economies available from the database described in Jordà and others (2016), figure 1 shows the average path of inflation and real GDP per capita across 17 advanced economies following past financial crises set against the average path of inflation after a non-financial crisis recession.[7,8] In addition, we show the path of inflation for economies hit by the 2008 financial crisis.

More specifically, panel (a) of figure 1 displays the path of headline CPI inflation for the crisis-hit countries since the Great Recession set against the average rate of inflation observed in normal recessions and past financial crises using data across the 17 advanced economies, observed from 1870 to 2015. The data are normalized to the rate of the average CPI inflation in 2017 across crisis-hit economies to facilitate the comparison. The panel shows that, on average, prices tend to remain subdued for several years after a financial crisis (dot-dash line). Interestingly, up to 2013, the behavior of average inflation (solid line) resembled that of a typical recession, not a financial crisis (short-dash line). Since 2013, however, prices have been slower to recover, more in line with financial crises.

Panel (b) shows the average path of real GDP per capita following typical recessions, past financial crises, and the actual experience for the crisis-hit countries since the financial crisis. The path of real GDP per capita is normalized to 100 at the start of each of these events. It shows that real GDP per capita in crisis-hit countries have been relatively slow to recover. Moreover, whereas in previous financial crises it appears that real GDP per capita recovered after about 10 years, the latest financial crisis appears to have had permanent effects on the economy.[9]

These observations have to be set against the behavior of the economy. Panel (b) of figure 1 shows that in typical recessions, real GDP per capita declines by about 2.5 percent in the first year of the recession, but by the second year, real GDP per capita is back to where it was at the peak, and continues to grow thereafter. This is shown with the dashed line. In financial crises, the economic toll is heavier and longer-lasting. Returning to peak levels takes four rather than two years, and only 10 years after the crisis begins does real GDP per capita catch up to the path we see for typical recessions. The world experience following the financial crisis is harrowing. The gap opened by the crisis has not only failed to narrow, but it appears to have continued to widen (albeit only slightly). There appears to be a permanent loss in economic welfare that is somewhat unprecedented in economic history.

7. The database is available at www.macrohistory.net/data.

8. The set of countries include: Australia, Belgium, Canada, Switzerland, Germany, Denmark, Spain, Finland, France, Italy, Japan, Netherlands, Norway, Portugal, Sweden, the United Kingdom and the Unites States.

9. Barnichon and others (2018) estimate a lifetime present-value loss of about $70,000 per person for the United States.

Figure 1. Headline CPI Inflation and Real GDP per Capita after Financial Crises and Normal Recessions

A. Headline CPI Inflation *B. Real GDP per capita*

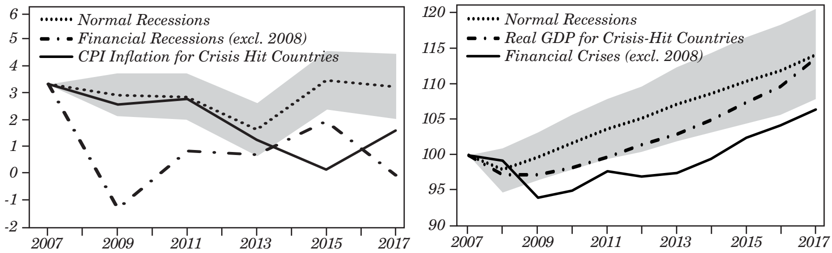

Source: Authors' calculations.
Notes: Headline CPI price index and real GDP per capita reported as an index normalized to 100 in 2007. 90% error bands reported for normal recessions as a gray shaded region.

What happened in 2008 and how does it compare with the historical record? Figure 2 provides a summary using different groupings of economies by region. Panel (a) of that figure shows the experience of the U.S. and the euro area[10] (the two largest economic zones hit by the crisis) against Australia (the reference economy selected as an example of an advanced economy that was largely unaffected by the crisis.) This will become clearer in figure 3, which shows the unemployment rate for these economies. Panel (b) simply summarizes the recent history of inflation for Asian and Latin American economies compared to advanced economies.[11]

Like figure 1, figure 2 shows that, following the crisis, inflation declined sharply in the year of the crisis, but stayed subdued until very recently (in the U.S. core PCE inflation is now near its long-run target of 2 percent). However, this time there was less outright deflation (except for the euro area, which experienced the initial hit of the financial crisis followed by the sovereign debt crisis). Australian inflation slowed down slightly, but it could be argued that this slowdown was part of a trend toward lower inflation that preceded the crisis. This is one of the possibilities that we investigate more carefully in our analysis below.

10. Euro area countries exclude Estonia, Latvia and Lithuania due to lack of data.

11. Asian economies include: China, India, Indonesia, Japan, South Korea, Malaysia, Philippines, Singapore, Thailand, and Vietnam. Latin America includes: Brazil, Chile, Colombia, Costa Rica, Mexico. Advanced economies include: Australia, Austria, Belgium, Canada, Croatia, Denmark, Estonia, Finland, France, Germany, Iceland, Ireland, Italy, Japan, Luxembourg, Netherlands, New Zealand, Norway, South Korea, Spain, Sweden, Switzerland, United Kingdom, United States, and Singapore.

Figure 2. CPI Inflation before and after the Global Financial Crisis

A. U.S., Euro area, Australia

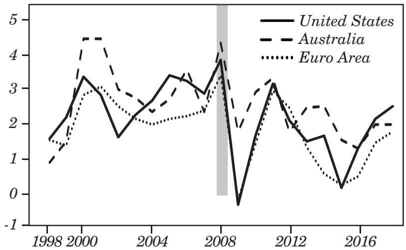

B. Asia, Advanced Economies, Latin America

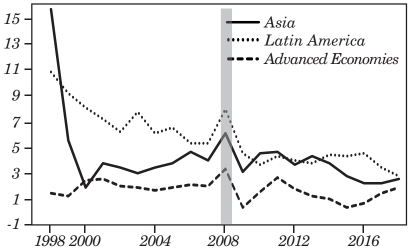

Source: Authors' calculations.
Notes: Sample: 1997–2017. Data for U.S., Euro area (Austria, Belgium, Cyprus, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, Malta, the Netherlands, Portugal, Slovakia, Slovenia, and Spain), Canada, and Australia.

Figure 3. Unemployment before and after the Global Financial Crisis

A. U.S., Euro area, Australia

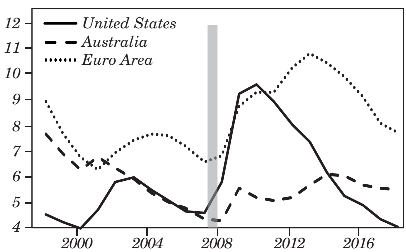

B. Asia, Advanced economies, Latin America

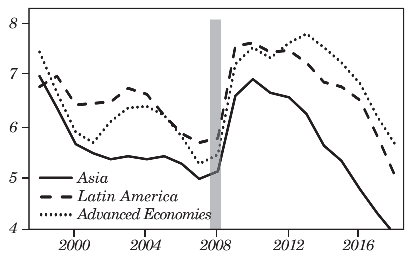

Source: Authors' calculations.
Notes: Sample: 1997–2017. Data for U.S., Euro area (Austria, Belgium, Cyprus, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, Malta, the Netherlands, Portugal, Slovakia, Slovenia, and Spain), Canada, and Australia.

Panel (b) of figure 2 also shows a general decline of inflation across the world. Even discounting some of the hyperinflation episodes early in the sample, inflation in Latin America had run consistently above 5 percent on average prior to the crisis, but has been running consistently below since then. Inflation in Asian economies, except for some high inflation episodes affecting the average at the start of the sample, has remained fairly stable throughout the period displayed.

Much, but not all of the behavior of inflation can be explained by cyclical fluctuations, here exemplified by the unemployment rate in figure 3. This figure is organized in the same manner as figure 2. Panel (a) shows the rapid and dramatic increases in the unemployment rate in the U.S. and in the euro area (especially following the sovereign debt crisis) in contrast to the considerable stability of the unemployment rate in Australia. Yet, Australia's inflation rate continued a mild decline similar to the economies of the U.S. and the euro area. Meanwhile, the declines in the inflation rate in Asian and Latin American economies visible in figure 2 seem to correspond to a generalized slowdown in economic activity in these regions. The unemployment rate also went up in these regions, in part because some countries also experienced the effects of the financial crisis outright, or indirectly through weaker global demand.

Some of the main themes to come are becoming apparent in figure 1 and figure 2. First, relative to the rapid increase in the unemployment rate depicted in figure 3, the puzzle seems to be why inflation did not fall by more early in the global financial crisis, as it had done in the previous 140 years. At this point it is probably good to remember that, at least in the U.S., the revival of productivity that stretched over the decade spanning 1995 to 2005, had already begun to quickly taper off before the crisis and had bottomed out during that time only to experience a very mild recovery more recently. Others have argued that the global trend to higher level of concentration across sectors provides a buffer against sharp declines in inflation, although this literature is still in its infancy.

Secondly, and perhaps as important for our purposes, Australia experienced a slowdown in inflation that matches the experience of crisis economies reasonably well. This pattern, however, seems to suggest little role for the traditional Phillips mechanism as Australia's unemployment rate did not move as much. Crisis countries did not see inflation dip as in previous eras, and some non-crisis countries followed a similar pattern of subdued inflation that seemed, at times, unrelated to domestic economic activity. On this evidence alone, it could be argued that, if anything, inflation expectations became increasingly well anchored throughout a very turbulent period.

2. International Inflation Dynamics: Basic Facts

As a way to summarize the dynamics of inflation globally, we pursue a straightforward strategy. Whenever we want to characterize the degree of interconnectedness of a given variable, say headline

CPI internationally, we report a 5-year rolling-window average of a country's correlation with another, that is:

$$\overline{r}_t = \frac{\sum_i \sum_j r_t^{i,j}}{N}; \quad N = \frac{(n-1)n}{2}, \tag{1}$$

where $r_t^{i,j}$ refers to the sample correlation coefficient between countries i and j, calculated over a 5-year rolling window that ends at time t. The number of countries in the sample is n, and hence N refers to the total number of distinct correlation pairs in a sample with n countries.

At times we will modify this statistic slightly. We will be interested, for example, on the correlation of headline CPI inflation for a given country against oil price inflation. And we want to report a summary statistic for the correlation observed across all countries. In that case we modify equation (1) as follows:

$$\overline{r}_t = \frac{\sum_i r_t^{i,x}}{n} \tag{2}$$

where $r_t^{i,x}$ refers to the correlation of country i against a given variable of interest x, in our example, oil price inflation. We think that these simple statistics help convey the main features of inflation globally more clearly than if we had used, for example, a factor model in which the factors would have to be "interpreted" as representing some quantity of interest.

With these preliminaries out of the way, we begin our data exploration with figure 4. The figure reports the cross-correlation of CPI inflation across economies (see equation (1) organized into two blocs: advanced vs. emerging economies.[12] A 95 percent significance band is provided as a reference in the figure. Figure 4 shows clearly that the financial crisis was an advanced-economies event. Before the crisis, inflation was loosely connected in both advanced and emerging economies (borderline significant at best), but as the crisis hit and inflation became subdued in crisis-hit economies, the separation between advanced and emerging economies became more visible.

Because we are using headline CPI inflation (the most widely available measure of inflation), we have to be mindful that some of the patterns that we observed could be explained by energy and

12. See the appendix table A1 for the list of countries under each grouping.

commodity prices. For this reason, figure 5 displays the correlation of CPI inflation with oil price inflation—using West Texas Intermediate (WTI) crude oil—and commodity price inflation. These are displayed respectively in panels (a) and (b) of the figure.

Figure 4. Headline CPI: Rolling Window average Cross-Correlation

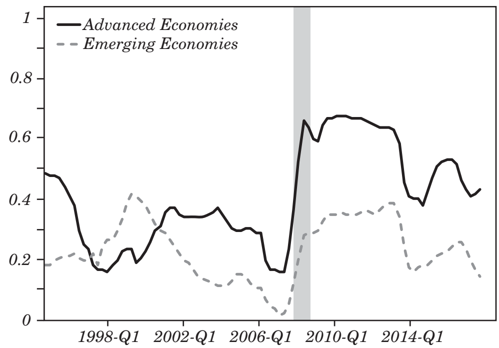

Source: Authors' calculations.
Notes: Cross-correlation calculated as in expression equation (1). Data organized into Advanced Economies and Emerging Economies according to the country grouping listed in the appendix. Dashed line is the 95% significance critical value. Shaded gray region indicates 2008 to 2009, the year of the financial crisis for most countries. See text for additional details.

Figure 5. Headline CPI: Country Cross-Correlation and Autocorrelation

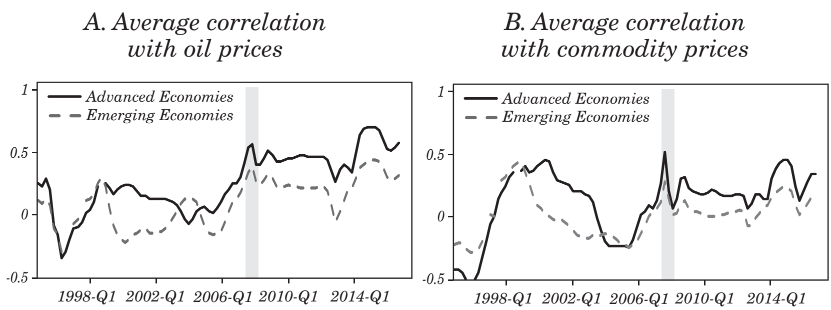

Source: Authors' calculations.
Notes: Average correlation calculated as in expression equation (2) using a 5-year rolling window. Data organized into Advanced Economies and Emerging Economies according to the country grouping listed in the appendix. Dashed line is the 95% significance critical value. Shaded grey region indicates 2008 to 2009, the year of the financial crisis for most countries. See text.

Figure 5 panel (a) shows that, although the correlation of CPI inflation with oil prices has grown over time, it was marginally significant for all economies before the financial crisis, but has grown over time more noticeably for advanced economies. It is good to remember that oil prices grew rapidly before the crisis and in the first few years thereafter, only to subsequently decline quite noticeably. Some commentators[13] in fact attribute the decline in inflation expectations to oil prices. Again, if inflation were well anchored, headline CPI inflation would primarily move with energy prices, and perhaps the increase over time of the correlation displayed in figure 5 is a reflection of that. Interestingly, later on we provide additional evidence pointing toward a better anchoring of inflation expectations since the crisis.

Our final set of charts is shown in figure 6. Panel (a) shows the cross-correlation of unemployment rates and could be interpreted as a measure of the amount of international business cycle synchronization. Meanwhile, panel (b) shows the cross-correlation in long-term bond yields (government securities with typical duration of 10 years) in an effort to examine the oft-cited synchronization in interest rates due to the decline of the real rate of interest in advanced economies reported, for example, in Holston and others (2017).

Figure 6. Headline CPI: Country Cross-Correlation and Autocorrelation

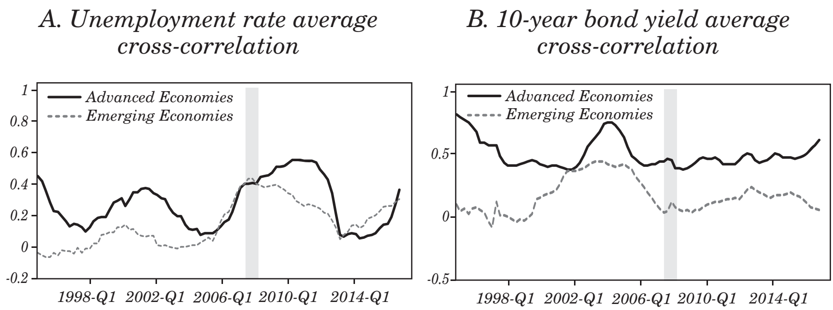

Source: Authors' calculations.
Notes: Cross correlations calculated as in expression equation (1) using a 5-year rolling window. Data organized into Advanced Economies and Emerging Economies according to the country grouping listed in the appendix. Dashed line is the 95% significance critical value. Shaded grey region indicates 2008 to 2009, the year of the financial crisis for most countries. See text for additional details.

13. See Cao and Shapiro (2016) and references therein.

Beginning with panel (a) of figure 6, evidence of synchronization in business cycles is surprisingly feeble for emerging economies, with a slight tick up during the crisis, but in any case small in economic and statistical terms. Synchronization in advanced economies is somewhat more visible, especially after the crisis, as to be expected, but it has clearly died down in recent times. Meanwhile, panel (b) suggests that, while there is considerable comovement of interest rates in advanced economies (averaging about 0.5 or higher for the entire sample), the same cannot be said for emerging economies, with almost near-zero correlation. We do not investigate this dichotomy further, but it strikes us as an interesting divergence that does not fit the neat experience of advanced economies, which have seen decline in long-term rates over the past 30–40 years.

Summing up, we use the results of these figures in a peculiar way perhaps. To anticipate our analysis, we will rely on the different experiences that advanced economies (primarily) endured during the crisis set against the experience of (mostly) emerging-market economies to try to tease out the manner in which both the crisis and the subsequent credit crunch affected the dynamics of inflation embodied by the Phillips curve.

The figures in this section show that there is a fair amount of synchronicity in advanced economies, not just in terms of inflation, but also in terms of the business cycle (as shown for unemployment rates) and in terms of interest rates. They also show that oil prices seem to be a less important source of variation in headline CPI inflation for emerging markets than for advanced economies. The contrasting experiences of advanced and emerging market economies will be helpful in getting a cleaner read on the causal effects of the crisis on inflation.

3. Identification of the Phillips Curve

Economies that endured the global financial crisis experienced rapid increases in unemployment, dramatic declines in economic activity, but remarkable stability in inflation rates. The missing disinflation shown earlier in figure 1 was a widespread phenomenon. Mechanically, such an outcome will tend to eliminate the correlation between inflation and economic slack—the unemployment rate fluctuated wildly but inflation remained relatively stable. This disconnect has been widely interpreted as a flattening of the

Phillips curve.[14] More recently, the missing disinflation has turned into the missing inflation. With most economies now recovered from the financial crisis and several running close to, or at full employment, inflation has remained relatively subdued.

Even an earlier literature has documented the empirical challenges in estimating the contributions of economic slack to inflation. These challenges include mismeasurement concerns regarding the measures of output gap, omitted variables, misspecification, estimation concerns, and difficulty in obtaining measures of inflation expectations.[15] In addition, more recently, McLeay and Tenreyro (2018) argue that most estimates of the Phillips curve insufficiently recognize the role of a history of "good monetary policy" in its estimation. In the context of a traditional New Keynesian framework and absent any supply-side shocks, an inflation-targeting central bank that conducts countercyclical policy to neutralize aggregate demand fluctuations will achieve constant inflation at the targeted level. As a result, the correlation between inflation and measures of economic slack will be zero.

These empirical challenges and the seemingly weak relationship between slack and inflation stand in sharp contrast to standard macroeconomic models, such as the New Keynesian ones.[16]

While the empirical literature presents varying reasonings behind the weak empirical properties of Phillips curves, most papers agree that the implication of their concerns is that identification of the slope of the Phillips curve requires exogenous shifts in policy that cannot be explained by aggregate demand fluctuations. In plain English, we need a valid instrumental variable. Authors have long recognized these issues and have used instrumental variable methods, including measures of oil- and commodity-price inflation to soak up cost-push fluctuations in inflation,[17] and have tried to account for changes in productivity and other supply-side factors.[18]

14. See Ball and Mazumder (2011), Blanchard and others (2015), Coibion and Gorodnichenko (2015), Dotsey and others (2017), Forbes and others (2018).

15. See Galí and Gertler (1999), Sbordone (2002), Mavroeidis (2005), Stock and Watson (2007, 2008), Coibion and Gorodnichenko (2015), Mavroeidis and others (2014), Nason and Smith (2008).

