

Evolution and Intelligent Design

Thomas J. Sargent[‡]

November 10, 2007

Abstract

This paper discusses two sources of ideas that influence monetary policy makers today. The first is a set of analytical results that impose the rational expectations equilibrium concept and do ‘intelligent design’ by solving Ramsey and mechanism design problems. The second is the adaptive learning process that first taught us how to anchor the price level with a gold standard, then how to replace the gold standard with a fiat currency wanting nominal anchors. Models of out-of-equilibrium learning tell us that such an adaptive evolutionary process will converge to a self-confirming equilibrium (SCE). In an SCE, a government’s probability model is correct about events that occur under the prevailing government policy, but possibly wrong about the consequences of other policies. That causes mistakes absent from a rational expectations equilibrium and enlarges the role of learning.

Keywords: Rational expectations equilibrium, optimal monetary policy, model misspecification, learning, evolution, observational equivalence, self-confirming equilibrium. (JEL).

*New York University and Hoover Institution. Email: ts43@nyu.edu.

[‡]This draft of my presidential address to the American Economic Association in January 2008 continues a long conversation I have had with Chris Sims (see Sims (1982)). After Bob Lucas read Sargent (1984), he wrote me that “With friends like you, Chris doesn’t need enemies.” Maybe it is more complicated than that. I thank Gadi Barlevy, Francisco Barillas, Marco Bassetto, Alberto Bisin, William Branch, In-Koo Cho, Timothy Cogley, Lars Peter Hansen, Larry Jones, Kenneth Kasa, Narayana Kocherlakota, Ramon Marimon, Athanasios Orphanides, Carolyn Sargent, Hyun Shin, Christopher Sims, Viktor Tsyrennikov, Harald Uhlig, François Velde, Carl Walsh, Noah Williams, Peyton Young, and Tao Zha for helpful comments. I thank the National Science Foundation for research support.

1 Introduction

... the ideas of economists and political philosophers, both when they are right and when they are wrong, are more powerful than commonly understood. Indeed the world is ruled by little else. Practical men, who believe themselves to be quite exempt from any intellectual influences, are usually the slaves of some defunct economist. Madmen in authority, who hear voices in the air, are distilling their frenzy from some academic scribbler of a few years back. Keynes (1936, p. 383)¹

Today leading practical men and women at important institutions in my field, the central banks, are distinguished academic economists. They have used prevailing academic ideas to choose policies. Some designed new institutions. For example, in 1997 Mervyn King and others created new decision-making protocols for Britain's monetary policy committee.

This essay is about two important sources of prevailing ideas in macroeconomics. The first is a collection of powerful theoretical results and empirical methods described in sections 2, 3, and 4 that apply the rational expectations equilibrium concept to estimate models and design optimal macroeconomic policies intelligently. The second is an adaptive evolutionary process, modelled in section 5 and illustrated in sections 6 and 7, that has left us an historical economic record littered with discarded ideas and policies.

The rational expectations equilibrium concept equates all subjective distributions with an objective distribution. It is useful to distinguish the step of equating all subjective distributions from the step of equating subjective distributions to the objective distribution that actually governs outcomes. By equating subjective distributions for endogenous variables to an equilibrium distribution implied by a model, the rational expectations hypothesis makes agents' beliefs disappear as extra components of a theory and sets the stage for a variety

¹A younger Keynes was less optimistic about the influence of economists' ideas:

Financiers of this type [Lord Rothschild, Lord Avebury, Lord Swaythling] will not admit the feasibility of anything until it has been demonstrated to them by practical experience. It follows, therefore, that they will seldom give their support to what is new. Keynes (1913, pp. 24-25)

of powerful theoretical results and intelligent policy design exercises described in section 2. Section 3 describes theoretical and practical reasons for equating subjective distributions to an objective one and how it facilitates the rational expectations econometrics described in section 4.

The assumption that agents share common beliefs underpins influential doctrines about whether observed inflation-unemployment dynamics can be exploited by policy makers, the time inconsistency of benevolent government policy, the capacity of reputation to substitute for commitment, the incentives for one type of policy maker to emulate another, and the wisdom of making information public. The common beliefs assumption is especially stressed in modern theories of optimal macroeconomic policy that focus on how a benevolent government optimally shapes expectations. This intelligent design approach to macroeconomic policy perfects an older econometric policy evaluation method that Lucas (1976) criticized because it imputed different beliefs to the government and the other agents.

Intelligent design is normative ('what should be') economics, but when it influences policy makers, it becomes positive ('what is') economics, as asserted in the epigraph from Keynes. Some researchers in the intelligent design tradition ignore the distinction between positive and normative economics from the start. Thus, a standard tool for understanding observed time series properties of government debt and taxes is to apply a normative analysis, e.g., Barro (1979), Lucas and Stokey (1983), and Aiyagari et al. (2002). It is also true that some policy advisors have enough faith that evolution produces good outcomes to recommend copying best practices (for example, see Keynes (1913)). If only good things survive the tests of time and practice, evolution produces intelligent design.

Theories of out-of-equilibrium learning tell us not always to expect that.² Section 5 extends a disturbing possibility that emerges from recent developments in the rational expectations econometrics of section 4 to describe how a system of adaptive agents converges to a self-confirming equilibrium in which all agents have correct forecasting distributions for

²Nevertheless, subsection 7.2 identifies special features of a model of Primiceri (2006) that makes outcomes converge to the best of all possible worlds.

events that occur often enough along an equilibrium path, but possibly mistaken views about policies and outcome paths that are not observed. This matters because intelligent design of rational expectations equilibria hinges on knowing and manipulating expectations about events that will not be observed. Self-confirming equilibria allow models to survive that imply mistaken policies even though they match historical data. Section 6 mentions examples from a millenium of monetary history. To tell some stories about the emergence of U.S. inflation in the 1970s and its conquest under Volcker and Greenspan, section 7 uses adaptive models in which the government solves intelligent design problems using probability models that are misspecified, either permanently or temporarily. While these stories differ in many interesting details, they all say that choices of the monetary authorities were determined by misunderstandings that do not occur within a rational expectations equilibrium.³ Section 8 concludes by commenting on some of the reasons that policy makers and students of macroeconomic history have thought in terms of adaptive processes that are beyond rational expectations equilibria.