16. See Galí (2008), Smets and Wouters (2007).

17. See Roberts (1995), Galí and Gertler (1999), Galí and others (2005).

18. See Ball and Moffitt (2001).

Figure 7. Identification of the Phillips Curve Through Exogenous Policy Implementation

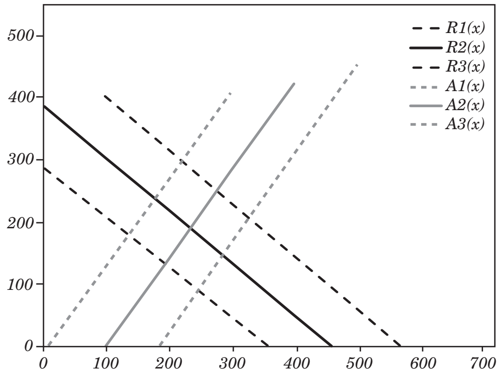

Source: Authors' calculations.
Notes: Graphical representation of equations (3) and (5) for different values of the cost-push and implementation shocks. "Incorrect Phillips Curve" refers to the Phillips curve that would be inferred from observing equilibrium outcomes. See text for additional details.

Figure 7 helps illustrate the need for instrumental variables when estimating the Phillips curve. Like McLeay and Tenreyro (2018), we borrow a simple New Keynesian model from Galí (2008) to communicate the main ideas. Consider a log-linearized New Keynesian curve given by:

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t + \varepsilon_t , \tag{3}$$

where π_t is the deviation of inflation from its target, x_t measures economic slack with the output gap, and ε_t is a cost-push shock. Assume that $\kappa > 0$.

Assume that the policymaker can directly target economic slack, so that we can set aside having to discuss the IS curve (which only provides the nexus between monetary policy and economic slack). Under discretion—McLeay and Tenreyro (2018) show that similar results can be obtained under commitment and under more general settings—, assume the policymaker minimizes the following quadratic loss:

$$L_t = \pi_t^2 + \lambda x_t^2, \tag{4}$$

where $\lambda \geq 0$ is the policymaker's preference parameter over output stabilization relative to inflation stabilization. Given equation (3), the policymaker's optimal target rule equals

$$\pi_t = -\frac{\lambda}{k} x_t + \eta_t, \tag{5}$$

where we have added the implementation error process η_t. This process naturally generates exogenous variation in the manner monetary policy is implemented.

The Phillips curve, equation (3), and the target rule, equation (5), are displayed in figure 7 in dotted and dot-dashed lines, respectively. Movements along the Phillips curve show a positive relationship between the output gap and inflation. Shocks to this curve, however, result in a negative relationship between the last two variables. If one only observes the unemployment gap and inflation equilibrium outcomes, the estimation would lead to the "incorrect Phillips Curve." The challenge is, hence, to identify movements along the Phillips curve. To do so, the econometrician can rely on shocks to the target rule, which would trace out movements along the correct Phillips curve. In very simple terms, this is what we aim for with our empirical strategy laid out in the next section.

4. Statistical Design

The previous section forcefully makes the case that, to properly estimate the Phillips curve, one requires instrumental variable methods. However, as figure 2 makes clear, declines in the rate of inflation were already evident in many economies that did not experience the financial crisis directly. If we were to estimate the Phillips curve before and after the financial crisis, we may be tempted to conclude that in its aftermath—like many times in the past—the slope has flattened. We think conclusive evidence requires that we go one step further. In addition to proposing an instrumental variable strategy, we consider a difference-in-differences approach. We discuss our choice of instrumental variable first.

4.1 The Trilemma as a Source of Exogenous Variation

One of the difficulties in extracting conclusions about the Phillips curve from raw data on inflation, expected inflation, and a measure of the economic activity gap (such as the deviation of the unemployment rate from its natural rate) is classical simultaneity bias, as we have just discussed. In other words, setting aside the role of expectations, the correlation between inflation and economic activity reflects movements along the Phillips curve, as well as shifts generated by supply shocks. We need to find an instrumental variable.

To address this identification issue in a manner that can be used across a wide swath of countries, we borrow from previous work by di Giovanni and others (2009), and Jordà and others (2015, 2017). The core idea in these papers is to exploit the trilemma of international finance. Simply stated, the trilemma suggests that countries that peg their exchange rate while allowing capital to flow freely across borders give up a great deal of monetary autonomy. The reason is that pegging the exchange rate (or even floating, but managing the exchange rate over a narrow band) mitigates exchange-rate risk considerably. Using absence of arbitrage and uncovered interest-rate-parity-based arguments, similar assets (in risk- and maturity-adjusted terms) should have similar returns across borders. Countries that peg to a base economy will therefore relinquish much of their ability to affect interest rates, since base country interest rates will largely determine domestic rates.

To implement the trilemma idea, we divide the sample into three subpopulations: base economies, pegs, and floats. Base economies in our sample will be the U.S. and Germany (for euro-area countries only). Pegs are economies that fixed their exchange rate to that of a base economy. For example, we interpret that euro-area economies peg to Germany, the largest economy in the bloc. Moreover, a cursory look at how policy rates are set in the euro area suggests that this may not be a bad approximation. We illustrate this point in figure 8.

This figure divides the euro area into core (think Germany as an example) and periphery (think Spain as an example) economies. It then calculates the prescribed policy rate based on a standard Taylor rule and compares the result with the euro-area target rate. The point that the figure makes is abundantly clear and fits with the intuition of many commentators: in the lead up to the crisis, monetary policy was too accommodative for countries like Spain (whose inflation rate doubled that of Germany during this period), but in line with what would have been optimal for Northern European economies.

Figure 8. The Trilemma in the Euro Area. An Illustration (percent)

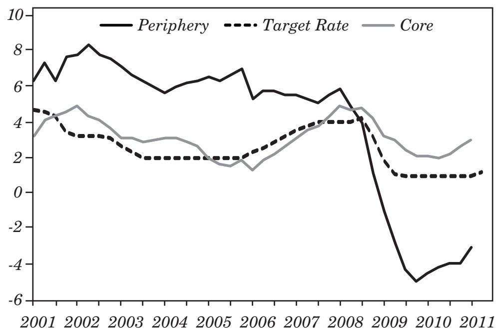

Source: OECD and Eurostat.

We also consider in our sample countries that peg to the U.S. dollar. We rely on Ilzetzki and others (2017) to sort those countries into the appropriate categories. All other economies, the floats, are assumed to allow their exchange rate to vary freely with market forces. For this reason, the pass-through of, say, U.S. monetary policy to domestic rates will be greatly attenuated if not altogether eliminated. Within this division into three subpopulations, we will also narrow our focus and further divide the data into OECD and non-OECD economies, for example. The reason is that one could argue that developing economies may be a poor control.

4.2 Difference-in-differences Estimates

The second feature that we bring to the analysis is designed to tackle the following situation. If we were to estimate the Phillips curve by using our panel of pegs before and after the financial crisis, we would have no guarantee that any changes in the Phillips curve reflected the aftereffects of the crisis as opposed to other forces that may have affected the dynamics of inflation over time. To account for that possibility, the strategy we pursue comes directly from the applied microeconomics literature. Basically, we will further divide the data into a treated group and a control group (countries that were hit by the crisis versus those that were not). Then we will compare each group's Phillips curve before and after the crisis.

The basic specification of the hybrid Phillips curve that we are interested in estimating is that given by

$$\pi_{it} = \alpha_i + \rho_1 \pi_{i,t-1} + \rho_e \pi^e_{it} + \beta x_{it} + \gamma w_t + \varepsilon_{it} \tag{6}$$

where α_i is a country-fixed effect, ρ_1 measures the accelerationist or persistent term of inflation, ρ_e measures the weight on inflation expectations, and β is the slope of the Phillips curve—x_{it} refers to the measure of slack, either the output gap or, in robustness checks provided in the appendix, the unemployment rate in deviations from the natural rate, $u_{it} - u^*_{it}$. In addition, w_t will capture fluctuations in the price of oil and in commodity prices. Because these variables are common across countries, they act as a pseudo-common factor for inflation. Below we describe these variables in more detail.

Next, define two indicator variables: One of them is $G_i \in \{0,1\}$ which selects those countries in our sample that experienced the financial crisis ($G_i = 1$) versus those that did not ($G_i = 0$). The other indicator variable is $D_t \in \{0,1\}$, which takes the value of 0 for observations preceding 2008Q1, and 1 thereafter. In other words, $D_t = I(\mathrm{t} \geq 2008\mathrm{Q}1)$. Using these indicator variables, we can now expand equation (6) as follows:

$$\begin{aligned}\pi_{it} = \alpha_i &+ \left\{\rho_1^0 \pi_{i,t-1} + \rho_e^0 \pi^e_{it} + \beta^0 x_{it}\right\} + \\ &+ \left\{\rho_1^c \pi_{i,t-1} + \rho_e^c \pi^e_{it} + \beta^c x_{it}\right\} G_i + \\ &+ \left\{\rho_1^a \pi_{i,t-1} + \rho_e^a \pi^e_{it} + \beta^a x_{it}\right\} D_i + \\ &+ \left\{\rho_1^{ac} \pi_{i,t-1} + \rho_e^{ac} \pi^e_{it} + \beta^{ac} x_{it}\right\} G_i D_t + \gamma w_t + \varepsilon_{it}.\end{aligned} \tag{7}$$

For example, focusing on the slope of the Phillips curve coefficient, $\hat{\beta}^0$ is the estimate of the baseline slope for the control group, that is, those countries that did not experience the crisis, evaluated before the crisis; $\hat{\beta}^0 + \hat{\beta}^a$ is an estimate of the slope for countries hit by the crisis, and a test of the null H_0: $\beta^c = 0$ is a test of the null that the slope of the Phillips curve between the treated group (crisis-hit economies) and the control group is the same in the period before the crisis. Next, $\hat{\beta}^0 + \hat{\beta}^a$ is an estimate of changes in the Phillips curve slope after 2008Q1 for the control group. A test of the null H_0: $\beta^a = 0$ would indicate that non-crisis economies saw no changes in the slope of the Phillips curve.

However, the key parameter that captures the effect that we pursue is $\hat{\beta}^{ac}$ and the corresponding null hypothesis is H_0: $\beta^{ac} = 0$. This null evaluates the effect of having being hit by the crisis and evaluates the slope in the post-2007 sample, stripping any changes in the slope of the Phillips curve that could have affected all economies even if the crisis had not occurred (the counterfactual). To save degrees of freedom, we assume that the fixed effects and the effect of oil and commodity prices remained unchanged throughout.

Our analysis differs from common applications of difference-in-differences estimators. Typical applications in applied microeconomics often take as the object of interest the introduction of a particular policy (the treatment), so that the change in policy is itself indicated with a binary variable. In our case, we are concerned about changes in two parameters that could have been affected by the financial crisis: the accelerationist term ρ_1 as well as the expectations term ρ^e, and the slope parameter β. In addition, there are often two periods involved, before and after the policy is implemented. In our application, we have two different samples rather than two points in time. The next few sections put all these methods to work.

5. Analysis

Below we take these ideas to the data by using a broad cross section of economies observed at quarterly frequency over the past 20 years or more. We will gradually build the analysis as follows. First we provide a description of the data, its sources and the main transformations. Next, we start from a full-sample, panel-IV estimate of the Phillips curve to draw the main features. The analysis uses a coarse breakdown of economies to spot where differences may be coming from. We follow this preliminary look at the data with a more careful difference-in-differences panel-IV analysis to examine more carefully the main hypothesis of our analysis: Did the financial crisis and subsequent credit crunch cause the Phillips curve to change? And if so, in what ways?

5.1 Data Description

The sample that we examine consists of a panel that includes 45 OECD and non-OECD countries across the world. Because our focus is on a narrow window of time, we need as large a cross section of countries to improve the precision of our estimates. For many of our economies, data is only available starting sometime in the mid-1990s. In particular, we focus on the 1986Q1 to 2018Q1 period and divide the sample in 2007Q4 to mark the financial crisis starting in 2008. We recognize that not all countries experienced the same starting date for the financial crisis, although figure 2 suggests that this characterization is not too far off the mark.

To achieve better sample sizes and consistent measures across countries, we rely on CPI as our inflation measure. Our measure of inflation expectations correspond to one-year-ahead expectations and are obtained directly from alternative sources or constructed as in Hamilton and others (2016) by using past inflation to generate out of sample forecasts. We measure economic slack with output gap, and provide results with unemployment gaps in the appendix. Output gaps are either obtained directly from alternative sources or constructed as the wedge between real GDP and its 4-quarter average (we use a similar approach to obtain unemployment gaps). Oil and commodity prices are obtained from WTI oil prices and commodity prices are measured from a future price index. Finally, because our sample includes countries that experienced periods of high inflation, we exclude from the sample country-dates in which inflation readings were above 25 percent.

The appendix provides extensive details on the samples of countries, data sources and methodologies applied to improve data quality.

5.2 Solving Attenuation Bias with Instrumental Variables

Using the trilemma logic, we preview the basic elements of the empirical strategy that is to follow with a simple example. Consider the hybrid specification of the Phillips curve:[19]

$$\pi_{it} = \alpha_i + \rho_1 \pi_{i,t-1} + \rho_e \pi^e_{it} + \beta x_{it} + \gamma w_t + \varepsilon_{it}$$

where π_{it} refers to quarterly headline CPI inflation expressed in percent annually, π^e_{it} refers to time t projected future expected inflation, and x_{it} is the output gap expressed in percent (in the appendix, the unemployment gap is given by the difference between the unemployment rate expressed in percent with respect to its natural rate). We do not constrain the coefficients on the lagged inflation and inflation expectations to add up to one although this is an interesting

19. See Galí and Gertler (1999), Galí and others (2005).

reference value to compare to. When $\rho_1 = 1$ and $\rho_e = 0$ we get the accelerationist version of the Phillips curve.[20] When $\rho_1 = 0$ and $\rho_e = 1$ we have a modern version of the expectations-based Phillips curve.

We estimate equation (6) by two-stage least squares (TSLS). We instrument the output gap using interest rates from the base economy to which each country pegs. Specifically, for countries in the euro area, we chose the German 10-year Bund rate. The reason is that this maturity never quite reached the zero lower bound, so it will be a natural stand-in for the policy rate for Germany. Naturally, movements in this rate reflect movements in the premiums. However, as the sovereign debt crisis showed, these premiums are probably small enough that movements in the rate are a reasonable proxy for policy movements.

All other pegging economies fixed their exchange rate to the U.S. Following Swanson and Williams (2014), we use the 2-year T-bond rate. Swanson and Williams (2014) show that this rate captures policy movements quite well during the period in which the funds rate hit the zero lower bound. Gertler and Karadi (2015) use a similar strategy based on the 1-year T-Bill rate instead, but the differences are minor. The 2-year T-bond rate will be our instrument for the subset of pegs to the dollar.

Table 1 provides a first look at the Phillips curve by reporting estimates of equation (6) by using full-sample panels using all countries (ALL), OECD economies except Hungary, Greece, Latvia and Lithuania (OECD) and a third group containing all remaining non-OECD economies (Non-OECD). A few findings deserve comment.

First, the IV first-stage F-statistic is quite high, thus denoting that the instruments are highly relevant. This is perhaps not too surprising since the endogenous variables are very persistent, so lagged information usually provides a good prediction. Second, there are visible differences between the OLS and the IV estimates, thus suggesting that OLS estimates are biased, as our previous discussion already intimated. Third, although the persistence and expectations terms do not add up to one, they are reasonably close in economic terms and certainly in statistical terms. Fourth, the expectations term is quite important across all economies. Fifth, as many before us have documented, estimates of the slope of the Phillips curve are generally economically and statistically close to zero, although with the correct sign. The appendix provides estimates based on the unemployment gap, which tell a similar story.

20. See Phelps (1967) and Friedman (1968).

Table 1. Hybrid Phillips Curve Using the Output Gap. Full Sample
(1986Q1-2018Q1)

	ALL		*OECD*		*Non-OECD*	
	OLS	*IV*	*OLS*	*IV*	*OLS*	*IV*
Persistence (ρ_1)	0.77***	0.66***	0.69***	0.87***	0.79***	0.41***
	(0.03)	(0.04)	(0.03)	(0.04)	(0.03)	(0.12)
Inflation expectations (ρ_e)	0.19***	0.34***	0.25***	0.06**	0.18***	0.78***
	(0.04)	(0.06)	(0.04)	(0.03)	(0.05)	(0.19)
Slack (β)	0.03***	0.03**	0.03***	0.02***	0.03*	0.00
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.04)
Observations	3,224	2,978	1,835	1,795	1,389	1,183
Adjusted R^2		0.91		0.92		0.73
K–P Stat		6.67		12.83		2.83
1st-stage F-stat *p*-value		0.00		0.00		0.00

Source: Authors' calculations.
Notes: Estimates based on equation (6) using OLS and TSLS. Sample includes all countries, OECD economies (excluding Hungary, Greece, Latvia, and Lithuania), and non-OECD economies. Panel estimates using fixed effects. Clustered robust standard errors. */**/*** indicates significance at the 90/95/99% confidence level. See text.

Table 2. Hybrid Phillips Curve Using the Output Gap Before Crisis
(1986Q1-2007Q4)

	ALL		*OECD*		*Non-OECD*	
	OLS	*IV*	*OLS*	*IV*	*OLS*	*IV*
Persistence (ρ_1)	0.76***	0.66***	0.71***	0.96***	0.77***	0.64***
	(0.03)	(0.05)	(0.04)	(0.03)	(0.04)	(0.09)
Inflation expectations (ρ_e)	0.21***	0.34***	0.24***	0.03	0.22***	0.45***
	(0.04)	(0.05)	(0.05)	(0.06)	(0.05)	(0.11)
Slack (β)	0.03***	0.03*	0.05***	0.02***	0.01	-0.02
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.04)
Observations	2,037	1,811	1,261	1,221	776	590
Adjusted R^2		0.89		0.91		0.87
K–P Stat		4.78		13.77		2.74
1st-stage F-stat *p*-value		0.00		0.00		0.00

Source: Authors' calculations.
Notes: Estimates based on equation (6) using OLS and TSLS. Sample includes all countries, OECD economies (excluding Hungary, Greece, Latvia, and Lithuania), and non-OECD economies. Panel estimates using fixed effects. Clustered robust standard errors. */**/*** indicates significance at the 90/95/99% confidence level. See text.

Table 3. Hybrid Phillips Curve Using the Output Gap After Crisis
(2008Q1-2018Q1)

	ALL		*OECD*		*Non-OECD*	
	OLS	*IV*	*OLS*	*IV*	*OLS*	*IV*
Persistence (ρ_1)	0.74***	0.50***	0.62***	0.49***	0.76***	0.54***
	(0.06)	(0.09)	(0.03)	(0.09)	(0.06)	(0.11)
Inflation expectations (ρ_e)	0.25**	0.80***	0.41***	0.74***	0.23*	0.72***
	(0.10)	(0.16)	(0.05)	(0.27)	(0.11)	(0.17)
Slack (β)	0.02**	0.03	0.01	0.01	0.03*	0.03
	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.02)
Observations	1,187	1,167	574	574	613	593
Adjusted R^2		0.83		0.88		0.84
K–P Stat		6.99		8.99		4.18
1st-stage F-stat *p*-value		0.00		0.00		0.00

Source: Authors' calculations.
Notes: Estimates based on equation (6) using OLS and TSLS. Sample includes all countries, OECD economies (excluding Hungary, Greece, Latvia, and Lithuania), and non-OECD economies. Panel estimates using fixed effects. Clustered robust standard errors. */**/*** indicates significance at the 90/95/99% confidence level.

Breaking down the sample before and after the financial crisis, table 2 and table 3 provide useful insights. Before the financial crisis, the persistence term is considerably larger and the expectations term much smaller. This is true across the board although more so in OECD economies. Consistent with this observation, estimates of the slope of the Phillips curve are bigger and almost always significant. In contrast, estimates based on the sample following the crisis (table 3) indicate that expectations became better anchored, and the persistence parameter became much smaller and the Phillips curve flatter. Thus, this first pass of the data provides already some support to the notion that the Phillips curve might have evolved around the time of the financial crisis, in part perhaps because of the credit crunch that followed it. The next section builds on this basic setup to obtain a more careful measure of the effect of the crisis on the Phillips curve.

5.3 Difference-in-differences Results

Having shown the attenuation bias from using OLS vs. IV when estimating the Phillips curve, we move now toward the main hypothesis of interest. That is, how did the financial crisis affect inflation dynamics? Did the credit crunch that followed the crisis boost the role of the slack component in the Phillips curve? Or did the relative stability of inflation throughout a period of considerable turmoil boost the public's confidence in the ability of central banks to keep inflation in check? Or did the crisis have no effect on the Phillips curve?

Table 4 reports estimates of equation (7) using the sample of all countries. The table is organized into three blocks of columns referring to estimates for each of the parameters of interest in the Phillips curve using a variety of methods. The column ALL-OLS uses panel fixed effect estimates using the float and the peg subpopulations (the subpopulation of floats is too small to provide reliable estimates), the column Peg-OLS uses panel fixed effects but using the subpopulation of pegging economies, and the column Peg-IV uses panel IV with fixed effects using the subpopulation of pegs. The reason for these three alternatives is to show that OLS estimates based on the entire population or the subpopulation of pegs are very similar to each other. However, the instrument only operates for the subpopulation of pegs. Hence the third column of each block is meant to display the attenuation bias of OLS vs. IV estimation.

Next, the table is divided into three blocks of rows. Panel (a) refers to estimates for the subpopulation of economies that did not experience the global financial crisis; panel (b) refers to the subpopulation of crisis-hit economies. Within each of these two blocks, we report estimates based on a sample preceding the financial crisis and labeled *Before*, and then using a sample following the crisis, labeled *After*. The row labeled *Diff* then collects the difference in the coefficients before and after.

The difference-in-differences measure of the treatment effect is reported in the third block of rows in panel (c). It measures the difference between changes in the parameters before and after the crisis for each of the subpopulations considered (crisis-hit vs. crisis-missed countries). One way to think about this measure is as a counterfactual. If the trends in non-crisis-hit economies had also been present in crisis-hit economies, what would we have expected the coefficients to look like? And hence, how do they differ relative to the coefficients we actually estimated?

Table 4 nicely sets the stage. Comparing estimates of the Phillips curve parameters between the ALL-OLS and Peg-OLS columns, it is fair to say that the differences are relatively small, in almost all the cases, well within the margin of error. It is safe to conclude that there are no major differences between the subpopulations of floats and pegs. Hence any conclusions obtained by using IV estimates can probably be extrapolated to characterize the subpopulation of economies that float their exchange rate. To use the language of the treatment evaluation literature, because the instrument only works for the subpopulation of pegs, what we are doing, properly speaking, is estimating a local average treatment effect, or LATE. We return to this issue momentarily.

Table 4. Difference in differences estimation. Phillips Curve Using Output Gap. All Countries

	Persistence (ρ_1)			*Inflation expectations* (ρ_e)			*Slack* (β)		
	All-OLS	*Peg-OLS*	*Peg-IV*	*All-OLS*	*Peg-OLS*	*Peg-IV*	*All-OLS*	*Peg-OLS*	*Peg-IV*
	(a) No-Crisis economies, $G_i = 0$								
Before	0.76***	0.79***	1.02***	0.22***	0.18**	-0.11	0.04	0.06***	0.09***
	(0.04)	(0.06)	(0.18)	(0.06)	(0.08)	(0.28)	(0.03)	(0.02)	(0.02)
After	0.70***	0.74***	0.60***	0.34***	0.25**	0.46***	0.04	0.07**	0.09***
	(0.06)	(0.07)	(0.04)	(0.09)	(0.12)	(0.12)	(0.03)	(0.03)	(0.01)
(i) Diff.	-0.06	-0.05	-0.42**	0.12*	0.07	0.56***	-0.00	0.01	0.01
	(0.04)	(0.03)	(0.20)	(0.06)	(0.05)	(0.21)	(0.01)	(0.01)	(0.02)
	(b) Crisis economies, $G_i = 1$								
Before	0.66***	0.63***	0.64***	0.35***	0.39***	0.40*	0.01	0.01	0.01
	(0.04)	(0.04)	(0.17)	(0.04)	(0.04)	(0.21)	(0.01)	(0.01)	(0.02)
After	0.60***	0.60***	0.53***	0.49***	0.49***	0.63***	0.01	-0.00	-0.00
	(0.02)	(0.02)	(0.05)	(0.03)	(0.05)	(0.13)	(0.01)	(0.01)	(0.02)
(ii) Diff.	-0.06	-0.03	-0.11	-0.13***	0.11**	0.23**	0.00	-0.01	-0.01
	(0.04)	(0.03)	(0.12)	(0.05)	(0.05)	(0.09)	(0.02)	(0.02)	(0.02)
	(c) Treatment effect: (ii) – (i)								
D-i-D:	-0.00	0.02	0.31	0.02	0.04	-0.33	0.00	-0.02	-0.02
	(0.06)	(0.05)	(0.28)	(0.08)	(0.07)	(0.26)	(0.02)	(0.02)	(0.03)
Observations	3,464	2,228	2,134						
Adjusted R^2	0.88	0.91	0.87						

Source: Authors' calculations.
Notes: Sample includes all countries. Clustered robust standard errors. */**/*** indicates significance at the 90/95/99% confidence level. See text.

For now, the table reveals, however, that not using an instrumental variable approach can result in considerable attenuation bias. Estimates for the persistence parameter for the crisis-missed economies go from 0.79 to 1.02, for example. We find similar attenuation in several other estimates, consistent with earlier results.

The more economically interesting results are hence contained in the Peg-IV columns. Note that, if we had considered only the crisis-hit economies (panel (b) of the table), we would have concluded that (i) persistence declined from 0.64 to 0.53, a decline of 0.11 and not statistically significant; (ii) the role of expectations measured by estimates of the parameter ρ_e went up in similar proportion, from 0.40 to 0.63, an increase of 0.23 and significant; and that (iii) in both periods the Phillips curve remained mostly flat. It would thus be tempting to conclude that, as a result of the crisis, expectations became better anchored. However, a similar story took place in economies that did not experience the financial crisis, as panel (a) of table 4 shows. If anything, the decline in the accelerationist term is much larger and so is the increase in the weight that inflation expectations now receive. As we will see shortly, this difference is driven by differences between OECD and non-OECD economies.

Hence, the estimate of the difference-in-differences (local) treatment effect of the financial crisis reported in panel (c) of the table 4 is critical. In terms of broad trends, we find no actionable statistical evidence indicating substantial changes in crisis-hit economies relative to economies that escaped the crisis. However, it is clear that both types of economies saw changes in the Phillips curve, indicative of better anchoring to inflation expectations. One should be careful though. Although the D-i-D estimates are not significant statistically speaking, the magnitudes are relatively sizable: a reduction of persistence of about 0.3 that translated into a boost of similar magnitude to the coefficient on expectations. That is our best measure of the effect of the crisis on the Phillips curve based on full-sample results.

To better visualize the results reported in table 4, figure 9 presents graphically the estimates of crisis/no-crisis economies, before and after the crisis, and the counterfactual path that allows us to measure the treatment effect. Panel (a) of the figure corresponds to estimates of the persistence parameter ρ_1 and panel (b) corresponds to estimates of the expectations parameter ρ_e. The line denoted "No crisis" corresponds to estimates before and after the crisis for countries that did not experience the crisis. Note the decline in the persistence parameter and the increase in the expectations parameter. The line denoted

"Crisis" corresponds to the subpopulations of crisis-hit economies. The decline in the persistence parameter is also visible but it is more muted. Similarly, the increase in the expectations parameter is still there, but is also more muted. The grey dashed line is the counterfactual path that crisis-hit countries would have been expected to follow had they shared the same trends as the crisis-missed countries. As we remarked earlier, the effect is quite sizable economically speaking, though imprecisely estimated. The dashed vertical line visually represents the difference-in-differences estimates reported in panel (c) of table 4.

In economic terms, figure 9 neatly shows how the crisis affected the Phillips curve. Generally speaking, leading into the crisis, all countries were experiencing a boost to the expectations term and a concomitant decline in the accelerationist term consistent with a flattening of the Phillips curve. A natural explanation of all these developments is that central banks were generally becoming more credible. Crisis-hit economies saw that trend slow down considerably relative to non-crisis-hit economies, consistent with Gilchrist and others (2017). The effects are economically sizable although not estimated precisely enough to provide statistically conclusive evidence.

Figure 9. Less Persistence, Better Anchoring. Difference-Difference Measures of the Effect of the Crisis. Full Sample.

A. Persistence paramenter ρ_1

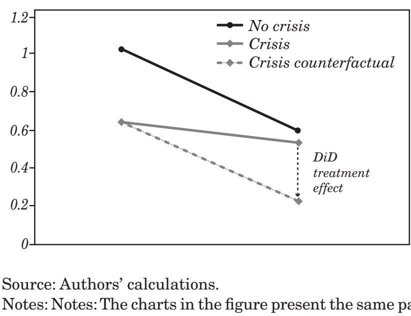

B. Expectations parameter ρ_e

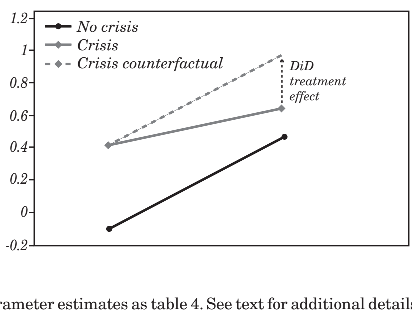

Source: Authors' calculations.
Notes: Notes: The charts in the figure present the same parameter estimates as table 4. See text for additional details.

Table 5. Difference in differences estimation. Phillips Curve Using Output Gap.
(OECD countries excluding Hungary, Greece, Latvia and Lithuania)

	Persistence (ρ_1)			*Inflation expectations* (ρ_e)			*Slack* (β)		
	All-OLS	*Peg-OLS*	*Peg-IV*	*All-OLS*	*Peg-OLS*	*Peg-IV*	*All-OLS*	*Peg-OLS*	*Peg-IV*
	(a) No-Crisis economies, G_i =0								
Before	0.54***	0.49***	0.68***	0.47***	0.51***	0.25	0.04***	0.05***	0.04***
	(0.05)	(0.06)	(0.14)	(0.07)	(0.13)	(0.18)	(0.01)	(0.01)	(0.01)
After	0.56***	0.56***	0.56***	0.51***	0.48***	0.44***	0.04***	0.05***	0.05***
	(0.04)	(0.07)	(0.05)	(0.07)	(0.13)	(0.09)	(0.01)	(0.02)	(0.02)
(i) Diff.	0.02	0.07***	-0.12	0.05	-0.02***	0.18*	0.00	0.00	0.00
	(0.06)	(0.01)	(0.11)	(0.06)	(0.01)	(0.11)	(0.01)	(0.01)	(0.01)
	(b) Crisis economies, G_i =1								
Before	0.70***	0.67***	0.81***	0.29***	0.32***	0.13	0.02**	0.03***	0.04***
	(0.03)	(0.05)	(0.07)	(0.03)	(0.04)	(0.09)	(0.01)	(0.01)	(0.01)
After	0.61***	0.61***	0.62***	0.46***	0.46***	0.39***	0.03**	0.01	0.03***
	(0.02)	(0.02)	(0.03)	(0.04)	(0.02)	(0.10)	(0.01)	(0.01)	(0.01)
(ii) Diff.	-0.10**	-0.06*	-0.19***	0.18***	0.14***	0.27***	0.01	-0.01	-0.01
	(0.04)	(0.03)	(0.06)	(0.05)	(0.03)	(0.07)	(0.02)	(0.02)	(0.01)
	(c) Treatment effect: (ii) – (i)								
D-i-D:	-0.11*	-0.13***	-0.07	0.13*	0.27***	0.08	0.00	-0.02	-0.01
	(0.07)	(0.03)	(0.11)	(0.08)	(0.03)	(0.10)	(0.02)	(0.02)	(0.02)
Observations	1,760	1,146	1,146						
Adjusted R^2	0.88	0.89	0.89						

Source: Authors' calculations.
Notes: Sample includes OECD countries except Hungary, Greece, Latvia, and Lithuania. Clustered robust standard errors. */**/*** indicates significance at the 90/95/99% confidence level. See text.

It is well known, however, that the crisis was primarily an advanced-economies event. Moreover, advanced economies—more than emerging markets—have a longer tradition of independent central banks. These features show up clearly in the results reported in table 5. The table is organized exactly as table 4. Hence we can focus directly on panel (a) of table 5, which shows that in the few OECD economies not hit by the crisis (think primarily of Canada, Australia, and New Zealand), there was a small decline in the weight of the accelerationist term (from 0.68 to 0.56 for a difference of 0.12), with a similarly small boost to the weight on the expectations terms (from 0.25 to 0.44 for a difference of 0.18).

The values reported for the crisis-hit economies in panel (b) of table 5 are very similar indeed. The weight on the persistence parameter declines from 0.81 to 0.62 or a drop of 0.19 and significant. The weight on the expectations term goes from 0.13 to 0.39 or an increase of 0.27 and significant. Interestingly, without constraining the coefficients to be so, the weight on lagged inflation and the expectations terms sum up to close to 1, as the theory prescribes. In addition, the slope of the Phillips curve appears far more stable. Both types of countries have a similar slope (in the order of 0.04) and it does not change meaningfully from one group to the other or before and after the crisis. It is no surprise that the treatment effect of the crisis in all cases is essentially zero. That is, for OECD economies the evidence suggests that the trends set in motion before the crisis explain changes in the Phillips curve and we see essentially no evidence that the crisis itself had any lasting effects on the Phillips curve.

The natural complement to table 4 and table 5 is provided in table 6. The results presented in the table make clear that the experiences of OECD and non-OECD economies were quite different. Persistence generally declines after the crisis, but it declines much more for non-crisis-hit economies and hence the differences-in-differences estimate is 0.27—sizable economically, although imprecisely estimated and of a similar magnitude to the effect estimated using all countries. This is almost the mirror image of what happens with the expectations term, which gains in importance after the crisis and the difference-in-differences effect at -0.30, which nearly matches, but with the opposite sign, what happened with the persistence parameter. Interestingly, although we cannot see any effects of the crisis on the difference-in-differences estimate, it is clear that non-crisis-hit economies have a much steeper Phillips curve (at 0.09 and significant) versus crisis-hit economies (-0.01 and not significant).

Table 6. Difference in differences estimation. Phillips Curve Using Output Gap. Non-OECD

	Persistence (ρ_1)			*Inflation expectations* (ρ_e)			*Slack* (β)		
	All-OLS	*Peg-OLS*	*Peg-IV*	*All-OLS*	*Peg-OLS*	*Peg-IV*	*All-OLS*	*Peg-OLS*	*Peg-IV*
	(a) No-Crisis economies, G_i =0								
Before	0.76***	0.80***	1.01***	0.22***	0.17**	-0.08	0.04	0.06**	0.09***
	(0.04)	(0.06)	(0.17)	(0.06)	(0.09)	(0.25)	(0.03)	(0.02)	(0.01)
After	0.71***	0.75***	0.64***	0.33***	0.24**	0.42***	0.04	0.07**	0.09***
	(0.06)	(0.08)	(0.06)	(0.10)	(0.12)	(0.09)	(0.04)	(0.03)	(0.02)
(i) Diff.	-0.05	-0.05	-0.36*	0.11*	0.07	0.50**	-0.00	0.01	0.00
	(0.04)	(0.03)	(0.21)	(0.07)	(0.05)	(0.23)	(0.02)	(0.01)	(0.02)
	(b) Crisis economies, G_i =1								
Before	0.56***	0.57***	0.62***	0.48***	0.47***	0.42***	-0.01	-0.01	-0.01
	(0.02)	(0.02)	(0.07)	(0.02)	(0.02)	(0.08)	(0.01)	(0.01)	(0.01)
After	0.56***	0.57***	0.53***	0.58***	0.56***	0.62***	-0.01	-0.01*	-0.01
	(0.02)	(0.02)	(0.03)	(0.06)	(0.06)	(0.04)	(0.01)	(0.01)	(0.02)
(ii) Diff.	0.00	0.00	-0.09*	0.10	0.09	0.20***	0.00	-0.00	-0.01
	(0.03)	(0.03)	(0.05)	(0.07)	(0.07)	(0.07)	(0.01)	(0.01)	(0.01)
	(c) Treatment effect: (ii) – (i)								
D-i-D:	0.06	0.05	0.27	-0.01	0.03	-0.30	0.00	-0.01	-0.01
	(0.05)	(0.05)	(0.23)	(0.10)	(0.08)	(0.24)	(0.02)	(0.02)	(0.03)
Observations	1,704	1,082	988						
Adjusted R^2	0.89	0.91	0.88						

Source: Authors' calculations.
Notes: Sample excludes OECD countries. Clustered robust standard errors. */**/*** indicates significance at the 90/95/99% confidence level. See text.

The increase in the expectations term in crisis-hit economies is quite large at 0.50, even though there is no evidence that the slope of the Phillips curve became any flatter. At 0.09, the slope of the Phillips curve is reminiscent of the values observed for the U.S. in the mid-1980s and earlier. Crisis-hit economies started from a much different position and perhaps this in part explains why the changes in the main coefficients are more muted. The coefficient on the expectations term was already at 0.42 before the crisis (0.62 for the persistence or accelerationist parameter, thus adding almost exactly to 1) and with a completely flat Phillips curve. After the crisis the expectations term increases from 0.42 to 0.62 (and the persistence parameter declines from 0.62 to 0.53), which reinforces the earlier message that the crisis mainly affected advanced economies, which in turn operated with a more modern Phillips curve to begin with.

6. Conclusion

The aftereffects of the global financial crisis cannot be overstated. Ten years after the crisis, many economies remain far from their pre-crisis trend growth paths. The permanent losses of income are quite dramatic and stand out in comparison to similar events in the history of modern finance, as figure 1 showed. Inflation after financial crises tends to fall sharply and remain subdued for extended periods of time. This time inflation declined far less than in previous times, but it has been difficult to arouse from its torpor. It is still early to know whether inflation is now turning a corner. Such features of the recent inflation experience have to be set against a background in which central banks were gaining increasing independence and credibility[21]. Our goal has been to examine what, if any, were the consequences to inflation dynamics from the unfolding collapse in credit and subsequent dip in demand, all within the global context.

The financial crisis was a global event. Even countries outside its destructive path were nevertheless buffeted by its downdraft. Economies are interconnected through a variety of channels that we discussed earlier, perhaps now to a greater extent than ever before. Inflation in advanced economies now moves in greater unison than in the past (figure 4), and it is more sensitive to movements in oil prices (figure 5). Such a development is consistent with modern central banking. The ability to better execute countercyclical policy to smooth shocks to aggregate demand probably makes inflation more correlated with oil prices and other transitory shocks.

Despite its international dimension, it is clear that advanced economies differ considerably from emerging markets. The business cycle, inflation, and interest rates are far more synchronized in the advanced world than in developing economies (figure 6). That said,

21. See Den Haan and others (2017).

some of the trends advanced economies experienced in earlier decades are now becoming more clearly visible in emerging markets.

That much is clear from our analysis. Such views influenced our empirical investigation of inflation dynamics globally. Our premise was to examine in what ways, if any, the financial crisis had affected the basic Phillips mechanism. But that investigation had to be set against the background of an evolving Phillips curve and in an environment of increasing central-bank independence. Central-bank credibility has largely made inflation less susceptible to aggregate demand shocks. Short-run fluctuations in inflation largely come from transitory factors to which central banks rarely feel the need to respond to.[22]

We showed the importance of being careful about identification along two fronts. First by recognizing that OLS estimates of the Phillips curve suffer from classical simultaneity bias. We addressed this bias by using instrumental variable methods based on the trilemma of international finance.[23] We showed that ignoring simultaneity can result in considerable attenuation bias of the main coefficients of interest.