2 Intelligent design with common beliefs

By solving Pareto problems in which a planner and all agents optimize in light of information and incentive constraints and a common probability model, what I call intelligent design is a coherent response to Lucas's (1976) indictment of pre-rational expectations macroeconomic policy design procedures. Lucas accused those procedures of incorporating private agents' decision rules that were not best responses to government policy under an equilibrium probability measure. The cross-equation restrictions of common belief models fix that problem.

I use f to denote a probability density and x^t to denote a history x_t, x_{t-1}, \dots, x_0 . Partition $x_t = [y_t \ v_t]'$, where v_t is a vector of decisions taken by a government and y_t is a vector of

³These adaptive models make room for a 'law of unintended consequences' cited by Friedman (1991) that is excluded from rational expectations equilibria.

all other variables. Let $f(y^\infty, v^\infty|\rho)$ be a joint density conditional on a parameter vector $\rho \in \Omega_\rho$. Government chooses a sequence h of functions

$$v_t = h_t(x^t|\rho), \quad t \geq 0, \quad (1)$$

to maximize a Pareto criterion that can be expressed as expected utility under density $f(x^\infty|\rho)$:

$$\int U(y^\infty, v^\infty|\rho) f(y^\infty, v^\infty|\rho) d(y^\infty, v^\infty). \quad (2)$$

Modern intelligent design in macroeconomics solves government programming problems (2) with models f that impute common beliefs and best responses to all of the agents who inhabit the model. The common beliefs assumption used to construct the macroeconomic model makes parameters describing agents' beliefs about endogenous variables disappear from the vector ρ .

The common beliefs assumption underlies a long list of interesting results in modern macro. The following four have especially influenced thinking within central banks.

1. *Expected versus unexpected government actions.* Lucas (1972b) drew a sharp distinction between the effects of foreseen and unforeseen monetary and fiscal policies when the government and the public share a probability model. That idea defines the terms in which central bankers now think about shocks and systematic policies.
2. *Optimal fiscal and monetary policy cast as Ramsey and mechanism design problems.* A literature summarized and extended by Clarida et al. (1999) and Woodford (2003) uses dynamic macroeconomic models with sticky prices to design monetary policy rules by solving Ramsey plans like (2) and finding good ways to represent and implement them. New dynamic models of public finance refine Ramsey plans by focusing on a tradeoff between efficiency and incentives that emerges from the assumption that each individual privately observes his own skills and effort, a feature that imposes constraints on the allocations that the planner can implement relative to ones he could

achieve if he had more information.⁴

3. *Time consistency.* The availability of the rational expectations equilibrium concept enabled Kydland and Prescott (1977) and Calvo (1978) to explain how alternative timing protocols affect a benevolent government's capacity to manipulate and then confirm prior expectations about its actions.⁵ The time consistency 'problem' is the observation that equilibrium outcomes in a representative-agent economy depend on the timing protocol for decision making that nature or the modeler imposes. Better outcomes emerge if a government chooses a history-contingent plan once-and-for-all at time 0 than if it chooses sequentially. By choosing future actions at time 0, the government can take into account how expectations about its actions at times $t > 0$ influence private agents' actions at all dates between 0 and t . A government must ignore those beneficial expectations effects if it is forced to choose sequentially.

4. *Reputation can substitute for commitment.* Under rational expectations, a government strategy plays two roles, first, as a *decision rule* for the government and, second, as a *system of private sector expectations* about government actions that the government always wants to confirm.^{6,7} A system of expectations is a history-dependent

⁴See for example Golosov et al. (2003), Kocherlakota (2005), and Golosov et al. (2007).

⁵While technical treatments of the time consistency problem rely heavily on the rational expectations equilibrium concept, all that is needed to spot the problem is that private agents care about future government actions. In a discussion at the U.S. Constitutional Convention about whether the Federal government should be prohibited from fiduciary currency on August 16, 1787, Gouverneur Morris, Oliver Ellsworth, and James Madison recognized a time consistency problem, while Edmund Randolph and George Mason raised doubts about tying the hands of the government because no one can foresee all contingencies. See Madison (1987, pp. 470-471).

⁶The theory is silent about who chooses an equilibrium system of beliefs, the government (after all, it *is* the *government's* decision rule) or the public (but then again, they *are* the private sector's expectations). This and the multiplicity of equilibria make it difficult to use this theory to formulate advice to policy makers about actions that can help it to earn a good reputation. Instead, the theory is about how a government comes into a period confronting a set of private sector expectations about its actions that it will want to confirm. Blinder (1998, pp. 60-62) struggles with this issue when he describes pressures on the Fed not to disappoint the market. While Blinder's discussion can be phrased almost entirely within the rational expectations paradigm, the account by Bernanke (2007) of the problems the Fed experiences in anchoring private sector expectations cannot. Bernanke argues in terms of objects outside a rational expectations equilibrium.

⁷The theory of credible public policy seems to explain why some policy makers who surely knew about better decision rules chose instead to administer ones supporting bad outcomes. Chari et al. (1998) and Albanesi et al. (2002) interpret the big inflation of the 1970s and its stabilization in the 1980s in terms of the

government strategy like (1). A credible public policy is an equilibrium system of expectations that gives a government incentives to confirm prior expectations about its future actions, actions to which it cannot commit because it chooses sequentially.⁸ There are multiple equilibrium strategies, i.e., multiple systems of common expectations that a government would want to confirm, with incentive constraints tying good and bad equilibria together.

These theoretical rational expectations results have influenced the way monetary policy is now discussed within central banks. Because central banks want to implement solutions of Ramsey problems like (2) in contexts like (1) in which the distinction between the effects of foreseen and unforeseen policy actions is important, a time consistency problem like (3) arises, prompting them to focus on ways like (4) of sustaining good reputations.⁹

3 Justifications for equating objective and subjective distributions

These and many other theoretical results hinge on the part of the rational expectations equilibrium concept that equates subjective distributions for endogenous variables to an equilibrium distribution. To gain empirical content, rational expectations models also takes the logically distinct step of equating the equilibrium distribution to the data generating distribution. I shall use asset pricing theory to illustrate two justifications for taking that step, one based on an argument that agents with beliefs closer to the truth will eliminate others, another on empirical convenience.

actions of benevolent and knowledgeable policy makers who became trapped within but, thanks to a sunspot, eventually managed to escape expectations traps within subgame perfect or Markov perfect equilibria.

⁸See the credible public policy models of Stokey (1989, 1991) and Chari and Kehoe (1993b,a). By making an intrinsically ‘forward-looking’ variable, a promised discounted value for the representative household, also be a ‘backward-looking’ state variable that encodes history, Abreu et al. (1986, 1990) tie past and future together in a subtle way that exploits the common beliefs equilibrium concept. For some applications, see Chang (1998), Phelan and Stacchetti (2001), and Ljungqvist and Sargent (2004, ch. 22).

⁹See Blinder (1998) and Bernanke et al. (2001).

Hansen and Singleton (1983) and many others have generated restrictions on the co-variation of consumption and a one-period return $R_{j,t+1}(x_{t+1})$ on asset j by starting with consumer i 's Euler equation

$$1 = \beta \int_{x_{t+1}} \frac{u'_i(c_{i,t+1}(x^{t+1}))}{u'_i(c_{i,t}(x^t))} R_{j,t+1}(x_{t+1}) f_i(x_{t+1}|x^t) dx_{t+1} \quad (3)$$

where $f_i(x_{t+1}|x^t)$ is consumer i 's subjective one-step-ahead transition density for a state vector x_{t+1} that determines both returns and time $t+1$ consumption, $c_{i,t+1}$, and $u'_i(c_{i,t+1}(x^{t+1}))$ is consumer i 's marginal utility of consumption.

3.1 Complete markets and survival

In a finite-horizon setting, Harrison and Kreps (1979) showed that when there are complete markets, the *stochastic discount factor*

$$m_{t+1} = \beta \frac{u'_i(c_{i,t+1}(x^{t+1}))}{u'_i(c_{i,t}(x^t))} \frac{f_i(x_{t+1}|x^t)}{f(x_{t+1}|x^t)} \quad (4)$$

is unique. Here $f(x_{t+1}|x^t)$ is a common physical conditional density. Because offsetting differences in marginal utility functions and probabilities leave the left side of (4) fixed, the uniqueness of the stochastic discount factor allows different densities f_i . Suppose that density f actually governs outcomes. Then Blume and Easley (2006) showed that in complete markets economies with Pareto optimal allocations and an *infinite* horizon, for the subset of agents with positive wealths in the limit, all $f_i(x^\infty)$'s merge to the density that is closest to the truth $f(x^\infty)$.¹⁰ Merging means that the densities agree about tail events.¹¹ If $f_i(x^\infty) = f(x^\infty)$ for some i , then for an infinite horizon complete markets economy with a Pareto optimal allocation, this survival result implies the rational expectations assumption,

¹⁰Closest as measured by Kullback and Leibler's relative entropy.

¹¹In the context of a complete markets economy with a Lucas tree, Sandroni (2000) argued that a disagreement about tail events would present some consumers with arbitrage opportunities that cannot exist in equilibrium.

provided that agents have access to an infinite history of observations at time 0.

3.2 Incomplete markets

Grossman and Shiller (1981), Hansen and Singleton (1983), and Hansen and Richard (1987) wanted an econometric framework for incomplete markets, in which case it is not enough to appeal to Blume and Easley's market survival justification for assuming that beliefs are common or eventually common. Hansen and Singleton (1983) and Hansen and Richard (1987) did not let that stop them. They simply imposed rational expectations directly and made enough stationarity assumptions to validate a Law of Large Numbers that gives GMM or maximum likelihood estimators good asymptotic properties. Under the rational expectations assumption, (3) imposes testable restrictions on the empirical joint distribution of returns and either individual or aggregate consumption.

3.3 An empirical reason to allow belief heterogeneity

Many have followed Hansen and Singleton (1983) and Hansen and Richard (1987) by imposing rational expectations, letting $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, and defining the stochastic discount factor as the intertemporal marginal rate of substitution

$$m_{t+1} = \frac{\beta u'(c_{t+1})}{u'(c_t)}. \quad (5)$$

The aggregate consumption data have treated (5) and

$$1 = \int_{x_{t+1}} m_{t+1}(x^{t+1}) R_{j,t+1}(x_{t+1}) f(x_{t+1}|x^t) dx_{t+1} \quad (6)$$

badly. One reaction has been to stick to rational expectations but to add backward-looking (see Campbell and Cochrane (1999)) or forward-looking (see Epstein and Zin (1989)) contributions to time t felicity. Another reaction has been to let disparate beliefs contribute to the stochastic discount factor. Hansen and Jagannathan (1991) opened the door to such an