However, trends set in motion in the decades leading to the crisis could also obfuscate its true impact on inflation dynamics. Here again we resorted to another tool from the treatment evaluation literature: a difference-in-differences strategy. Economically speaking, a persuasive mechanism that could explain the relative firmness of inflation in the early days of the crisis has to do with how the collapse in credit affected pricing decisions, as suggested, for example, by Gilchrist and others (2017). We find that the evidence is somewhat mixed. Statistically speaking, estimates for this mechanism are imprecise though economically sizable (figure 9). However, one has also to contend with the observation that inflation dynamics in advanced economies were already determined, to a great degree, by an expectations mechanism. The accelerationist term had already declined considerably and the slope of the Phillips curve was basically flat, as many others had documented.[24]

What should policymakers take away from our analysis? Inflation is currently subdued, often running below levels many central banks target explicitly or implicitly. It would be tempting to take advantage of this circumstance to try to stimulate the economy above its potential.

22. See Ball and Mazumder (2018).
23. See Ball and Mazumder (2018).
24. See Ball and Mazumder (2011).

However, this seems ill-advised. The decline of inflation persistence and the increasing weight of inflation expectations, together with a relatively flat Phillips curve, are all consistent with central banks' being generally seen as credible inflation-fighters. If central banks start behaving differently than the public expects, it would be reasonable to then expect a reversal of these trends. That would make inflation misses and transitory factors more persistent, and it would make it costlier to rein in runaway inflation expectations.

Inflation globally also appears to have a stronger common component. In part this probably reflects the crisis itself. In part it probably also reflects common transitory factors (such as oil prices) that affect inflation across borders in an environment where central banks effectively implement countercyclical policy. Meanwhile, inflation has been declining almost everywhere (figure 2). No doubt part of this decline reflects more effective policymaking. However, others have pushed forward alternative explanations based on demographic factors,[25] secular stagnation,[26] and so on. We leave to others to contrast these theories against one another.

25. See Carvalho and others (2016).
26. See Summers (2014).

References

Auer, R., C. Borio, and A.J. Filardo. 2017. "The Globalisation of Inflation: The Growing Importance of Global Value Chains." Technical report, CEPR.

Ball, L. and S. Mazumder. 2011. "Inflation Dynamics and the Great Recession". *Brookings Papers on Economic Activity* 42(1): 337–405.

———. 2018. "A Phillips Curve with Anchored Expectations and Short-term Unemployment. *Journal of Money, Credit and Banking* 51(1).

Ball, L. and R. Moffitt. 2001. "Productivity Growth and the Phillips Curve. Working Papers No. 8421, National Bureau of Economic Research.

Barnichon, R., C. Matthes, and A. Ziegenbein. 2018. "The Financial Crisis at 10: Will We Ever Recover? FRBSF Economic Letter No. 2018-19.

Blanchard, O., E. Cerutti, and L.H. Summers. 2015. "Inflation and Activity: Two Explorations and Their Monetary Policy Implications." Working Paper Series WP15-19, Peterson Institute for International Economics.

Cao, Y. and A.H. Shapiro. 2016. "Energy's Impact on Inflation Expectations. FRBSF Economic Letter No. 2016-19.

Carvalho, C., A. Ferrero, and F. Nechio. 2016. "Demographics and Real Interest Rates: Inspecting the Mechanism. *European Economic Review* 88(C): 208–26.

Coibion, O. and Y. Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics* 7(1): 197– 232.

Den Haan, W., E. Ilzetzki, M. McMahon, and R. Reis. 2017. "Is the Era of Central Bank Independence Drawing to a Close? LSE Business School Blog. Available at http://eprints.lse.ac.uk/79768/1/blogs.lse.ac.uk-Is%20the%20era%20of%20central%20bank%20independence%20drawing%20to%20a%20close.pdf.

Di Giovanni, J., J. McCrary, and T. von Wachter. 2009. "Following Germany's Lead: Using International Monetary Linkages to Estimate the Effect of Monetary Policy on the Economy." *Review of Economics and Statistics* 91(2): 315–31.

Dotsey, M., S. Fujita, and T. Stark. 2017. "Do Phillips Curves Conditionally Help to Forecast Inflation?" Working Papers No. 17-26, Federal Reserve Bank of Philadelphia.

Eggertsson, G.B., N.R. Mehrotra, and J.A. Robbins. 2017. "A Model of Secular Stagnation: Theory and Quantitative Evaluation." Working Papers No. 23093, National Bureau of Economic Research.

Eickmeier, S. and M. Kühnlenz. 2018. "China's Role in Global Inflation Dynamics." *Macroeconomic Dynamics* 22(2): 225–54.

Fernández, A., S. Schmitt-Grohé, and M. Uribe. 2017. "World Shocks, World Prices, and Business Cycles: An Empirical Investigation." *Journal of International Economics* 108(S1): S2–S14.

Forbes, K., L. Kirkham, and K. Theodoridis. 2018. "A Trendy Approach to UK Inflation Dynamics." CEPR Discussion Papers No. 12652.

Friedman, M. 1968. "The Role of Monetary Policy." *American Economic Review* 58: 1–17.

Galí, J. 2008. "Introduction to Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework." In *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press.

Galí, J. and M. Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44(2): 195–222.

Galí, J., M. Gertler, and J.D. López-Salido. 2005. "Robustness of the Estimates of the Hybrid New Keynesian Phillips Curve." *Journal of Monetary Economics* 52(6): 1107–18. The Econometrics of the New Keynesian Price Equation.

Gertler, M. and P. Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7(1): 44–76.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek. 2017. "Inflation Dynamics during the Financial Crisis." *American Economic Review* 107(3): 785–823.

Gürkaynak, R.S., A. Levin, and E. Swanson. 2010. "Does Inflation Targeting Anchor Long-run Inflation Expectations? Evidence from the U.S., UK, and Sweden." *Journal of the European Economic Association* 8(6): 1208–42.

Gürkaynak, R.S., B. Sack, and E. Swanson. 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking* 1(1): 55–93.

Hamilton, J.D., E.S. Harris, J. Hatzius, and K.D. West. 2016. "The Equilibrium Real Funds Rate: Past, Present, and Future." *IMF Economic Review* 64(4): 660–707.

Holston, K., T. Laubach, and J.C. Williams. 2017. "Measuring the Natural Rate of Interest: International Trends and Determinants." *Journal of International Economics* 108: S59–S75. Thirty-nineth Annual NBER International Seminar on Macroeconomics.

Ilzetzki, E., C.M. Reinhart, and K.S. Rogoff. 2017. "The Country Chronologies to Exchange Rate Arrangements into the 21st Century: Will the Anchor Currency Hold?" Working Papers No. 23135, National Bureau of Economic Research.

Jordà, Ò., M. Schularick, and A.M. Taylor. 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking* 45(s2): 3–28.

———. 2015. "Betting the House." *Journal of International Economics* 96: S2–S18. Thirty-seventh Annual NBER International Seminar on Macroeconomics.

———. 2016. "Macrofinancial History and the New Business Cycle Facts." University of Chicago Press.

———. 2017. "The Effects of Quasi-random Monetary Experiments." Working Papers No. 23074, National Bureau of Economic Research.

Laeven, L. and F. Valencia. 2013. "Systemic Banking Crises Database." *IMF Economic Review* 61(2): 225-70.

Mavroeidis, S. 2005. "Identification Issues in Forward-looking Models Estimated by GMM, with an Application to the Phillips Curve." *Journal of Money, Credit and Banking* 37(3): 421–48.

Mavroeidis, S., M. Plagborg-Møller, and J.H. Stock. 2014. "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve." *Journal of Economic Literature* 52(1): 124–88.

McLeay, M. and S. Tenreyro. 2018. "Optimal Inflation and the Identification of the Phillips Curve." Mimeograph, London School of Economics.

Nason, J.M. and G.W. Smith. 2008. "Identifying the New Keynesian Phillips Curve." *Journal of Applied Econometrics* 23(5): 525–51.

Nechio, F. (2011). "Monetary policy when one size does not fit all." Economic Letter 2011-18, Federal Reserve Bank of San Francisco.

Obstfeld, M., J.C. Shambaugh, and A.M. Taylor. 2004. "Monetary Sovereignty, Exchange Rates, and Capital Controls: The Trilemma in the Interwar Period." IMF Staff Papers No. 75–108.

———. 2005. "The Trilemma in History: Tradeoffs among Exchange Rates, Monetary Policies, and Capital Mobility." *Review of Economics and Statistics* 87(3): 423–38.

Obstfeld, M. and A. Taylor. 1998. "The Great Depression as a Watershed: International Capital Mobility over the Long Run." In *The Defining Moment: The Great Depression and the American Economy in the Twentieth Century*. National Bureau of Economic Research.

———. 2003. "Globalization and Capital Markets." In *Globalization in Historical Perspective*. University of Chicago Press.

Orphanides, A. and J.C. Williams. 2005. "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations." *Journal of Economic Dynamics and Control* 29(11): 1927–50.

Phelps, E.S. 1967. "Phillips Curves, Expectations of Inflation, and Optimal Employment over Time." *Economica* 63: 675–94.

Phillips, A.W. 1958. "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957." *Economica* 25(100): 283–99.

Roberts, J.M. 1995. "New Keynesian Economics and the Phillips Curve." *Journal of Money, Credit and Banking* 27(4): 975–84.

Sbordone, A.M. 2002. "Prices and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49(2): 265–92.

Shambaugh, J. 2004. "The Effect of Fixed Exchange Rates on Monetary Policy." *Quarterly Journal of Economics* 119(1): 301–52.

Shimer, R. 2017. "Labor Markets." *Journal of Political Economy* 125(6): 1834–40.

Smets, F. and R. Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97(3): 586–606.

Stock, J.H. and M.W. Watson. 2007. "Why Has US Inflation Become Jarder to Forecast?" *Journal of Money, Credit and Banking* 39(s1): 3–33.

———. 2008. "Phillips Curve Inflation Forecasts." Conference Series [Proceedings] 53.

Summers, L.H. 2014. "Reflections on the 'New Secular Stagnation Hypothesis'." In *Secular Stagnation: Facts, Causes and Cures*, Chapter 1. VoxEU.org.

Swanson, E.T. and J.C. Williams. 2014. "Measuring the Effect of the Zero Lower Bound on Medium and Longer-term Interest Rates." *American Economic Review* 104(10): 3154–85.

Woodford, M. 2010. "Optimal Monetary Stabilization Policy." *Handbook of Monetary Economics* 3: 723–828.

Appendix

Data Sources and Calculations

The set of countries included in the analysis are: Australia (AUS), Austria (AUT), Belgium (BEL), Brazil (BRA), Canada (CAN), Chile (CHL), China (CHN), Colombia (COL), Costa Rica (CRI), the Czech Republic (CZE), Denmark (DEN), Estonia (EST), Finland (FIN), France (FRA), Germany (DEU), Greece (GRC), Hungary (HUN), Iceland (ISL), India (IND), Indonesia (IDN), Ireland (IRL), Israel (ISR), Italy (ITA), Japan (JPN), Korea (KOR), Latvia (LTA) Lithuania (LTU), Luxembourg (LUX), Mexico (MEX), the Netherlands (NLD), New Zealand (NZL), Norway (NOR), Poland (POL), Portugal (PRT), Russia (RUS), Saudi Arabia (SAU), Slovakia (SVK), Slovenia (SVN), South Africa (ZAF), Spain (ESP), Sweden (SWE), Switzerland (CHE), Turkey (TUR), the United Kingdom (GBR), and the United States (USA).

Inflation and Expected Inflation

We use year-over-year inflation calculated from the quarterly headline consumer price index. See table A1 for inflation sources. We use primarily IFS and OECD data, though we use Swiss National Bank data for Switzerland and Tao Zha's compiled dataset for China.

Quarterly expected inflation is often available from the OECD as year-over-year CPI growth. We use the Society of Professional Forecasters (SPF) expected year-over-year CPI inflation, which is also quarterly, for the United States. If expected inflation is not available, we estimate it as in Hamilton and others (2016):

First, we take annual year-over-year CPI inflation. Then, we estimate a rolling regression of current on lagged inflation with a twenty-year window separately for each country. With the rolling coefficients we then predict inflation by using the coefficients from the previous twenty years of data to obtain a reasonable estimate of expected inflation for the next year. Finally, we linearly interpolate the estimated annual expected inflation to quarterly.

Unemployment and NAIRU

When available, we use seasonally adjusted harmonized unemployment rates from the OECD (BLS for the U.S.); however, for a number of countries seasonally adjusted data is not available. For those, we implement an x12 seasonal adjustment on the raw series.

Country-specific NAIRU is also typically available through the OECD (CBO for the U.S.); if not we use a four-quarter moving average of seasonally adjusted unemployment rates as a country's NAIRU. Thus unemployment gap is calculated as seasonally adjusted unemployment minus NAIRU.

Interest Rates

We use the two-year treasury constant maturity rate for the U.S, which comes from the Board of Governors. We use the ten-year main long-term government bond yield for Germany, which comes from the OECD. Both series are available on Federal Reserve Economic Data (FRED).

We use the natural rate of interest from Holston and others (2017) for the U.S., U.K., Germany, and Canada to calculate the Taylor rule and its residual for each country. Data for Brazil come from the Central Bank of Brazil.

Actual and Potential GDP

Actual GDP is nearly always provided by the OECD. The exceptions are the U.S., for which we use CBO measures; and Saudi Arabia, which comes from its own Central Department of Statistics and Inflation with exchange rates from the Saudi Arabian Monetary Authority.

Potential GDP also comes from the OECD or CBO when available; otherwise we use a four-quarter moving average.

Output gap is calculated as actual minus potential GDP, divided by potential GDP.

Oil and Commodity Inflation

We use the West Texas Intermediate crude oil dollars per barrel. The index is monthly; we calculate inflation as year-over-year growth based on the quarterly average of the index.

Commodity inflation comes from the Continuous Commodity Future Price Index in Bloomberg, which is compiled by Thomson Reuters. The series is an "equal-weighted geometric average of commodity-price levels relative to the base year average price." The index is daily; we calculate inflation as year-over-year growth based on the quarterly average of the index.

Global Financial Crisis

We use Laeven and Valencia (2013), which lists the start years for systemic banking crises across countries, to define countries that experienced the global financial crisis. We list a country as having experienced the global financial crisis if it experienced a systemic banking crisis starting in 2007 or 2008. The countries in our sample that experienced the global financial crisis are: Austria, Belgium, Denmark, France, Germany, Greece, Hungary, Israel, Italy, Luxembourg, Latvia, the Netherlands, Portugal, Slovenia, Spain, Sweden, Switzerland, the United Kingdom, and the United States.

Table A1. Data Sources

Country	Inflation	Unemp. Rate	Exp. Inflation	NAIRU	GDP	Potential GDP	Interest Rate	GFC
			OECD economies					
AUS	OECD	OECD	OECD	OECD	OECD	OECD	OECD 3M	0
AUT	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
BEL	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
CAN	OECD	OECD	OECD	OECD	OECD	OECD	OECD 10Y	0
CHE	SNB	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
CZE	OECD	OECD	OECD	OECD	OECD	OECD	OECD 3M	0
DEU	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
DNK	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
ESP	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
EST	IFS	OECD	OECD	OECD	OECD	OECD	OECD 3M	0
FIN	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	0
FRA	OECD	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
GBR	IFS	OECD	OECD	OECD	OECD	OECK	OECD 10Y	1
GRC	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
HUN	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
IRL	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
ISL	IFS	OECD	OECD	OECD	OECD	OECD	OECD 3M	1
ISR	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	0
ITA	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
JPN	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	0
KOR	IFS	OECD	OECD	OECD	OECD	OECD	OECD 3M	0
LUX	IFS	OECD	OECD	OECD	OECD	OECD	GFD 10Y	1
LVA	IFS	OECD	OECD	OECD	OECD	OECD	GFD 10Y	1
MEX	IFS	OECD	OECD	OECD	OECD	4Q	OECD 3M	0
NLD	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
NOR	IFS	OECD	OECD	OECD	OECD	OECD	OECD 3M	0
NZL	OECD	OECD	OECD	OECD	OECD	OECD	OECD 3M	0
POL	IFS	OECD	OECD	OECD	OECD	OECD	OECD 3M	0
PRT	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
SVK	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	0
SVN	OECD	OECD	OECD	OECD	OECD	OECD		1
SWE	IFS	OECD	OECD	OECD	OECD	OECD	OECD 10Y	1
TUR	IFS	OECD	OECD	4Q	OECD	4Q	Bloomberg 2Y	0
USA	BLS	BLS	SPF	CBO	CBO	CBO	Bloomberg 2Y	1

Table A1. (continued)

Country	Inflation	Unemp. Rate	Exp. Inflation	NAIRU	GDP	Potential GDP	Interest Rate	GFC
			Non-OECD economies					
BRA	BCB	OECD/x12	OECD	4Q	OECD	4Q	Bloomberg 2Y	0
CHL	OECD	OECD	OECD	OECD	OECD	4Q	OECD 3M	0
CHN	ZHA	IFS/x12	OECD	4Q	OECD	4Q	OECD 3M	0
COL	IFS	IFS/x12	OECD	4Q	OECD	4Q	Bloomberg 2Y	0
CRI	IFS	IFS/x12	Rolling Reg.	4Q	OECD	4Q	OECD 3M	0
IDN	IFS	IFS/x12	OECD	4Q	OECD	4Q	OECD 3M	0
IND	IFS	Eurostat/ x12	OECD	4Q	OECD	4Q	GFD 10Y	0
LTU	OECD	IFS/x12	Rolling Reg.	4Q	OECD	4Q	GFD 10Y	0
RUS	IFS	IFS/x12	OECD	4Q	OECD	4Q	OECD 3M	0
SAU	IFS	IFS/x12	Rolling Reg.	4Q	CDSI	4Q		0
ZAF	IFS	IFS/x12	OECD	4Q	OECD	4Q	OECD 3M	0

Notes: Lithuania the OECD in 2018.

Estimates of the Phillips Curve Using the Unemployment Gap (Equation 6)

Table A2. Unemployment Gap: Full Sample (1986Q1-2018Q1)

	ALL		*OECD*		*Non-OECD*	
	OLS	*IV*	*OLS*	*IV*	*OLS*	*IV*
Persistence (ρ_1)	0.79***	0.66***	0.74***	0.90***	0.79***	0.41**
	(0.03)	(0.04)	(0.03)	(0.03)	(0.03)	(0.19)
Inflation expectations (ρ_e)	0.17***	0.35***	0.20***	0.03	0.17***	0.79**
	(0.04)	(0.06)	(0.03)	(0.03)	(0.04)	(0.32)
Slack (β)	-0.03***	-0.02	-0.03**	-0.02*	-0.03*	0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.05)
Observations	3,181	2,977	1,863	1,795	1,318	1,182
Adjusted R^2		0.91		0.92		0.73
K–P Stat		4.99		12.63		2.99

Source: Authors' calculations.

Table A3. Unemployment Gap: Before Crisis (1986Q1-2007Q4)

	ALL		*OECD*		*Non-OECD*	
	OLS	*IV*	*OLS*	*IV*	*OLS*	*IV*
Persistence (ρ_1)	0.77***	0.63***	0.77***	0.99***	0.78***	0.67***
	(0.04)	(0.06)	(0.03)	(0.06)	(0.05)	(0.07)
Inflation expectations (ρ_e)	0.19***	0.39***	0.17***	-0.06	0.20***	0.***
	(0.04)	(0.06)	(0.04)	(0.06)	(0.05)	(0.10)
Slack (β)	-0.04***	-0.04*	-0.05***	-0.03**	-0.02	-0.01
	(0.01)	(0.02)	(0.02)	(0.01)	(0.03)	(0.05)
Observations	2,006	1,811	1,289	1,221	717	590
Adjusted R^2		0.89		0.90		0.88
K–P Stat		4.56		14.01		2.99

Source: Authors' calculations.