approach when they treated the stochastic discount factor m_{t+1} as an unknown nonnegative random variable and deduced what observed returns $R_{j,t+1}$ and restriction (6) imply about the first and second moments of admissible stochastic discount factors (with incomplete markets, there exist multiple stochastic discount factors). Their idea was that before specifying a particular theory about the utility function and beliefs linking m to real variables like consumption, it could be useful to characterize the mean and standard deviation that an empirically successful m must have. This approach leaves open the possibility that a successful theory of a stochastic discount factor will assign a role to a fluctuating probability ratio $\frac{f_i(x_{t+1}|x^t)}{f(x_{t+1}|x^t)} \neq 1$ even for an economy in which agent i is a single representative agent. The likelihood ratio creates a wedge relative to the Euler equation that has usually been fit in the rational expectations macroeconomic tradition originating in Hansen and Singleton (1983) and Mehra and Prescott (1985). Likelihood ratio wedge approaches have been investigated by Bossaerts (2002, 2004), Hansen (2007), and Hansen and Sargent (2006), among others. The art in Hansen (2007) is to relax rational expectations enough to understand the data better while also retaining the econometric discipline that rational expectations models acquire by economizing on free parameters that characterize agents' beliefs.¹²

3.4 Another empirical reason to allow belief heterogeneity

Applied macroeconomists know that data can be weakly informative about parameters and model features. Ultimately, this is why differences of opinion about how the economy works can persist. The philosophy of Hansen (2007) is to let agents inside a model have views that diverge from the truth in ways about which the data speak slowly and quietly.

¹²Hansen (2007) bears only one new free parameter that governs how much a representative agent's beliefs are exponentially twisted vis-a-vis the data generating mechanism.

4 Rational expectations econometrics

Ideas from rational expectations econometrics will help me to set the stage for some stories and models that feature gaps between an objective distribution and the temporary subjective distributions used by a government that solves a sequence of intelligent design problems. I review econometric methods that allow an outsider to learn about a rational expectations equilibrium and introduce some objects and possibilities that are in play in about models containing agents who are learning an equilibrium.

A rational expectations equilibrium is a joint probability distribution $f(x^t|\theta_o)$ over histories x^t indexed by free parameters $\theta_o \in \Theta$ that describe preferences, technologies, endowments, and information. For reasons that will become clear, I have called the parameter vector θ rather than ρ as in section 2. Rational expectations econometrics tells an econometrician who is outside the model how to learn θ . The econometrician knows only a parametric form for the model and therefore initially knows less about the equilibrium joint probability distribution than nature and the agents inside the model. The econometrician's tools for learning θ are (1) a likelihood function, (2) a time series or panel of observations drawn from the equilibrium distribution, and (3) a Law of Large Numbers, a Central Limit Theorem, and some large deviations theorems that can be used to characterize limits, rates of convergence, and tail behaviors of estimators. With enough data and a correct likelihood function, an econometrician can learn θ_o .

A *likelihood function* is another name for a rational expectations equilibrium evaluated at a particular history is

$$L(\theta|x^t) = f(x^t|\theta) = f(x_t|x^{t-1}; \theta)f(x_{t-1}|x^{t-2}; \theta) \cdots f(x_1|x_0; \theta)f(x_0|\theta). \quad (7)$$

The most confident and ambitious branch of rational expectations econometrics recommends maximizing a likelihood function or combining it with a Bayesian prior $p(\theta)$ to construct a

posterior $p(\theta|x^t)$.¹³ In choosing θ to maximize a likelihood function, a rational expectations econometrician in effect searches for a system of expectations that prompts the forward-looking artificial agents inside the model to make decisions that best fit the data.¹⁴ Taking logs in (7) gives

$$\log L(\theta|x^t) = \ell(x_t|x^{t-1}; \theta) + \ell(x_{t-1}|x^{t-2}; \theta) + \dots + \ell(x_1|x_0; \theta) + \ell(x_0|\theta) \quad (8)$$

where $\ell(x_t|x^{t-1}; \theta) = \log f(x_t|x^{t-1}; \theta)$. Define the score function as $s_t(x^t, \theta) = \frac{\partial \ell(x_t|x^{t-1}, \theta)}{\partial \theta}$. In population, the first-order conditions for maximum likelihood estimation are the conditional moment conditions

$$E[s_t|x^{t-1}] = 0, \quad (9)$$

which imply that the score vector is a martingale difference sequence, the starting point for a theory of statistical inference. By replacing the mathematical expectation E in equation (9) with its sample analogue $T^{-1} \sum_{t=1}^T$, an econometrician finds a θ that allows him to approximate the equilibrium density very well as $T \rightarrow +\infty$.

4.1 Using a misspecified model to estimate a better one

Lucas (1976) warned researchers not to use non-structural models for policy analysis. But the first-order conditions for estimating a good fitting nonstructural model *can* help to make good inferences about parameters of a structural economic model.

Indirect estimation assumes that a researcher wants to estimate a parameter vector ρ of a structural rational expectations model for which (1) analytical difficulties prevent directly evaluating a likelihood function $f(x^t|\rho)$, and (2) computational methods allow simulating time series from $f(x^t|\rho)$ at given vector ρ . See Gourieroux et al. (1993), Smith (1993), and Gallant and Tauchen (1996). Indirect estimation carries along two models, a model

¹³For early applications of this empirical approach, see Sargent (1977), Sargent (1979), Hansen and Sargent (1980), Taylor (1980), and Dagli and Taylor (1984).

¹⁴As the econometrician searches over probability measures indexed by θ , he imputes to the agents inside the system of expectations implied by the θ under consideration.

of economic interest with an intractable likelihood function, and an auxiliary model with a tractable likelihood function that fits the historical data well. The parameters of the economist's model ρ are interpretable in terms of preferences, technologies, and information sets, while the parameters θ of the auxiliary model $f(x^t|\theta)$ are data fitting devices. The idea of Gallant and Tauchen (1996) is first to estimate the auxiliary model by maximum likelihood, then to use the score functions for the auxiliary model and the first-order conditions in equation (9) to define a criterion for a GMM estimator that can be used in conjunction with simulations of the economic model to estimate the parameters ρ . Thus, let the auxiliary model have a log likelihood function given by equation (8) and, for the data sample in hand, compute the maximum likelihood estimate $\hat{\theta}$. For different ρ 's, simulate paths $x_\tau(\rho)$ for $\tau = 1, \dots, N$ from the economic model. Think of using these artificial data to evaluate the score function for the *auxiliary* model $s_t(x^t(\rho), \hat{\theta})$ for each t . Gallant and Tauchen estimate ρ by setting the average score for the auxiliary model¹⁵

$$\frac{1}{N} \sum_{\tau=1}^N s_\tau(x^\tau(\rho), \hat{\theta}) \quad (10)$$

as close to zero as possible when measured by a quadratic form of the type used in GMM. If the auxiliary model fits well, this method gives good estimates of the parameters ρ of the economic model. In particular, the indirect estimator is as efficient as maximum likelihood in the ideal case that the economic and auxiliary models are observationally equivalent.

4.2 A troublesome possibility

This ideal case raises the following question: what happens if macroeconomic policy makers incorrectly use what from nature's point of view is actually an auxiliary model? Data give the government no indication that it should abandon its model. Nevertheless, the government can make major policy design mistakes because it misunderstands the consequences of poli-

¹⁵This description fits their Case 2.

cies that it has not chosen.¹⁶ The possibility that the government uses what, unbeknownst to it, is just an auxiliary model, not a structural one, sets the stage for the self-confirming equilibria that play an important role in the adaptive learning models of the following section and in the stories to be told in sections 6 and 7.

5 Adaptive learning models and their limits

Section 3 described Blume and Easley’s demonstration that a survival argument for equating objective and subjective distributions falls short in many economies. This section takes up where that discussion left off by describing transient and limiting outcomes in models in which agents make decisions by using statistical models that at least temporarily are misspecified. I summarize findings from a literature that studies systems of agents who use forward-looking decision algorithms based on temporary models that they update using recursive least squares algorithms (see Marcet and Sargent (1989a), Evans and Honkapohja (1999, 2001), Woodford (1990), and Fudenberg and Levine (1998)).¹⁷ These adaptive systems can have limiting outcomes in which objective and subjective distributions are equal over commonly observed events, but not over rarely observed events. That causes problems for intelligent macroeconomic policy design. I shall use examples of such adaptive systems to tell some stories in section 7. It is useful to begin by defining some population objects that suppose that agents have finished learning.

5.1 Self-confirming equilibrium

A true data generating process and an approximating model, respectively, are

$$f(y^\infty, v^\infty | \rho) \text{ and } f(y^\infty, v^\infty | \theta). \tag{11}$$

¹⁶See Lucas (1976), Sargent (1999, ch. 7), and Fudenberg and Levine (2007).

¹⁷Appendix A describes a related literature on learning in games.

A decision maker has preferences ordered by

$$\int U(y^\infty, v^\infty) f(y^\infty, v^\infty | \theta) d(y^\infty, v^\infty) \quad (12)$$

and chooses a history-dependent plan

$$v_t = h_t(x^t | \theta), \quad t \geq 0 \quad (13)$$

that maximizes (12). This gives rise to the sequence of decisions $v(h|\theta)^\infty$. The difference between this choice problem and the canonical intelligent design problem in section 2 is the presence of the approximating model $f(y^\infty, v^\infty | \theta)$ in (12) rather than the true model that appeared in (2). I call maximizing (12) a ‘‘Phelps problem’’ in honor of a government control problem that Phelps (1967) solved and that will play an important role in section 7.

Definition 5.1. *A self-confirming equilibrium (SCE) is a parameter vector θ_o for the approximating model that satisfies the data-matching conditions*

$$f(y^\infty, v(h|\theta_o)^\infty | \theta_o) = f(y^\infty, v(h|\theta_o)^\infty | \rho). \quad (14)$$

An SCE builds in, first, optimization of (12) given beliefs indexed by θ_o , and, second, a $\theta = \theta_o$ that satisfies the data matching conditions (14). Data matching prevails for events that occur under the equilibrium policy $v(h|\theta_o)^\infty$, but it is possible that

$$f(y^\infty, v^\infty | \theta_o) \neq f(y^\infty, v^\infty | \rho) \quad (15)$$

for $v^\infty \neq v(h|\theta)^\infty$. In an SCE, the approximating model is observationally equivalent with the true model for events that occur under the policy implied by equilibrium decisions, but not necessarily under other policies.

5.2 Learning converges to an SCE

An SCE is a possible limit point an adaptive learning process. Bray and Kreps (1987) distinguish between learning *about* an equilibrium and learning *within* an equilibrium.¹⁸ By saying *about* and not *within*, Bray and Kreps emphasize that the challenge is to analyze how a system of agents can come to learn an endogenous objective distribution by using adaptive algorithms that do not simply apply Bayes' law to a correct probability model.¹⁹ We can't appeal to the same econometrics that lets a rational expectations econometrician learn an equilibrium because an econometrician is *outside* the model and his learning is a side-show that does not affect the data generating mechanism. It is different when the people learning about an equilibrium are inside the model. Their learning affects decisions and alters the distribution of endogenous variables over time, making them aim at moving targets.