Table A4. Unemployment Gap: After Crisis
(2008Q1-2018Q1)

	ALL		*OECD*		*Non-OECD*	
	OLS	*IV*	*OLS*	*IV*	*OLS*	*IV*
Persistence (ρ_1)	0.73***	0.45***	0.63***	0.47***	0.74***	0.51***
	(0.06)	(0.10)	(0.03)	(0.10)	(0.07)	(0.11)
Inflation expectations (ρ_e)	0.30***	0.97***	0.42***	0.83***	0.28**	0.82***
	(0.11)	(0.17)	(0.05)	(0.30)	(0.12)	(0.17)
Slack (β)	-0.01	0.03	0.02**	0.03***	-0.03	0.01
	(0.02)	(0.03)	(0.01)	(0.01)	(0.03)	(0.04)
Observations	1,175	1,166	574	574	601	592
Adjusted R^2		0.78		0.87		0.82
K–P Stat		7.10		9.15		4.07

Source: Authors' calculations.

Difference-in-differences Results Using the Unemployment Gap

Table A5. Difference in Differences Estimation. Phillips Curve Using Unemployment Gap

(All countries)

	Persistence (ρ_1)			*Inflation expectations* (ρ_e)			*Slack* (β)		
	All-OLS	*Peg-OLS*	*Peg-IV*	*All-OLS*	*Peg-OLS*	*Peg-IV*	*All-OLS*	*Peg-OLS*	*Peg-IV*
	(a) No-Crisis economies, G_i =0								
Before	0.76***	0.80***	1.04***	0.24***	0.19*	-0.10	-0.01	-0.02	-0.08
	(0.04)	(0.08)	(0.18)	(0.06)	(0.11)	(0.28)	(0.01)	(0.02)	(0.05)
After	0.68***	0.72***	0.62***	0.38***	0.30**	0.45***	-0.06***	-0.06***	-0.07***
	(0.06)	(0.08)	(0.04)	(0.09)	(0.13)	(0.13)	(0.02)	(0.02)	(0.02)
(i) Diff.	-0.08*	-0.08***	-0.42**	0.14**	0.10***	0.54**	-0.05*	-0.04	0.01
	(0.04)	(0.03)	(0.20)	(0.05)	(0.04)	(0.22)	(0.03)	(0.03)	(0.05)
	(b) Crisis economies, G_i =1								
Before	0.66***	0.63***	0.71***	0.36***	0.39***	0.29**	-0.01	0.00	-0.02
	(0.04)	(0.04)	(0.10)	(0.04)	(0.04)	(0.13)	(0.01)	(0.01)	(0.02)
After	0.60***	0.61***	0.59***	0.48***	0.48***	0.51***	-0.01	-0.00	-0.00
	(0.02)	(0.02)	(0.05)	(0.04)	(0.05)	(0.12)	(0.01)	(0.01)	(0.02)
(ii) Diff.	-0.05	-0.02	-0.13**	0.12**	0.09*	0.21***	-0.00	-0.00	0.02
	(0.04)	(0.03)	(0.06)	(0.05)	(0.05)	(0.06)	(0.02)	(0.02)	(0.02)
	(c) Treatment effect: (ii) – (i)								
D-i-D:	0.03	0.06	0.30	-0.01	-0.01	-0.33	0.05	0.03	0.01
	(0.06)	(0.04)	(0.23)	(0.07)	(0.06)	(0.23)	(0.03)	(0.03)	(0.06)
Observations	3,376	2,162	2,134						
Adjusted R^2	0.88	0.90	0.87						

Source: Authors' calculations.

Notes: Sample includes all countries. Clustered robust standard errors. */**/*** indicates significance at the 90/95/99% confidence level. See text.

Table A6. Difference in Differences Estimation. Phillips Curve Using Output Gap
(OECD countries excluding Hungary, Greece, Latvia, and Lithuania)

	Persistence (ρ_1)			*Inflation expectations* (ρ_e)			*Slack* (β)		
	All-OLS	*Peg-OLS*	*Peg-IV*	*All-OLS*	*Peg-OLS*	*Peg-IV*	*All-OLS*	*Peg-OLS*	*Peg-IV*
	(a) No-Crisis economies, G_i =0								
Before	0.54***	0.50***	0.71***	0.50***	0.58***	0.29*	-0.01	-0.01	-0.09
	(0.05)	(0.06)	(0.13)	(0.08)	(0.12)	(0.15)	(0.05)	(0.09)	(0.09)
After	0.58***	0.60***	0.59***	0.51***	0.51***	0.46***	-0.03	0.06	-0.00
	(0.04)	(0.08)	(0.06)	(0.07)	(0.15)	(0.10)	(0.04)	(0.04)	(0.03)
(i) Diff.	0.04	0.09***	-0.12	0.02	-0.07**	0.16*	-0.02	0.07	0.08
	(0.06)	(0.02)	(0.09)	(0.06)	(0.03)	(0.08)	(0.04)	(0.05)	(0.09)
	(b) Crisis economies, G_i =1								
Before	0.61***	0.68***	0.77***	0.30***	0.33***	0.20*	-0.01	-0.00	0.02
	(0.03)	(0.05)	(0.08)	(0.03)	(0.05)	(0.11)	(0.01)	(0.01)	(0.02)
After	0.61***	0.62***	0.64***	0.46***	0.46***	0.37***	-0.01	0.00	-0.01
	(0.03)	(0.02)	(0.04)	(0.04)	(0.03)	(0.12)	(0.01)	(0.01)	(0.01)
(ii) Diff.	-0.10**	-0.06	-0.12*	0.16***	0.13***	0.18**	-0.01	0.01	0.02
	(0.05)	(0.04)	(0.07)	(0.05)	(0.04)	(0.08)	(0.01)	(0.01)	(0.02)
	(c) Treatment effect: (ii) – (i)								
D-i-D:	-0.14*	-0.15***	-0.01	0.14*	0.20***	0.01	0.01	-0.06	-0.07
	(0.07)	(0.04)	(0.09)	(0.08)	(0.05)	(0.08)	(0.05)	(0.05)	(0.09)
Observations	1,760	1,146	1,146						
Adjusted R^2	0.88	0.89	0.89						

Source: Authors' calculations.
Notes: Sample includes OECD countries except Hungary, Greece, Latvia, and Lithuania. Clustered robust standard errors. */**/*** indicates significance at the 90/95/99% confidence level. See text.

Table A7. Difference in Difference Estimation. Phillips Curve Using Unemployment Gap. Non-OECD countries
(excluding Hungary, Greece, Latvia, and Lithuania)

	Persistence (ρ_1)			*Inflation expectations* (ρ_e)			*Slack* (β)		
	All-OLS	*Peg-OLS*	*Peg-IV*	*All-OLS*	*Peg-OLS*	*Peg-IV*	*All-OLS*	*Peg-OLS*	*Peg-IV*
	(a) No-Crisis economies, $G_i = 0$								
Before	0.76***	0.81***	1.04***	0.24***	0.19*	-0.10	0.01	-0.02	-0.07
	(0.05)	(0.08)	(0.18)	(0.06)	(0.11)	(0.27)	(0.02)	(0.02)	(0.05)
After	0.69***	0.73***	0.67***	0.37***	0.28**	0.39***	-0.06***	0.06***	0.07***
	(0.06)	(0.09)	(0.05)	(0.10)	(0.13)	(0.11)	(0.02)	(0.02)	(0.02)
(i) Diff.	-0.07	-0.07**	-0.37*	0.13**	0.10**	0.49**	-0.06*	-0.04	0.00
	(0.04)	(0.03)	(0.21)	(0.06)	(0.04)	(0.24)	(0.03)	(0.03)	(0.06)
	(b) Crisis economies, $G_i = 1$								
Before	0.58***	0.59***	0.75***	0.46***	0.43***	0.24***	0.02	0.02**	-0.04
	(0.04)	(0.04)	(0.05)	(0.03)	(0.04)	(0.05)	(0.01)	(0.01)	(0.03)
After	0.57***	0.58***	0.58***	0.56***	0.53***	0.49***	0.00	0.00	-0.01
	(0.02)	(0.02)	(0.02)	(0.05)	(0.06)	(0.06)	(0.02)	(0.03)	(0.03)
(ii) Diff.	-0.01	-0.01	-0.17***	0.10	0.09	0.25***	-0.02	-0.02	0.03
	(0.04)	(0.03)	(0.04)	(0.08)	(0.07)	(0.07)	(0.03)	(0.03)	(0.02)
	(c) Treatment effect: (ii) – (i)								
D-i-D:	0.06	0.06	0.21	-0.03	-0.00	-0.24	0.04	0.02	0.04
	(0.06)	(0.05)	(0.22)	(0.09)	(0.08)	(0.24)	(0.04)	(0.04)	(0.06)
Observations	1,616	1,016	988						
Adjusted R^2	0.88	0.91	0.87						

Source: Authors' calculations.
Notes: Sample excludes OECD countries. Clustered robust standard errors. */**/*** indicates significance at the 90/95/99% confidence level. See text.

Trend, Seasonal, and Sectorial Inflation in the Euro Area

James H. Stock
Harvard University

Mark W. Watson
Princeton University

A central focus of monetary policy is the underlying rate of inflation that might be expected to prevail over a horizon of one or two years. Because inflation is estimated from noisy data, the estimation of this underlying rate of inflation, which we refer to as trend inflation, requires statistical methods to extract the inflation "signal" from the noise. The task of measuring trend inflation is further complicated by the large seasonal fluctuations in many prices, so that attempts to estimate core or trend inflation at a frequency higher than annual must additionally either use seasonally adjusted data or undertake seasonal adjustment as part of the effort to measure trend inflation.

The challenge of estimating trend inflation is particularly acute for the euro-area Harmonized Index of Consumer Prices (HICP) inflation, official values of which are only reported seasonally unadjusted. Figure 1 plots quarterly values of euro-area HICP inflation (in percentage points at an annual rate) from 2001 to 2018. The quarter-to-quarter variation in inflation is large: the standard deviation of quarterly changes in inflation is 2.5 percentage points. HICP inflation is also highly seasonal: over the entire sample period, inflation averaged 1.6 percent, but averaged 4.8 and 2.2 percent in the second and fourth quarters respectively, and 0.1 and 0.3 over the first and third quarters. While some long-run, low-frequency variation in HICP inflation is evident, that variation —the "signal"— is small compared

Changing Inflation Dynamics, Evolving Monetary Policy edited by Gonzalo Castex, Jordi Galí, and Diego Saravia, Santiago, Chile.

to the seasonal variation and what appears to be transient, one-off movements in the rate of inflation. The question, "What is the value of trend inflation today?" is an important one for monetary policy, but the answer to it arguably requires more than just staring at figure 1.

One approach to estimating trend inflation is to exploit variation across the components of inflation (across sectors) to reduce noise. The most prominent such estimates are "core" measures (e.g., Gordon, 1975 Eckstein, 1981) that exclude inflation from the volatile food and energy sectors. Alternative core measures include trimmed mean or median of sectorial inflation rates; for example, see the early work by Bryan and Cecchetti (1994) or the paper by Ball and Mazumder in this volume. Ehrmann and others (2018) provide an up-to-date summary of work at the European Central Bank (ECB) involving underlying and sectorial inflation.

The HICP has 12 second-tier components, which we modify to create 13 components by pooling the energy components of housing and transportation into a separate "energy" component. These 13 inflation components are plotted in figure 2. The heterogeneity of the time-series properties of these components is striking. Some sectors exhibit large seasonal variation (for example, clothing), others exhibit large non-seasonal quarterly variation (energy) or outliers (healthcare), and relative price movements impart different lower-frequency trends in each sector. Almost as striking is the apparent variation over time in those time-series properties, for example, the seasonal components of furnishing, clothing, and transportation have increased markedly over this period. The heterogeneity of these components suggests that there could be considerable gains from using a multivariate approach that allows the components to have distinct time-series properties and uses both time-series smoothing and cross-sectional weighting to estimate aggregate HICP trend inflation.

This paper makes three contributions towards measuring trend HICP inflation. First, we estimate an unobserved components (UC) model with stochastic volatility (UCSV), which extends the UCSV model in Stock and Watson (2007) to include a seasonal component. This univariate model is an extension of the textbook unobserved components model[1] to incorporate stochastic volatility to capture

1. Chapter 1 of Nerlove and others (1979) offers a historical survey of UC models in economics. The textbook by Harvey (1989) is a classic reference on analyzing UC models by using Kalman filter methods.

the time-varying importance of the trend, seasonal, and irregular components.[2]

Second, we extend the multivariate unobserved components/stochastic volatility model of Stock and Watson (2016) to allow each component to have separate seasonals, also with stochastic volatility. We apply this extended model to the 13 HICP components in figure 2 to obtain multivariate estimates of the trend. We find that doing so produces trend estimates that are more precise than those based on the univariate model of aggregate HICP. We also find that this measure of core inflation moves cyclically with real economic activity.

Third, as a byproduct, we also obtain quarterly estimates of seasonally adjusted HICP. Another approach to handling seasonals is simply to use the four-quarter average of quarterly inflation; however, that measure tends to respond sluggishly. Compared with four-quarter rolling inflation, the new seasonally adjusted HICP series has the potential to provide more timely insights into movements of inflation.

Section 1 presents the univariate and multivariate model that we use for aggregate and sectorial inflation. Section 2 uses these models to estimate trend and seasonal factors for euro-area HICP inflation. Section 3 examines the relation between seasonally adjusted inflation and real activity.

Figure 1. HICP Inflation for the Euro Area

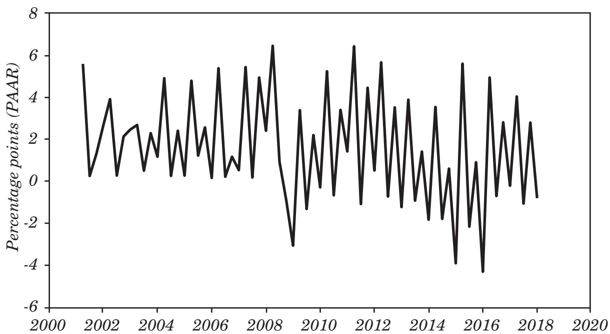

Source: Authors' calculations.

2. Several papers have used related univariate UC models to study the evolution of prices and inflation. Examples include Ball and Cecchetti (1990); Ceccheti and others (2007); Cogley and Sargent (2015); Cogley and others (2015); and Kang and others (2009).

Figure 2. 13 HICP Sectors

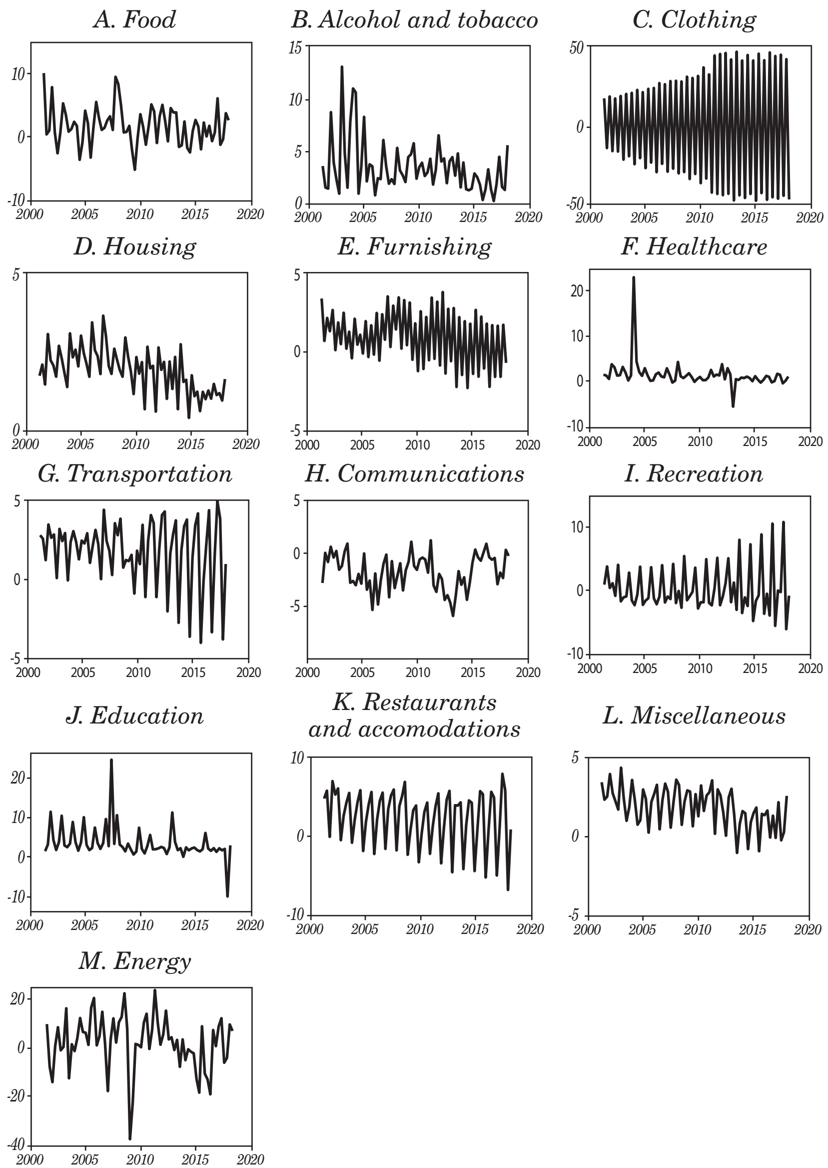

Source: Authors' calculations.
Notes: These are the 12 HICP tier-two sectors, with energy excluded from the housing and transportation sectors, and shown separately as the 13th sector.

1. Seasonal UCSV Models

Unobserved components models have a long history in economic time series and have been used for, among other things, data description, forecasting, structural analysis, and seasonal adjustment. Here we present versions of the UC model that can be used to seasonally adjust aggregate inflation and to estimate its trend value. One version of the model is univariate and uses only aggregate inflation; the other is multivariate and models the joint dynamics of sectoral inflation. Both models incorporate stochastic volatility and are known by their acronym UCSV.

1.1 Univariate Seasonal UCSV Model

Inflation is observed quarterly and is denoted by π_t. The UC model decomposes π_t into three unobserved components: trend (τ_t), seasonal (s_t), and irregular (ε_t).

$$\pi_t = \tau_t + s_t + \varepsilon_t. \tag{1}$$

The components are separately identified because they follow distinct stochastic processes. Let $\eta_{\tau,t}$, $\eta_{s,t}$, and $\eta_{\varepsilon,t}$ denote three martingale-difference processes; the trend component follows a martingale:

$$(1 - L)\tau_t = \eta_{\tau,t} \tag{2}$$

so it is dominated by low-frequency, or "trend", variation; s_t follows the quarterly seasonal process:

$$(1 + L + L^2 + L^3)\, s_t = \eta_{s,t} \tag{3}$$

so is dominated by variation at the seasonal frequencies with periods 2 and 4 quarters; and the irregular component is unforecastable:

$$\varepsilon_t = \eta_{\varepsilon,t} \tag{4}$$

The unobserved components model (1)–(4) is a version of Harvey's (1989) "local-level" model, augmented by the seasonal component s_t. Versions of the model (often with more flexible models for the components) are the backbone of model-based seasonal adjustment methods—e.g., Hillmer and Tiao (1982), Hausman and Watson (1985), and Maravall (1995).

In the non-seasonal version of the local-level model, the estimate of τ_t based on observations of π through date t is the forecast of the future rate of inflation:

$$E(\pi_{t+h}|\{\pi_i\}_{i=1}^{t}) = E(\tau_{t+h} + \varepsilon_{t+h}|\{\pi_i\}_{i=1}^{t}) = E(\tau_t|\{\pi_i\}_{i=1}^{t}) = \tau_{t|t}, \tag{5}$$

where the final equality follows from the martingale assumption for τ_t and the martingale-difference assumption for ε_t.