Suppose that an adaptive learner begins with an initial estimate $\hat{\theta}_0$ at time 0 and uses a recursive least squares learning algorithm

$$\hat{\theta}_{t+1} - \hat{\theta}_t = e_{\theta}(\hat{\theta}_t, y^t, v^t, t). \quad (16)$$

As in the models of learning in games of Foster and Young (2003) and Young (2004, ch. 8), we assume that decision makers mistakenly regard their time t model indexed by $\hat{\theta}_t$ as permanent and form the sequence of decisions²⁰

¹⁸A difficult challenge in the machine learning literature is to construct an adaptive algorithm that learns dynamic programming. For a recent significant advance based on the application of the adjoint of a resolvent operator and a law of large numbers, see Meyn (2007, ch. 11).

¹⁹Bray and Kreps's 'about' versus 'within' tension also pertains to Bayesian theories of convergence to Nash equilibria. Marimon (1997) said that a Bayesian knows the truth from the beginning. Young (2004) pointed out that the absolute continuity assumption underlying the beautiful convergence result of Kalai and Lehrer (1993, 1994) requires that players have substantial prior knowledge of their opponents' strategies. Young is skeptical that Kalai and Lehrer have answered the question "... can one identify priors [over opponents strategies] whose support is wide enough to capture the strategies that one's (rational) opponents are actually using, without assuming away the uncertainty inherent in the situation?" Young (2004, p. 95)

²⁰Cho and Kasa (2006) create a model structure closer to the vision of Foster and Young (2003). In particular, Cho and Kasa's model has the following structure: (1) one or more decision makers take actions at time t by solving a dynamic programming problem based on a possibly misspecified time t model, (2) the actions of some of those decision makers influence the data-generating process; (3) the decision maker shows that he is aware of possible misspecifications of his model by trying to detect them with an econometric specification test, (4) if the specification test rejects the model, the decision maker selects an improved

$$\hat{v}(h)_t = h_t(x^t|\hat{\theta}_t) \tag{17}$$

where $h_t(x^t|\theta)$ is the same function (13) that solves the original Phelps problem (12) under the model $f(y^\infty, v^\infty|\theta)$. Under this scheme for making decisions, the joint density of $(y^\infty, v^\infty, \hat{\theta}^\infty)$ is

$$f(y^\infty, \hat{v}(h)^\infty, \hat{\theta}^\infty|\rho). \tag{18}$$

The learning literature states restrictions on the estimator e and the densities $f(\cdot|\theta)$ and $f(\cdot|\rho)$ that imply that

$$\hat{\theta}_t \rightarrow \theta_o, \tag{19}$$

where convergence can be either almost surely or in distribution, depending on details of the estimator e in (16).²¹

5.3 Uses of adaptive learning models in macroeconomics

One important use of adaptive models in macroeconomics has been to select among multiple rational expectations equilibria (see Evans and Honkapohja (2001) for many useful examples). Another has been to choose among alternative representations of policy rules from Ramsey problems, a subset of which are stable under adaptive learning (see Evans and Honkapohja (2003)). Another has been to improve the fits of models of asset pricing and big inflations by positing gaps between the objective density and asset holders' subjective densities (e.g., Adam et al. (2006) and Marcet and Nicolini (2003)). In the remainder of this paper, I focus on yet another application, namely, to situations in which a government solves an intelligent

model, while (5) if the current model is not rejected, the decision maker formulates policy using the model under the assumption (used to formulate the dynamic programming problem) that he will retain this model forever. Cho and Kasa define useful mathematical senses in which the same stochastic approximation and large deviations results that pertain to a least-squares learning setup also describe the outcomes of their model-validation setup.

²¹For example, so-called 'constant gain' algorithms give rise to convergence in distribution, while estimators whose gains diminish at the proper rates converge almost surely. See Williams (2004). A few papers have studied rates of convergence. There are examples in which convergence occurs at a \sqrt{T} rate, but also examples where convergence occurs markedly more slowly.

design problem by using a misspecified model.

5.4 REE or SCE?

Some builders of adaptive models have specified an approximating model to equal a true one, meaning that there exists a value θ_o for which $f(y^\infty, v^\infty|\rho) = f(y^\infty, v^\infty|\theta_o)$ for *all* plans v^∞ , not just equilibrium ones. This specification prevails in adaptive models in which least squares learning schemes converge to rational expectations equilibria, like those used by Woodford (1990) and Marcet and Sargent (1989b). When $f(y^\infty, v^\infty|\rho) \neq f(y^\infty, v^\infty|\theta_o)$ for some choices of v , the most that can be hoped for is convergence to an SCE.²²

5.5 SCE-REE gaps and policy design

Why is a gap between a rational expectations equilibrium and a self-confirming equilibrium important for a macroeconomist? Macroeconomists build models with many small agents and some large agents called governments. Small agents can take aggregate laws of motion as given within a recursive competitive equilibrium. It doesn't matter to a small agent that his views may be incorrect views off the equilibrium path.

But it can matter very much when a large agent like a government has incorrect views off an equilibrium path because in designing its policy we suppose that a government solves a Ramsey problem in which it contemplates the consequences of off-equilibrium path experiments. Wrong views about off-equilibrium path events shape government policy and the equilibrium path.

²²Sargent (1999, ch. 6) works with a weaker notion of an SCE that Branch and Evans (2005, 2006) call a misspecification equilibrium. Branch and Evans construct misspecification equilibria in which agents i and j have different models parameterized, say, by θ_i and θ_j , and in which $f(x^t|\theta_i) \neq f(x^t|\theta_j) \neq f(x^t|\rho)$, where again ρ parameterizes the data generating mechanism. A misspecification equilibrium imposes moment conditions on agents' approximating models that imply parameters θ_i that give equal minimum mean square error forecast errors $E_{\theta_j}[(x_{t+1} - E_{\theta_j}(x_{t+1}|x^t))(x_{t+1} - E_{\theta_j}(x_{t+1}|x^t))']$ for all surviving models. Branch and Evans model equilibria in which beliefs and forecasts are heterogeneous across agents, though they have equal mean squared errors. They provide conditions under which recursive least squares learning algorithms converge to a subset of the possible misspecification equilibria. The models of Brock and Hommes (1997) and Brock and de Fontnouvelle (2000) are early versions of misspecification equilibria.

As my laboratory, I focus on more or less formal descriptions of some processes that have taught central bankers. Section 6 is a literary account of a millenium of monetary theory and experiments. I see self-confirming equilibria lurking in this story, for example, in the way a commodity money standard hid the quantity theory of money from empiricists by suppressing the variation in price levels and money supplies needed to identify it. That story takes us to the threshold of the 20th century experiment with fiat currency. Section 7 fast-forwards to the 1960s and 1970s and uses statistical models to describe how the U.S. monetary authorities struggled to understand inflation-unemployment dynamics as they sought to meet their dual mandate of promoting high output growth and low inflation.

6 Learning monetary policy over a millenium

Central bankers are preoccupied with nominal anchors.

6.1 From commodity to token to fiat money

Appendix B describes a 700 year process of theorizing and experimenting that transformed the European commodity money system from one with *many* nominal anchors – mint-melt price pairs (i.e., gold or silver points) for full bodied coins of all denominations – to a *one* nominal anchor system that retained gold points for only one standard full bodied coin and used government-issued convertible token coins and notes for other denominations.²³ After another 100 years, governments abolished the gold points for the standard coin too, leaving the nominal anchor to be the monetary authorities' good intentions and their knowledge of the quantity theory of money. The appendix notes how a commodity money system concealed the quantity theory of money because the gold and silver points made the price level a low variance, small trend exogenous variable by making the money supply into a low variance, small trend endogenous variable. Eventually, some atypical policy experiments generated

²³Fetter (1978, p. 16) and Friedman (1991, pp. 150-151) discuss how concerns about small denomination coins shaped the gold standard.

data with sufficient variance in price levels and money supplies to reveal the quantity theory to empiricists, a theory that eventually induced monetary experts like Keynes and Fisher to advocate a well-managed fiat system.

6.2 Two threats to a well managed fiat money system

Friedman (1991, pp. 249-252) said that our present fiat money system is historically unprecedented and repeated the warning of Fisher (1926, p.131) that “Irredeemable paper money has almost invariably proved a curse to the country employing it” because two obstacles obstruct the path to managing a fiat currency well: (1) political pressures to use fiat money to finance the government expenditures, and (2) temptations to exploit a Phillips curve (Friedman (1991, p. 207)). Empirical learning models have been used to interpret monetary authorities’ struggles to understand and avoid these obstacles. Marcet and Nicolini (2003) and Sargent et al. (2006a) constructed adaptive models that focus on (1) and feature private agents’ learning. Those papers both select among rational equilibria and modify their outcomes enough to fit data from big inflations in Latin America. In the remainder of this paper, I focus on statistical models that feature monetary authorities’ struggles with Friedman’s obstacle 2.

7 Learning inflation-unemployment dynamics

This section describes three stories about how the U.S. monetary authorities learned about inflation-unemployment dynamics after World War II. These stories accept that a monetary authority can control inflation if it wants.²⁴ Then why did the U.S. monetary authority allow inflation to rise in the late 1960s and 1970s, and why did it bring inflation down in the 1980s and 1990s? If we assume that its purposes did not change, and that it always disliked inflation and unemployment, then it is natural to focus on changes over time in

²⁴Appendix C discusses a monetary policy rules literature that focuses on the gap between the monetary authorities’ instruments and inflation outcomes.

the monetary authority's understanding of inflation-unemployment dynamics. I'll describe stories associated with three empirical models that feature either temporary or permanent discrepancies between a government's model and a true data generating mechanism, and a government that each period solves a misspecified intelligent design problem and revises its parameter estimates to align them with new data.²⁵

It is natural to impute popular contemporary models to the government. The 'revisionist history' of the U.S. Phillips curve by King and Watson (1994) provides a good source for these. King and Watson studied how econometric directions of fit (i.e., should you regress inflation on unemployment or unemployment on inflation?) affect government decisions. To make contact with studies from the 1970s, King and Watson call inflation on unemployment the Keynesian direction of fit and unemployment on inflation the classical direction.²⁶ I impute stripped down versions of more completely articulated models to the government.²⁷ These simplified models can capture the substantially different operating characteristics that drive our stories.

The three stories have monetary authorities solve adaptive intelligent design problems with misspecified models that induce them to make decisions that are influenced by erroneous views about taking decisions other than the ones they actually take. The stories differ in the nature of those misunderstandings. In the first story, the monetary authority's misspecified model misses a chain of causation linking its decisions first to the private sector's expectations of inflation and then to the position of an unemployment-inflation trade-off. In the second story, there exists a parameter vector $\theta_o = \rho$ that aligns the monetary authority's model with the data generating mechanism on *and* off the chosen stochastic monetary policy path, but

²⁵For testimony that policy authorities in the U.S. are concerned about related issues, see Bernanke (2007) and Mishkin (2007). See Evans and Honkapohja (2003), Orphanides and Williams (2005, 2007), and Bullard and Mitra (2007) for applications of models of this type to evaluating the stability and performance of alternative monetary policy rules. See Cogley (2005) and Piazzesi and Schneider (2007) for applications to the yield curve.

²⁶Sargent (1999, ch. 7) described how those specification decisions can affect self-confirming equilibrium outcomes.