The seasonal model (3) is specified so that this definition of the trend as the long-run forecast continues to hold for annual averages. Specifically, Harvey (1989), in subsection 6.2, defines a seasonal process to be any time-series process with predicted values that (i) repeat seasonally and (ii) sum to zero over a one-year period. The seasonal process (3) satisfies these two conditions, specifically (i) $s_{T'+j|T} = s_{T'+j+4|T}$ and (ii) $\Sigma_{j=1}^{4} s_{T'+j|T} = 0$, where $s_{r|T}$ is the predicted value of s_r made by using data through time T, for any $T' \geq T$. The seasonal model (3) yields a similar interpretation of $\tau_{t|t}$, but now for annual averages of future values of π: letting $\bar{x}_{i:j}$ denote the sample average of an arbitrary variable x between time i and j,

$$E\left(\bar{\pi}_{t+j:t+j+3}|\{\pi_k\}_{k=1}^{t}\right) = E\left(\bar{\tau}_{t+j:t+j+3} + \bar{s}_{t+j:t+j+3} + \bar{\varepsilon}_{t+j:t+j+3}|\{\pi_k\}_{k=1}^{t}\right) \tag{6}$$
$$= E\left(\tau_t|\{\pi_k\}_{k=1}^{t}\right) = \tau_{t|t}$$

for $j > 0$, where the penultimate equality follows from the random walk model for τ, $\Sigma_{j=1}^{4} s_{T'+j|T} = 0$, and the unpredictability of future ε's. Thus, as in the model without seasonality, $\tau_{t|t}$ measures the (non-seasonal) forecastable level of inflation.

Examination of the inflation series in figure 1 and figure 2 highlights the need for two modifications of the basic UCSV model. The first modification allows for time variation in the variances of the unobserved components, and the second allows for outliers. We discuss these in turn.

Time-varying variances are added to the model by allowing the shocks in (2), (3), and (4) to follow stochastic volatility processes, say $\eta_t = \sigma_t e_t$, where $e_t \sim$ i.i.d N(0,1) and σ_t^2 evolves through time as a logarithmic random walk: $(1-L)\ln(\sigma_t^2) = \nu_t$ with $\nu_t \sim$ i.i.d $N(0, \sigma_\nu^2)$. Kim and others (1998) show how this stochastic volatility model can be estimated using Gibbs sampling methods by using a mixture of normal densities to approximate the log-χ_t^2 density together with standard Kalman smoothing recursions; Omori and others (2007) provide improved approximations. Stock and Watson (2007) incorporate these methods together with ideas in Carter and Kohn (1994), and Kim and Nelson (1999) to estimate a non-seasonal version of the UCSV model.

Outliers are incorporated in the model through additional random multiplicative factors linking the η_t innovations to the i.i.d. N(0,1) shocks e_t. As in Stock and Watson (2016), we use a formulation with $\eta_t = o_t\sigma_t e_t$ where o_t is an i.i.d. outlier term with $o_t = 1$ with probability 1–p and $o_t \sim$ U(2,10) with probability p. When $o_t = 1$, there is no outlier, and when $o_t \sim$ U(2,10) there is an outlier with a standard deviation that is between 2 and 10 times larger than in the no-outlier case. In the model for euro-area inflation, we allow outliers only in the irregular component ε_t, as this seems consistent with outliers evident in figure 2; in other applications, outliers might also be appropriate for τ_t and/or s_t.

In summary, the complete UCSV model is (1)–(4) and

$$\eta_{\tau,t} = \sigma_{\tau,t}e_{\tau,t} \;;\eta_{s,t} = \sigma_{s,t}e_{s,t} \;;\eta_{\varepsilon,t} = o_t\sigma_{\varepsilon,t}e_{\varepsilon,t} \tag{7}$$

$$(1-\mathrm{L})\ln(\sigma_{x,t}) = \nu_{x,t} \quad \text{for } x = \tau, s, \varepsilon, \tag{8}$$

where $(e_{\tau,t}, e_{s,t}, e_{\varepsilon,t}, \nu_{\tau,t}, \nu_{s,t}, \nu_{x,t})$ are mutually independent i.i.d. normal random variables with mean zero, the e terms have unit variance, and each of the ν terms has a component-specific variance, say $\sigma_{\nu(\tau)}$, $\sigma_{\nu(s)}$, and $\sigma_{\nu(\varepsilon)}$.

1.2 Multivariate Seasonal UCSV Model

The multivariate model is a generalization of the univariate that includes common and sector-specific versions of the three unobserved components. For each of the $i = 1,\ldots, n$ sectors, the rate of price inflation in sector i, $\pi_{i,t}$ follows:

$$\pi_{i,t} = \alpha_{i,\tau}\tau_{c,t} + \alpha_{i,s}s_{c,t} + \alpha_{i,\varepsilon}\varepsilon_{c,t} + \tau_{i,t} + s_{i,t} + \varepsilon_{i,t} \tag{9}$$

where $(\tau_{c,t}, s_{c,t}, \varepsilon_{c,t})$ are common to all sectors, $(\tau_{i,t}, s_{i,t}, \varepsilon_{i,t})$ are sector specific, and $(\alpha_{i,\tau}, \alpha_{i,s}, \alpha_{i,\varepsilon})$ are time-invariant coefficients (factor loadings). The τ, s, ε components follow processes as in the univariate model, with component/sector-specific parameters. The components are mutually independent, so that dependence across sectors comes from the common components τ_c, s_c, and ε_c. Outliers are allowed in each of the sector-specific $\varepsilon_{i,t}$ components and in the common $\varepsilon_{c,t}$ component.

The multivariate sectorial model is designed so that it (approximately) aggregates to univariate UCSV model. Because of its symmetric structure, aggregation in the multivariate model is straightforward: letting $w_{i,t}$ denote the share weight for sector i at time t

$$\pi_t = \sum_{i=1}^{n} w_{i,t}\pi_{i,t} = \tau_t^a + s_t^a + \varepsilon_t^a \tag{10}$$

where

$$\tau_t^a = {}_{c,t}\sum w_{i,t}\alpha_{i,\tau} + \sum w_{i,t}\tau_{i,t} \tag{11}$$

and similarly for the other components. When the share weights are time-invariant, τ_t^a evolves as a martingale, s_t^a follows the seasonal process in (3), and ε_t^a is a martingale difference. And, as in the univariate model, filtered values of (τ_t^a, $\tau_{c,t}$, $\tau_{i,t}$) constructed from the multivariate model summarize the forecastable levels in both sectorial and aggregate inflation:

$$E\left(\bar{\pi}_{i,t+j:i,t+j+3} | \left\{\pi_{l,k}\right\}_{k=1}^{t}, l = 1,\ldots,n\right) = \alpha_{i,\tau}\tau_{c,t|t} + \tau_{i,t|t} \tag{12}$$

and

$$E\left(\bar{\pi}_{t+j:t+j+3} | \left\{\pi_{l,k}\right\}_{k=1}^{t}, l = 1,\ldots,n\right) = \tau_{t|t}^a . \tag{13}$$

1.3 Estimation and Inference

We estimate the univariate and multivariate UCSV models by using Bayes' methods that are generalizations of the methods outlined in online appendix to Stock and Watson (2016). We provide an overview here.

The univariate UCSV model is characterized by four sets of parameters: (i) the stochastic volatility innovation standard deviations,

$\sigma_{\nu(\tau)}$, $\sigma_{\nu(s)}$, and $\sigma_{\nu(\varepsilon)}$; (2) the outlier probability parameter p; (3) the initial values for the standard deviations $\sigma_{\tau,0}$, $\sigma_{s,0}$, and $\sigma_{\varepsilon,0}$; and (4) the initial values of the components τ_0 and $(s_0, s_{-1}, s_{-2}, s_{-3})$. We used independent priors for the parameters:

- $\sigma_\nu \sim U(0,0.10)$. (A value of $\sigma_{\nu(\tau)} = 0.10$ implies that the standard deviation of $\ln(\sigma_{\tau+t+40}/\sigma_{\tau,t})$ is approximately 0.3, that is a standard deviation of 30 percent over 40 quarters).
- $p \sim \text{Beta}(a,b)$ with $a = 2.5$ and $b = 37.5$. (This implies that an outlier is expected to occur every four years).
- $\ln(\sigma_{t,0})$, $\ln(\sigma_{s,0})$, $\ln(\sigma_{\varepsilon,0})$, and τ_0 follow independent diffuse Gaussian priors.
- $(s_0, s_{-1}, s_{-2}, s_{-3})$ follow a diffuse singular Gaussian distribution, where the singularity enforces $s_0 + s_{-1} + s_{-2} + s_{-3} = 0$.

The multivariate model requires two normalizations. First, the factor structure requires a normalization to separately identify the scales of the factor loadings $(\alpha_\tau, \alpha_s, \alpha_\varepsilon)$ and the common factors (τ_c, s_c, and ε_c). We normalize the standard deviations of the common factors to be unity for $t = 0$. The second normalization is needed because the initial values of the common and idiosyncratic factors (e.g., $\tau_{c,0}$ and $\tau_{i,0}$) are not separately identified. To identify the model, we normalize the common factors to be zero for $t = 0$; that is $\tau_{c,0} = 0$ and $(s_{c,0}, s_{c,-1}, s_{c,-2}, s_{c,-3}) = 0$.

The multivariate model also requires a prior distribution for the factor loadings. Let α_τ denote the $n\times1$ vector of factor loadings for $\tau_{c,t}$; we use the prior $\alpha_\tau \sim N(0, 10^2\iota\iota' + 0.4^2 I_n)$, where ι is an $n\times1$ vectors of ones. This prior is essentially uninformative about the average value of $\alpha_{i,\tau}$ (the first term in the variance), but shrinks the factor loadings toward a common value (the second term in the factor variance). Independent priors of the same form were used for α_s and α_ε.

The empirical results in the next section are based on 60,000 Markov chain Monte Carlo (MCMC) draws from the posterior (discarding the first 10,000 draws) by using the algorithm outlined in Stock and Watson (2016), modified to incorporate the seasonal factor. Error bands are from 68-percent equal-tail credible sets. The 95-percent error bands, which are unreported, are approximately twice as wide as the reported 68-percent bands.

2. The Data and Estimation Results

2.1 Data

There are twelve tier-two components for the euro-area HICP. These consumer spending components are organized by purpose (transportation, housing, recreation, etc.) rather than by type of product (motor vehicles, gasoline, recreational goods, etc.), which is the organizing principle used in the U.S. PCE and CPI data. Because the euro-area sectors are organized by purpose, they contain a mix of both goods and services. For example, the transportation component contains both motor vehicles (a good) and airline transport (a service). Energy is not a separate sector in the HICP tier-two categorization. Because energy prices historically behave differently from other prices, including large outliers and different seasonal patterns, we extracted the major energy components from housing (electricity, gas, liquid fuels, solid fuels, heat energy) and transportation (fuels and lubricants for personal transportation equipment) to form a separate energy component. Thus the 13 components we analyze are energy, housing excluding energy, transportation excluding fuels and lubricants for personal transportation, and the ten remaining unaffected components of the HICP. These are the thirteen sectors shown in figure 2.

The data are available monthly. We temporally aggregated the monthly price indices to quarterly averages and computed sectoral inflation rates as $\pi_{i,t} = 400\times\ln(p_{i,t} / p_{i,t-1})$, where $p_{i,t}$ is the quarterly price index for sector i in quarter t. Data are available for all sectors as from 2001:Q1, and the first quarterly inflation value is for 2001:Q2. Our sample ends in 2018:Q1.

Spending shares for each sector are available annually. We interpolated the annual average shares to construct quarterly shares by using a random walk interpolator.[3] Table 1 lists the 13 sectors, shows the average share weights over the entire sample period and over the first and second subsamples. Shares vary little over the sample period; the largest sector is food (16%) and smallest is education (1%); the energy share is 10 percent.

3. That is, we modeled the unobserved quarterly shares as a random walk, the observed annual shares as the annual average of the quarterly shares, and estimated the quarterly shares by using the Kalman smoother.

Table 1. The 12 Tier-two Sectors of the Euro-Area HICP Plus the Energy Sector

	Average expenditures shares		
	2001-2018	*2001-2009*	*2010-2018*
Food	0.16	0.16	0.15
Alcohol and tobacco	0.04	0.04	0.04
Clothing	0.07	0.07	0.06
Housing (excl. energy)	0.10	0.10	0.10
Furnishing	0.07	0.08	0.07
Healthcare	0.04	0.04	0.04
Transportation (excl. energy)	0.11	0.11	0.11
Communications	0.03	0.03	0.03
Recreation	0.10	0.10	0.09
Education	0.01	0.01	0.01
Restaurants and accommodations	0.09	0.09	0.09
Miscellaneous	0.08	0.08	0.09
Energy	0.10	0.09	0.10

Source: Authors' calculations.
Notes: Energy components of housing (electricity, gas, liquid fuels, solid fuels, heat energy) and transportation (fuels and lubricants for personal transportation equipment) were removed from those components and collected into the separate "Energy" category, given in the final row.

2.2 Results

Univariate HICP. The univariate model produces estimates of the volatilities $\sigma_{\tau,t}$, $\sigma_{s,t}$, $\sigma_{\varepsilon,t}$ and the components τ_t, s_t and ε_t. Table 2 shows the estimated values (posterior medians) and 68-percent credible sets for these variables at the beginning, middle, and end of the sample.

The estimated standard deviations of the innovations in τ, s, and ε are relatively constant over the sample period. The level of trend inflation is estimated to have fallen from 2.5 percent in 2001 to 1.5 percent in 2018. The estimated seasonal component shows that aggregate HICP inflation tends to be low in the first and third quarters and high in the second; the seasonal amplitude increased over the sample period.

Table 2. Parameter Estimates for the Univariate UCSV Model for Aggregate Inflation
Posterior medians and 68-percent equal-tail posterior credible intervals

(a) Estimated volatilities and trends from the univariate model

	2001:Q2	*2009:Q4*	*2018:Q1*
	Standard deviations of shocks to components		
σ_τ	0.44 (0.25, 0.70)	0.55 (0.34, 0.85)	0.52 (0.32, 0.81)
σ_s	0.29 (0.18, 0.45)	0.26 (0.17, 0.40)	0.26 (0.15, 0.43)
σ_ε	0.61 (0.32, 0.90)	0.67 (0.36, 0.99)	0.65 (0.35, 0.98)
	Estimates of trend component		
τ_t	2.54 (2.01, 3.12)	1.39 (0.96, 1.84)	1.45 (0.94, 2.03)

(b) Estimates of seasonal factors

	Q1	*Q2*	*Q3*	*Q4*
2002	-0.06 (-0.67, 0.48)	2.16 (1.74, 2.62)	-1.76 (-2.12, -1.39)	-0.14 (-0.50, 0.22)
2009	-1.78 (-2.20, -1.34)	3.24 (2.85, 3.63)	-2.23 (-2.59, -1.86)	0.81 (0.43, 1.19)
2017	-2.42 (-2.91, -1.97)	3.45 (2.95, 3.95)	-2.22 (-2.64, -1.78)	1.19 (0.77, 1.64)

Source: Authors' calculations.

Figure 3 shows estimated values of τ_t and seasonally adjusted inflation, $\pi_t - s_t$. The upper panels show the posterior estimates based on the full sample (the smoothed estimates) and the lower panel shows estimates based on data through date t (the filtered estimates).[4] As desired, the estimates of seasonally adjusted inflation evidently eliminate the largest seasonal swings. The 68-percent error bands for seasonally adjusted inflation are wide (1.0 percentage points at the end of the sample). The time path of trend inflation is also uncertain, but, as shown below, the estimates closely track real activity in the euro area.

4. For computational simplicity, the filtered estimates are based on the full-sample estimates of the variance parameters, and are therefore approximations the true one-sided estimates. The filtered estimates are plotted beginning in 2004 because of the diffuse prior for the $t = 0$ values.

Figure 3. Smoothed and Filtered Estimates from Univariate UCSV Model for Trend (τ_t) and Seasonally Adjusted ($\pi_t - s_t$) HICP Inflation

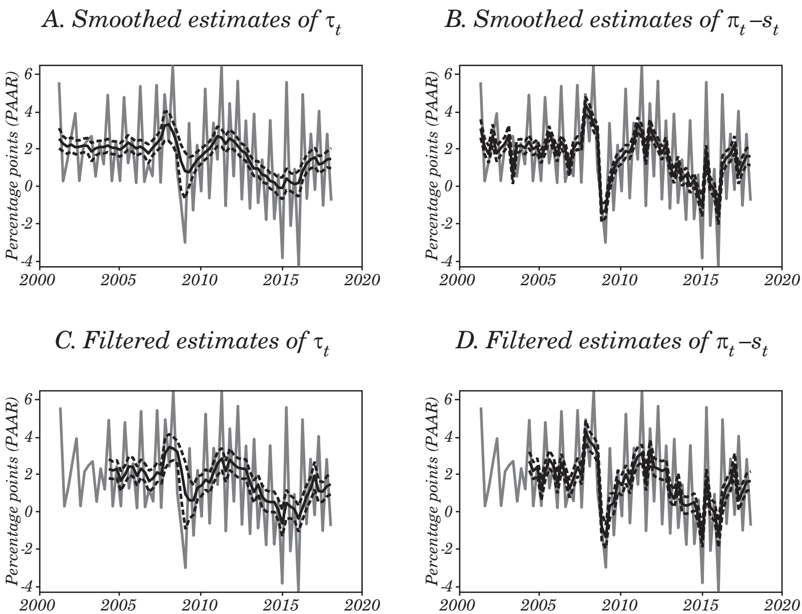

Source: Authors' calculations.
Notes: The values shown are the posterior median and 68-percent equal-tail posterior credible intervals for the dates shown.

The estimates of τ_t and s_t are weighted averages of the π_{t+j}. For example, the full-sample posterior estimates of τ_t are given by $\tau_{t|T} = \sum_{j=-t+1}^{T-t} a_{t,j}\pi_{t+j}$, where the weights $a_{t,j}$ depend on the parameters $\left\{\sigma_{\tau,i}, \sigma_{s,i}, \sigma_{\varepsilon,i}\right\}_{i=1}^{T}$. When these parameters are time-invariant and t is not close to the beginning or end of the sample, the weights are time-invariant, that is, $a_{t,j} \approx a_j$. Figure 4 plots these weights constructed by using the sample average of $\left\{\sigma_{\tau,i}, \sigma_{s,i}, \sigma_{\varepsilon,i}\right\}_{i=1}^{T}$ for both the one-sided (filtered) and two-sided (smoothed) estimates of τ_t. By construction, these weights sum to unity (because the zero-frequency pseudo-spectrum of π is determined solely by variation in τ) and the figure indicates that nearly all of the weight is placed on values of $\pi_{t,j}$ for $|j| \le 4$. These short moving-average weights are optimal because of the relatively high signal-to-noise ratio for the trend ($\sigma_\tau / \sigma_\varepsilon \approx 0.80$).

Figure 4. Weight Placed on π_{t+j} for Estimating τ_t.

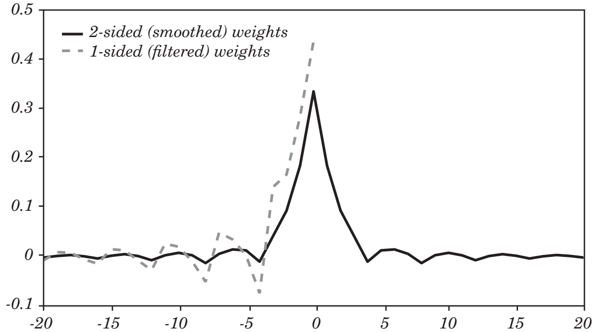

Source: Authors' calculations.
Notes: The weights are computed from the Kalman filter and smoother for a univariate trend + seasonal + irregular model with constant variances computed as the average values of the UCSV model variances.

Multivariate. The univariate model implicitly applies the same time-series filter to each of the 13 sectors making up the aggregate, with the component-wise results aggregated by using share weights. Yet it is clear from figure 2 that the components follow highly heterogeneous time-series processes. For example, the clothing sector appears to be dominated by seasonality, healthcare by a few large outliers but little seasonality, energy by large irregular variation, and the housing sector by components with roughly equal variation. Thus, there plausibly is considerable variation in the UCSV parameters across the 13 components.

These visual impressions are confirmed by the posterior estimates for 13-sector model. Table 3 summarizes some key results. Consider the standard deviations of the innovations in the idiosyncratic components: the estimated values of the $\sigma_\tau/\sigma_\varepsilon$ signal-to-noise ratios range from a high of 1.8 (furnishing) to a low of 0.2 (food and energy). Seasonal signal-to-noise ratios ($\sigma_s/\sigma_\varepsilon$) vary from nearly 4 (clothing) to 0.05 (energy). Most of these standard deviations are reasonably stable over the 2001–2018 sample, but there are exceptions: for example, seasonal fluctuations have become larger in recreation, and irregular fluctuations have become smaller in alcohol and tobacco.

The multivariate model captures the covariance across sectors through the common factors τ_c, s_c, and ε_c. The estimated standard deviation of the innovations in these factors fell by roughly 40 percent from 2001 to 2018; this implies a reduction in the co-variability across the sectors. The estimated factor loadings suggest that much of the comovement arises from the common trend component, less from common seasonals, and very little from common irregular variation.

Table 3. Parameter Estimates from the 13-sector Multivariate UCSV Model

(a) Standard deviation of shocks to common components (τ_c, s_c, ε_c)

	2001	*2018*
σ_τ	0.99 (0.91, 1.04)	0.57 (0.30, 1.00)
σ_s	0.98 (0.89, 1.03)	0.62 (0.32, 1.00)
σ_ε	0.99 (0.91, 1.03)	0.67 (0.36, 1.00)

(b) Sector-specific parameters

	Factor loadings		
Sector	α_τ	α_s	α_ε
Food	0.72 (0.41, 1.04)	0.29 (0.17, 0.42)	0.05 (-0.32, 0.42)
Alcohol and tobacco	0.06 (-0.15, 0.34)	0.03 (-0.06, 0.14)	0.05 (-0.29, 0.41)
Clothing	0.29 (0.12, 0.47)	0.12 (0.05, 0.19)	0.03 (-0.28, 0.36)
Housing (xE)	0.03 (-0.03, 0.12)	0.01 (-0.01, 0.05)	-0.01 (-0.12, 0.09)
Furnishing	0.24 (0.11, 0.45)	0.10 (0.04, 0.18)	0.02 (-0.15, 0.19)
Healthcare	0.28 (0.16, 0.45)	0.12 (0.06, 0.18)	-0.02 (-0.22, 0.21)
Transportation (xE)	0.36 (0.24, 0.53)	0.15 (0.10, 0.22)	-0.04 (-0.40, 0.33)
Communications	-0.13 (-0.38, 0.23)	-0.05 (-0.16, 0.09)	0.02 (-0.37, 0.39)
Recreation	0.35 (0.18, 0.54)	0.14 (0.07, 0.22)	0.03 (-0.51, 0.53)
Education	0.26 (0.11, 0.46)	0.11 (0.04, 0.19)	-0.01 (-0.23, 0.21)
Restaurants and accommodations	0.47 (0.32 0.69)	0.19 (0.13, 0.28)	-0.04 (-0.41, 0.40)
Miscellaneous	0.18 (0.08, 0.30)	0.07 (0.03, 0.12)	0.01 (-0.20, 0.23)
Energy	0.39 (-0.02, 0.78)	0.16 (-0.01, 0.32)	-0.01 (-0.45, 0.42)

Source: Authors' calculations.

Notes: The values shown are the posterior median and 68-percent equal-tail posterior credible intervals for the dates shown.

Table 3. Parameter Estimates from the 13-sector Multivariate UCSV Model

(b) Sector-specific parameters *(continued)*

	Standard deviation of shocks to sector-specific components $(\tau_i, s_i, \varepsilon_i)$					
	σ_τ		σ_s		σ_ε	
Sector	*2001*	*2018*	*2001*	*2018*	*2001*	*2018*
Food	0.30 (0.12, 0.94)	0.28 (0.12, 0.74)	0.17 (0.10, 0.31)	0.17 (0.09, 0.31)	1.60 (1.15, 2.16)	1.54 (1.19, 1.92)
Alcohol and tobacco	0.28 (0.14, 0.53)	0.29 (0.15, 0.50)	0.50 (0.17, 0.89)	0.40 (0.16, 0.69)	1.41 (0.69, 2.60)	0.92 (0.53, 1.34)
Clothing	0.13 (0.07, 0.23)	0.12 (0.07, 0.20)	1.56 (1.11, 2.13)	1.14 (0.79, 1.60)	0.31 (0.15, 0.51)	0.31 (0.15, 0.49)
Housing (xE)	0.13 (0.09, 0.18)	0.13 (0.09, 0.18)	0.10 (0.07, 0.14)	0.11 (0.08, 0.15)	0.13 (0.08, 0.19)	0.13 (0.08, 0.19)
Furnishing	0.21 (0.14, 0.27)	0.22 (0.17, 0.30)	0.15 (0.11, 0.20)	0.15 (0.11, 0.20)	0.12 (0.07, 0.20)	0.12 (0.07, 0.20)
Healthcare	0.12 (0.07, 0.20)	0.11 (0.07, 0.19)	0.13 (0.08, 0.21)	0.13 (0.08, 0.21)	0.77 (0.60, 1.02)	0.47 (0.32, 0.66)
Transportation (xE)	0.11 (0.07, 0.19)	0.11 (0.07, 0.20)	0.26 (0.15, 0.39)	0.28 (0.17, 0.40)	0.37 (0.21, 0.52)	0.39 (0.22, 0.55)
Communications	0.69 (0.46, 1.01)	0.69 (0.48, 1.01)	0.14 (0.08, 0.26)	0.14 (0.08, 0.25)	0.93 (0.59, 1.21)	0.87 (0.53, 1.15)
Recreation	0.17 (0.10, 0.27)	0.17 (0.10, 0.26)	0.35 (0.21, 0.54)	0.68 (0.46, 1.00)	0.28 (0.14, 0.50)	0.32 (0.15, 0.65)
Education	0.15 (0.09, 0.24)	0.15 (0.09, 0.24)	0.23 (0.11, 0.38)	0.22 (0.11, 0.39)	0.71 (0.51, 0.89)	0.77 (0.61, 0.97)
Restaurants and accommodations	0.15 (0.09, 0.23)	0.14 (0.09, 0.22)	0.23 (0.14, 0.37)	0.38 (0.22, 0.57)	0.24 (0.16, 0.35)	0.22 (0.13, 0.36)
Miscellaneous	0.16 (0.11, 0.22)	0.17 (0.12, 0.25)	0.17 (0.11, 0.25)	0.16 (0.10, 0.26)	0.22 (0.14, 0.34)	0.27 (0.18, 0.43)
Energy	1.38 (0.43, 2.63)	1.47 (0.45, 2.77)	0.40 (0.14, 1.03)	0.41 (0.15, 1.08)	7.15 (4.78, 9.28)	8.12 (6.25, 10.32)

Source: Authors' calculations.
Notes: The values shown are the posterior median and 68-percent equal-tail posterior credible intervals for the dates shown.

The multivariate model produces a rich set of results. Figures 5 and 6 illustrate a few of these results. The first four panels of figure 5 show selected results for the transportation sector: the raw data and seasonally adjusted values $(\pi_{i,t} - s_{i,t})$ are plotted in panel (a), the trend and seasonally adjusted values are plotted in panel (b), the seasonals are shown panel (c), and the estimated seasonal standard deviations, $\sigma_{i,\mathrm{s},t}$, are shown in panel (d). Evidently, the multivariate UCSV model accommodates the increased dispersion in the seasonal evident in panel (c) with increases in $\sigma_{i,s,t}$ in panel (d), and provides a reasonably sharp decomposition into trend, seasonal, and irregular components in panel (b). Panels (e)–(h) show the same results for the clothing sector. From panel (e), seasonal variation in clothing price inflation is so large that it is difficult to discern any variation in the seasonally adjusted series. A change of scale in panel (f) makes the variation in the seasonally adjusted series visible and shows an outlier in 2011. Panel (g) shows that the variance of the seasonal component increases in the first half of the sample, but remains large and approximately constant, in the second half of the sample. The estimates of $\sigma_{i,s,t}$ shown in panel (h) are consistent with this changing seasonal variability. Panel (i) plots healthcare inflation and shows two large outliers. Panel (j) shows the posterior mean estimates of the outlier factor $o_{i,t}$ for healthcare, which successfully pinpoints the outliers in panel (i). Panels (k) and (l) show the analogous results for the energy sector, where outliers are also an important source of variability.

Figure 6 shows the trend estimates for each of the 13 sectors. The sectorial trends differ, but comovement is apparent, most notably during the cyclical downturns in 2008–10 and 2014–15.

As discussed above, the estimates of τ_t from the univariate model are constructed by using weighted averages of aggregate inflation, where the weights sum to unity; the one- and two-sided weights were plotted in figure 4. In the multivariate model, estimates of τ_t are also weighted averages of leads and lags of inflation for each of the sectors. When share weights and variances are time-invariant, lead-lags weights on each sector sum to that sector's share weight. For sectors with low signal-to-ratios, substantial weight is placed on distant leads and lags, but for sectors with high signal-to-ratios, most of the weight is concentrated near the contemporaneous value of $\pi_{i,t}$. figure 7 plots the sector-specific optimal weights from the 13-sector model, and compares these to the weights for the 1-sector model (which are identical for all sectors). Relative to the 13-sector weights, the 1-sector model puts too much weight on contemporaneous values of food, alcohol, and energy inflation (which have a low signal-to-noise ratio) and too little weight on sectors like furnishing and restaurants (which have relatively high signal-to-noise ratios). An implication is that the estimates of the aggregate seasonal and trend components constructed from the sectorial model and data are more precise than the estimates by using only the aggregate data.

Figure 5. Selected Results from the 13-sector UCSV Model

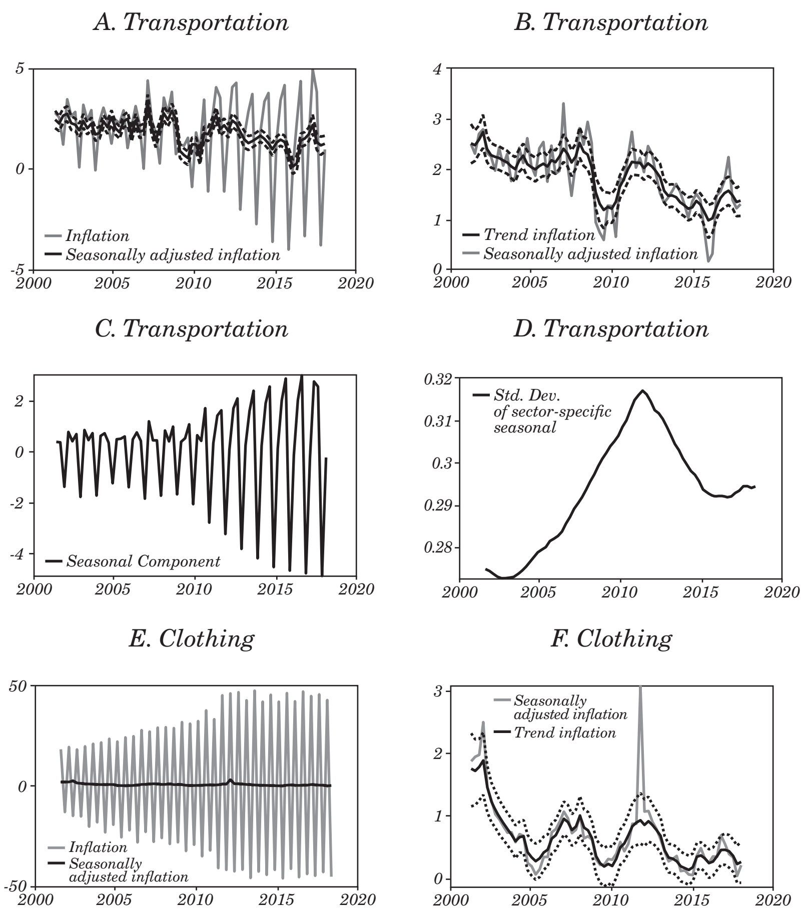

Figure 5. (continued)

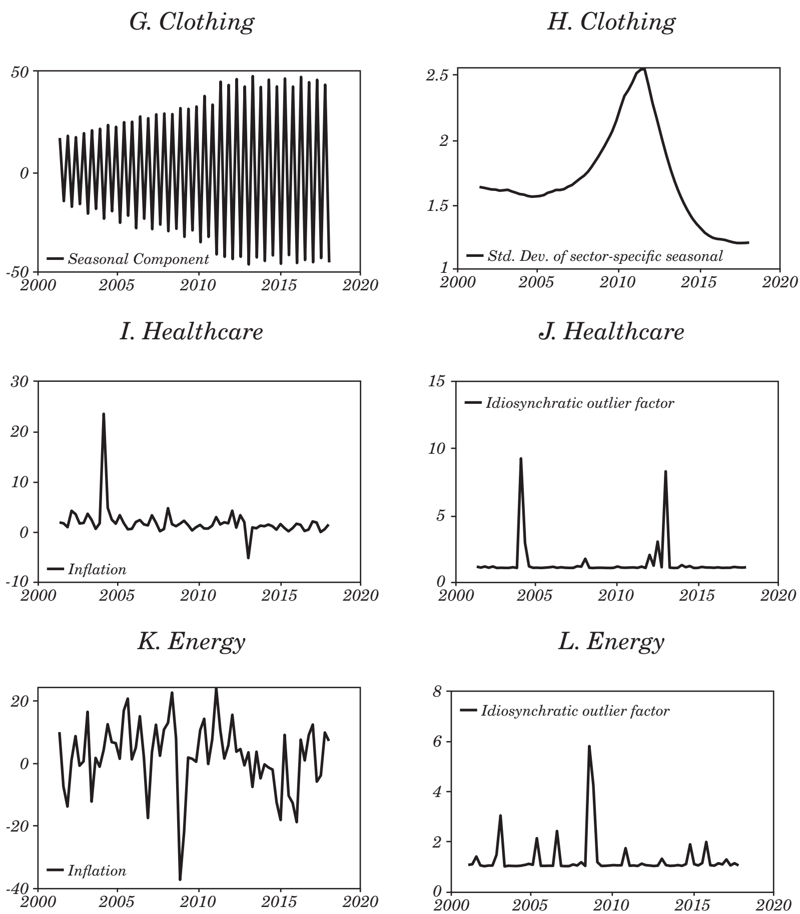

Source: Authors' calculations.
Notes: See text for description of the panels. Error bands are 68-percent posterior credible intervals.

Figure 6. Trend Estimates from the 13-sector UCSV Model

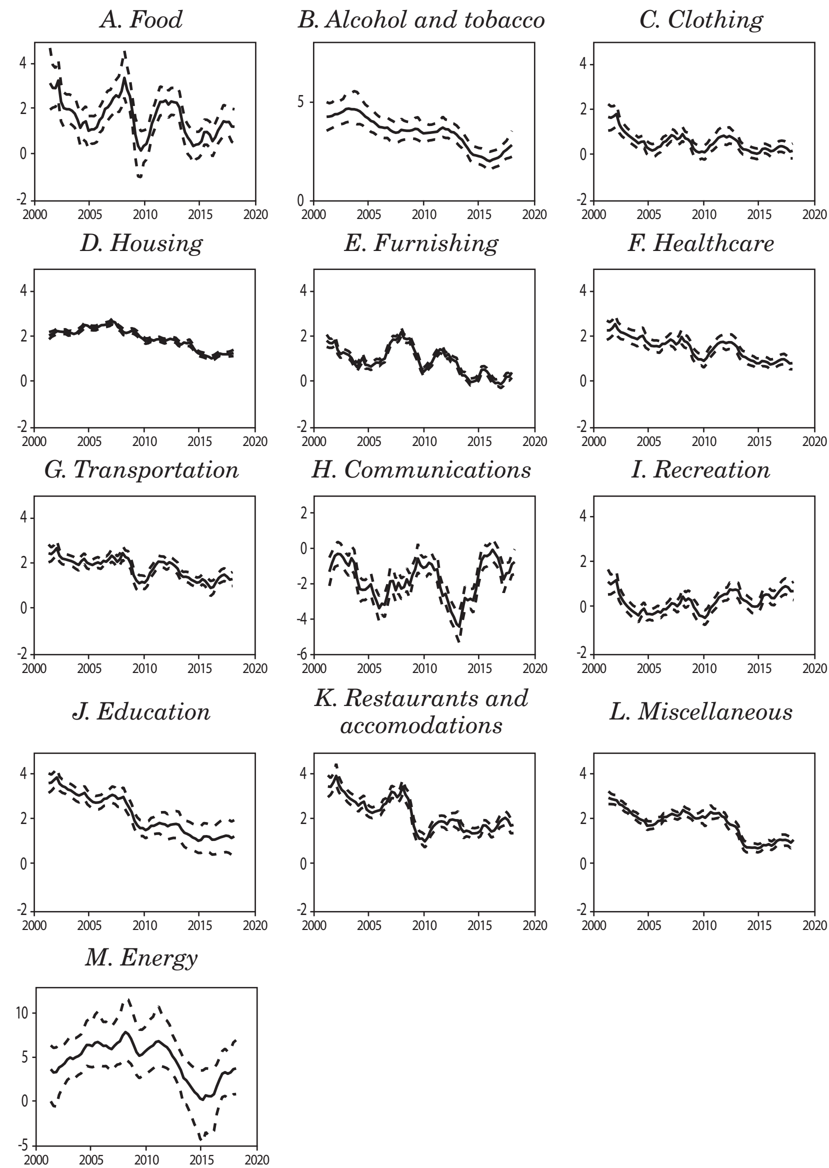

Source: Authors' calculations.
Notes: These are the 12 HICP tier-two sectors, with energy excluded from the housing and transportation sectors, and shown separately as the 13th sector.

Figure 7. Weight Placed on $\pi_{i,t+j}$ for Estimating Aggregate τ_t

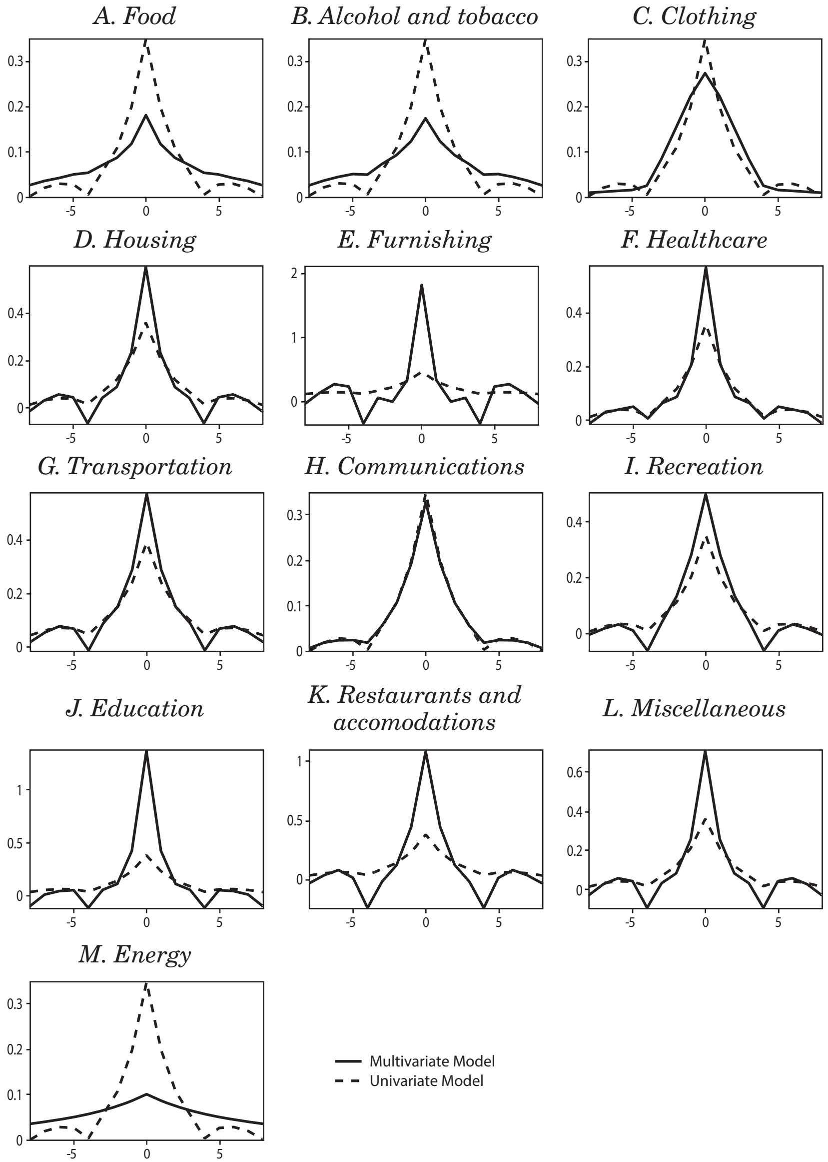

Source: Authors' calculations.

Notes: Values shown are the sum of the Kalman smoother weights on $\pi_{i,t+j}$ for estimating π_t. The results in the first row are from the univariate model for aggregate inflation. Weights are normalized by expenditure shares, so the weights for all sectors sum to unity over all leads and lags.

This improved precision from the multivariate model can be seen in table 4 and figure 8, which show aggregate estimates constructed as share-weighted averages of the sectorial components. Comparing the error bands in table 4 with the corresponding error bands for the univariate model in table 2 shows a tightening of the bands for the multivariate model. For example, the multivariate errors bands for $\tau_{2018:Q1}$ are roughly 80 percent as wide as the univariate bands, and the multivariate error bands for $s_{2018:Q1}$ are roughly 60 percent as wide as the univariate bands.

Figure 8. Smoothed and Filtered Estimates from 13-sector Multivariate UCSV Model for Aggregate HICP Inflation

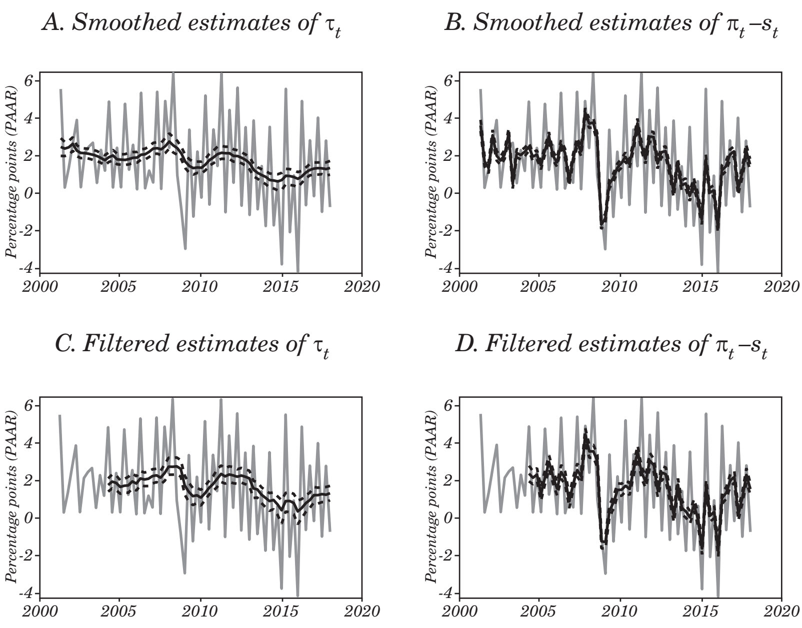

Source: Authors' calculations.
Notes: The values shown are the posterior median and 68-percent equal-tail posterior credible intervals for the dates shown. Aggregate values are computed as share-weighted averages of the sectorial values.

Table 4. Selected Results for Aggregate Inflation from the 13-sector UCSV Model

(a) Estimated trends from the multivariate model

	2001:Q2	*2009:Q4*	*2018:Q1*
τ_t	2.43 (1.95 2.92)	1.35 (1.05 1.66)	1.32 (0.91 1.73)

(b) Estimated seasonal factors

	Q1	*Q2*	*Q3*	*Q4*
2002	-0.44 (-0.75 -0.14)	1.95 (1.69 2.22)	-1.48 (-1.71 -1.24)	0.07 (-0.18 0.32)
2009	-1.65 (-1.89 -1.40)	2.94 (2.69 3.17)	-2.13 (-2.37 -1.89)	0.89 (0.66 1.13)
2017	-2.42 (-2.71 -2.13)	3.87 (3.53 4.18)	-2.09 (-2.38 -1.81)	0.52 (0.19 0.88)

Source: Authors' calculations.
Notes: The values shown are the posterior median and 68-percent equal-tail posterior credible intervals for the dates shown. Aggregate values are computed as share-weighted averages of the sectorial values.

2.3 Different Levels of Disaggregation

The results presented thus far show that the 13-sector multivariate trend and seasonal estimates are more accurate than estimates that only use aggregate inflation. A natural question to ask is how much of these gains could be achieved by using a coarser disaggregation scheme, for example by using a three-sector decomposition of food, energy, and the aggregate of all of the other sectors. Using data for the U.S., Stock and Watson (2016) found that much of the gain from using a 17-sector decomposition of U.S. PCE inflation could be achieved by using this three-sector decomposition. Can similar gains be achieved from the euro-area HICP?

To answer this question, we estimated three additional multivariate UCSV models. The first is a two-sector model composed of energy and HICP excluding energy. The second is a three-sector decomposition composed of food, energy, and HICP excluding food and energy. The third is a four-sector decomposition that uses third-tier components to further decompose the non-food-and-energy HICP into goods and services. The two- and three-sector models are special cases of the 13-sector model; the four-sector model is not: as discussed above, the second-tier decomposition in the 13-sector model includes goods and services jointly in many of the sectors.

Figure 9. Estimates of Trend Inflation from the Various UCSV Models

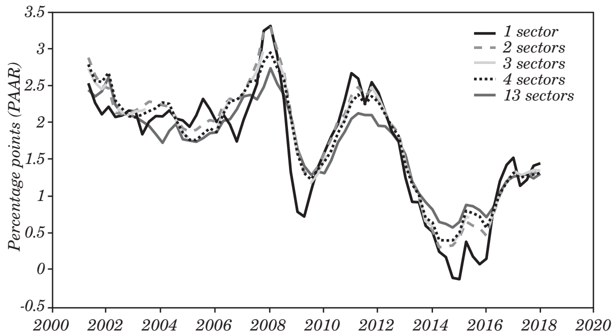

Source: Authors' calculations.
Notes: Values shown are full-sample posterior medians.

Table 5. Width of Credible Intervals, Final Quarter

Model	*68% credible interval*		*90% credible interval*	
	τ	$\pi - s$	τ	$\pi - s$
Univariate	1.09	1.00	1.89	1.77
2 sectors	0.96	0.81	1.63	1.39
3 sectors	0.87	0.68	1.53	1.19
4 sectors	0.82	0.67	1.48	1.19
13 sectors	0.82	0.62	1.45	1.09

Source: Authors' calculations.
Notes: The values are the widths of 68-percent and 90-percent credible intervals for τ and $\pi - s$ for the final quarter in the sample (2018:Q1).

Figure 9 plots the estimates of trend inflation computed for each model. The estimated trends are generally similar, although there are noteworthy differences between the one- and multi-sector trends during 2009 and 2015.[5] Table 5 summarizes the accuracy of these alternative models by showing the final quarter (2018:Q1) width of the 68-percent and 90-percent error bands for trend and seasonally adjusted inflation. Each decomposition yields marginal improvements,

5. This paper has taken a multivariate approach to trend (and seasonal adjustment) of aggregate inflation by using sectorial inflation rates. Other series beyond sectorial inflation rates may also help identify trend inflation. Mertens (2016) provides an interesting application by using inflation expectations and nominal interest rates as additional indicators.

but much of the gain can be achieved by using the three-sector decomposition; this is consistent with the results for the U.S. reported in Stock and Watson (2016).

3. Inflation and Real Activity

The multivariate estimates of trend inflation suggest a large variation in the trend level of inflation over the 2001–2018 sample period. Figure 10 shows how this variation in inflation was related to variation in real economic activity, where real activity is measured as an average of three coincident indicators for the euro area: the unemployment gap (inverted), capacity utilization, and the logarithm of industrial production, each band-pass filtered to isolate business-cycle variation (6-32 quarters) and standardized to have zero mean and unit variance. Over 2001–2018, changes in trend inflation closely mirrored changes in real activity: trend inflation increased to nearly 3 percent in early 2008 as activity was near its cyclical peak, fell by 1.5 percent during the 2009 recession, returned to 2 percent during the recovery, but fell again to under 1 percent as real activity weakened during 2013–2016.

Figure 10. Inflation and Real Activity

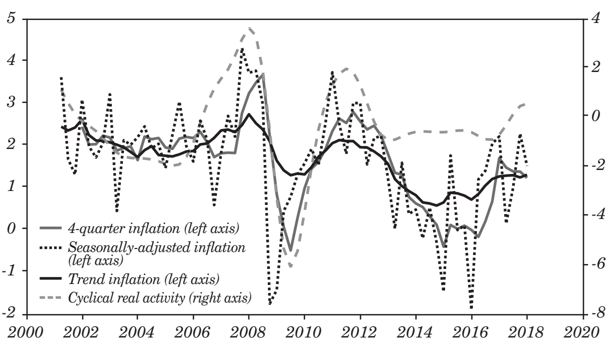

Source: Authors' calculations.
Notes: The trend and seasonally adjusted inflation values are the full-sample posterior medians from the 13-sector UCSV model. The cyclical activity index is the average of standardized band-pass filtered values of the unemployment gap (inverted), the capacity utilization rate, and the logarithm of industrial production, for a pass band of 6–32 quarters.

Table 6 presents correlations between the cyclical activity index and various measures of HICP inflation. The lowest correlation is with seasonally unadjusted quarterly inflation, and the highest (0.55) is with four-quarter inflation. As can be seen in figure 10, 4-quarter inflation falls sharply with economic activity in the 2009 recession, whereas trend inflation falls less, hence it has a somewhat lower correlation with the cyclical activity index. These correlations are all substantial and are consistent with a Phillips relation being present in euro-area inflation.

Table 6. Width of Credible Intervals, Final Quarter

Inflation measure	*Correlation*
Quarterly inflation	0.20
4-quarter inflation (100Δln(Pt/Pt-4))	0.55
Seasonally adjusted HICP	0.42
Univariate trend	0.43
3-sector trend estimate	0.47
13-sector trend estimate	0.44

Source: Authors' calculations.
Notes: Seasonally adjusted HICP is the smoothed estimate of $\pi_t - s_t$ computed by using the univariate UCSV model. The three trend estimates are computed by using the UCSV model (univariate or multivariate, depending on the estimate). The cyclical activity index is the average of standardized band-pass filtered values of the unemployment gap (inverted), the capacity utilization rate, and the logarithm of industrial production, for a pass band of 6-32 quarters).

References

Ball, L.M. and S.G. Cecchetti. 1990. "Inflation and Uncertainty at Short and Long Horizons." *Brookings Papers on Economic Activity* 1: 215–54.

Bryan, M.F. and S.G. Cecchetti. 1994. "Measuring Core Inflation." In *Monetary Policy,* edited by N. G. Mankiw. Chicago, IL: University of Chicago Press.

Carter, C.K. and R. Kohn 1994. "On Gibbs Sampling for State Space Models." *Biometrika* 81: 541–53.

Cecchetti, S., P. Hooper, B. Kasman, K. Schoenholtz, and M. Watson. 2007. "Understanding the Evolving Inflation Process." *U.S. Monetary Policy Report 1*.

Cogley, T. and T.J. Sargent. 2015. "Measuring Price-Level Uncertainty and Instability in the United States, 1850-2012." *Review of Economics and Statistics* 97(4): 827–38.

Cogley, T., T.J. Sargent, and P. Surico. 2015. "Price-Level Uncertainty and Instability in the United Kingdom." *Journal of Economic Dynamics and Control* 52: 1–16.

Eckstein, O. 1981. *Core Inflation*. New York: Prentice Hall.

Ehrmann, M., G. Ferrucci, M. Lenza, and D. O'Brien. 2018. "Measures of Underlying Inflation in the Euro Area." ECB *Economic Bulletin* No. 4.

Gordon, R.J. (1975), "Alternative Responses of Policy to External Supply Shocks." *Brookings Papers on Economic Activity* 1: 183–206.

Harvey, A.C. 1989. *Forecasting, Structural Time Series Models, and the Kalman Filter.* Cambridge MA: Cambridge University Press.

Hausman, J.A. and M.W. Watson. 1985. "Errors in Variables and Seasonal Adjustment Procedures." *Journal of the American Statistical Association* 80(391): 531–40.

Hillmer, S.C. and G.C. Tiao. 1982. "An ARIMA-Model-Based Approach to Seasonal Adjustment." *Journal of the American Statistical Association* 77: 63–70.

Kang, K. H., C.J. Kim, and J. Morley. 2009. "Changes in U.S. Inflation Persistence." *Studies in Nonlinear Dynamics and Econometrics* 13(4): 1–21.

Kim, C.J. and C.R. Nelson. 1999. *State-Space Models with Regime Switching: Classical and Gibbs Sampling Approaches with Applications.* Cambridge, MA: MIT Press.

Kim, S., N. Shephard, and S. Chib. 1998. “Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models.” *Review of Economic Studies* 65: 361–93.

Maravall, A. 1995. “Unobserved Components in Economic Time Series.” In *The Handbook of Applied Econometrics*, edited by M.H. Pesaran and M. Wickens. Oxford: Basil Blackwell.

Mertens, E. 2016. “Measuring the Level and Uncertainty of Trend Inflation.” *Review of Economic Statistics* 98(5): 950–67.

Nerlove, M., D.M. Grether, and J.L. Carvalho. 1979. *Analysis of Economic Time Series.* New York, NY: Academic Press.

Omori, Y., S. Chib, N. Shephard, and J. Nakajima. 2007. “Stochastic Volatility with Leverage: Fast and Efficient Likelihood Inference.” *Journal of Econometrics* 140: 425–49.

Stock, J.H. and M.W. Watson. 2007. “Why Has U.S. Inflation Become Harder to Forecast?” *Journal of Money, Banking, and Credit* 39(1): 3–33. (Also see Erratum in vol. 39, No. 7, 2007)

————. 2016. “Trend and Core Inflation.” *Review of Economics and Statistics* 98(4): 770–84.

Series on Central Banking, Analysis, and Economic Policies

The Book Series on "Central Banking, Analysis, and Economic Policies" of the Central Bank of Chile publishes new research on central banking and economics in general, with special emphasis on issues and fields that are relevant to economic policies in developing economies. Policy usefulness, high-quality research, and relevance to Chile and other open economies are the main criteria for publishing books. Most research published by the Series has been conducted in or sponsored by the Central Bank of Chile.

Volumes in the series:

1. Análisis empírico del ahorro en Chile
   Felipe Morandé and Rodrigo Vergara, editors
2. Indexation, Inflation, and Monetary Policy
   Fernando Lefort and Klaus Schmidt–Hebbel, editors
3. Banking, Financial Integration, and International Crises
   Leonardo Hernández and Klaus Schmidt–Hebbel, editors
4. Monetary Policy: Rules and Transmission Mechanisms
   Norman Loayza and Klaus Schmidt–Hebbel, editors
5. Inflation Targeting: Design, Performance, Challenges
   Norman Loayza and Raimundo Soto, editors
6. Economic Growth: Sources, Trends, and Cycles
   Norman Loayza and Raimundo Soto, editors
7. Banking Market Structure and Monetary Policy
   Luis Antonio Ahumada and J. Rodrigo Fuentes, editors
8. Labor Markets and Institutions
   Jorge Enrique Restrepo and Andrea Tokman R., editors
9. General Equilibrium Models for the Chilean Economy
   Rómulo Chumacero and Klaus Schmidt–Hebbel, editors
10. External Vulnerability and Preventive Policies
    Ricardo J. Caballero, César Calderón, and Luis Felipe Céspedes, editors
11. Monetary Policy under Inflation Targeting
    Frederic S. Mishkin and Klaus Schmidt–Hebbel, editors
12. Current Account and External Financing
    Kevin Cowan, Sebastián Edwards, and Rodrigo Valdés, editors
13. Monetary Policy under Uncertainty and Learning
    Klaus Schmidt–Hebbel and Carl E. Walsh, editors
14. Banco Central de Chile 1925-1964, Una Historia Institucional
    Camilo Carrasco, editor
15. Financial Stability, Monetary Policy, and Central Banking
    Rodrigo A. Alfaro, editor
16. Monetary Policy under Financial Turbulence
    Luis Felipe Céspedes, Roberto Chang, and Diego Saravia, editors
17. Fiscal Policy and Macroeconomic Performance
    Luis Felipe Céspedes and Jordi Galí, editors
18. Capital Mobility and Monetary Policy
    Miguel Fuentes D., Claudio E. Raddatz, and Carmen M. Reinhart, editors
19. Macroeconomic and Financial Stability:
    Challenges for Monetary Policy
    Sofía Bauducco, Lawrence Christiano, and Claudio Raddatz, editors
20. Global Liquidity, Spillovers to Emerging Markets and Policy Responses
    Claudio Raddatz, Diego Saravia, and Jaume Ventura, editors
21. Economic Policies in Emerging-Market Economies
    Festschrift in Honor of Vittorio Corbo
    Ricardo J. Caballero and Klaus Schmidt–Hebbel, editors
22. Commodity Prices and Macroeconomic Policy
    Rodrigo Caputo and Roberto Chang, editors
23. 25 Años de Autonomía del Banco Central de Chile
    Alberto Naudon D. and Luis Álvarez V., editors
24. Monetary Policy through Asset Markets: Lessons
    from Unconventional Measures and Implications
    for an Integrated World
    Elías Albagli, Diego Saravia, and Michael Woodford
25. Monetary Policy and Global Spillovers: Mechanisms,
    Effects, and Policy Measures
    Enrique G. Mendoza, Ernesto Pastén, and Diego Saravia
26. Monetary Policy and Financial Stability: Transmission Mechanisms
    and Policy Implications
    Álvaro Aguirre, Markus Brunnermeier, and Diego Saravia

Changing Inflation Dynamics, Evolving Monetary Policy

Understanding the dynamics of inflation is one of the most important challenges that macroeconomists face. The long period of declining unemployment combined with persistently low inflation has proved a puzzle for traditional macroeconomic frameworks. This outstanding collection of papers digs deeply into the issue, providing refreshing new insights, both empirical and theoretical. The volume is a must read for anyone interested in researching inflation dynamics, as well as policy makers interested in knowing the state of the art on the subject. I recommend it enthusiastically.

Mark Gertler
Henry and Lucy Moses Professor of Economics at New York University

This book was ahead of its time in tackling a key issue that continues to puzzle policymakers and academics: why has inflation remained muted despite sharp falls in unemployment? The chapters explore a wide range of hypotheses and draw on the diverse country experience from an impressive group of authors. Their analysis provides important new insights and explores several theories that have previously received little attention —some of which appear to go some way to resolving the "missing inflation" puzzle.

Kristin J. Forbes
Jerome and Dorothy Lemelson Professor of Management and Global Economics at MIT's Sloan School of Management

Recent inflation dynamics have been a major puzzle for academics and policymakers. Missing disinflation during the Great Recession, an apparently weakened Phillips curve, missing inflation in the aftermath of the Global Financial Crisis are just some of the riddles. In the "*Changing Inflation Dynamics, Evolving Monetary Policy*" volume, world-class scholars give their expert takes and shed new light on these issues. Methodological diversity, creative thinking, and lucid presentation are the hallmarks of this volume.

Yuriy Gorodnichenko
Quantedge Presidential Professor of Economics at University of California, Berkeley

BANCO CENTRAL DE CHILE