²⁷Some economists nowadays use the slang 'reduced form' to refer to such incompletely articulated models. I prefer to reserve 'reduced form' for its original meaning in Cowles commission econometrics, namely, a particular statistical representation associated with a well articulated structural model.

except in the limit as $t \rightarrow \infty$, the government's temporary misestimates $\hat{\theta}_t$ of θ_o induce it to misunderstand the consequences of policies that it chooses not to implement. In the third model, the government mixes across submodels with operating characteristics that give very different readings about the consequences of following a no-feedback low inflation policy.

7.1 The (temporary) conquest of U.S. inflation

This story is about generating sufficient variation in the data to allow a government's misspecified model to detect that there is no exploitable trade-off between inflation and unemployment. The story is that the only way the government's model let's it discover that there truly is no *exploitable* tradeoff is for it falsely to infer that there is no tradeoff *whatsoever*. That imperfection dooms any stabilization of inflation to be at best temporary.

This story uses specifications $f(y^\infty, v^\infty | \rho) \neq f(y^\infty, v^\infty | \theta)$ to capture how a monetary authority misrepresents how its decisions impact on private agents' expectations about inflation and, therefore, on the joint distribution of unemployment and inflation. I illustrate the forces at work with the following simplified version of the type of model that Sims (1988), Cho et al. (2002), and Sargent and Williams (2005) have studied and that Chung (1990), Sargent (1999), and Sargent et al. (2006b) have fit to U.S. data. The true model is

$$U = \rho_0 - \rho_1 \sigma_2 w_2 + \sigma_1 w_1, \tag{20}$$

$$\pi = v + \sigma_2 w_2 \tag{21}$$

where U is the unemployment rate, π is the rate of inflation, v is the systematic part of the inflation rate chosen by the monetary authority, w is a 2×1 Gaussian random vector with mean zero and identity covariance, and $\rho_0 > 0, \rho_1 > 0$, where ρ_0 is the natural rate of unemployment and ρ_1 is the slope of the Phillips curve. Through equation (20), which is the aggregate supply curve proposed by Lucas (1973), the model captures a rational expectations version of the natural unemployment rate hypothesis that asserts that the

systematic component of inflation v does not affect the distribution of the unemployment rate conditional on v . For an objective function $-E(U^2 + \pi^2)$, the government's best policy under the true model is $v = 0$.

The government's approximating model denies the natural rate hypothesis by asserting that v affects the probability distribution of U according to:

$$U = \theta_0 + \theta_1(v + \sigma_2\tilde{w}_2) + \sigma_1\tilde{w}_1 \quad (22)$$

$$\pi = v + \sigma_2\tilde{w}_2, \quad (23)$$

where the random vector \tilde{w} has the same distribution as w . Under the true model and the timing protocol that the government chooses target inflation before the private sector sets its expectation of inflation, the government's best policy is $v = 0$. However, under the approximating model (22)-(23), the government's best policy is

$$v = h(\theta) = \frac{-\theta_1\theta_0}{1 + \theta_1^2}. \quad (24)$$

There exists a self-confirming equilibrium in which

$$(\theta_0)_o = \rho_0 - \rho_1 h(\theta_o) \quad (25)$$

$$(\theta_1)_o = -\rho_1. \quad (26)$$

The self-confirming equilibrium equals the time-consistent equilibrium of Kydland and Prescott (1977).²⁸ An adaptive government's estimates $\hat{\theta}_t$ converge to the self-confirming equilibrium vector θ_o , and the systematic part of inflation converges to $v = h(\theta_o)$.

The data-matching restriction (25) pinpoints how the government mistakenly ignores the effect of its policy choice v , which equals the public's expected rate of inflation, on the

²⁸The same suboptimal outcome occurs, but for a different reason than the inferior timing protocol isolated by Kydland and Prescott (1977). Here the source of sub optimality of the government's choice originates in its misunderstanding of the economic structure. The timing protocol is such that if the government knew the correct model, it would attain Ramsey outcomes, in the language of Stokey (1989).

position of the Phillips curve. If v were generated randomly with enough variance, then even though it fits the wrong model, the government would estimate a Phillips curve slope θ_1 of approximately zero and according to (24) would set v approximately to its optimal value of 0 under the true model. But within an SCE, v doesn't vary enough for the government to estimate a θ_1 close enough to zero for that to happen. Furthermore, the outcome that $\hat{\theta}_t \rightarrow \theta_o$ means that the variation of v_t that occurs along transient paths converging to an SCE is insufficient to allow the government's model to approximate the data in a way that tells it to implement the optimal policy under the true model.

However, that is not the end of the story because the adaptive model's endogenous stochastic dynamics occasionally make v vary enough for the government to discover a (too strong) version of the natural rate hypothesis, too strong because it mistakenly asserts that there is no tradeoff whatsoever between π and U . The adaptive system is destined to experience recurrent episodes in which 'a most likely unlikely' sequence of w 's lowers the unconditional correlation between U and π , which in turn prompts the government's estimates $\hat{\theta}_t$ to induce the government to push v_t downward from its self-confirming value.²⁹ This process generates data that further weakens the unconditional correlation between inflation and unemployment and drives v even lower. The ultimate destination of this 'escape' from a self-confirming equilibrium is that the government estimates that θ_1 is 0, prompting it to set v_t at the optimal value 0. These escapes are more likely when the government's estimator (16) discounts past data more heavily, for example, by using a so-called constant gain algorithm. An escape is temporary because the mean dynamics that drive the system toward the SCE vector θ_o are bound to reassert themselves and push inflation back toward the suboptimal SCE value of $h(\theta_o)$. If this is a good parable for the Volcker-Greenspan stabilization, we should be worried.

²⁹See Cho et al. (2002) for an explanation of 'most likely unlikely' sequences of shocks.

7.1.1 Details

Simulations of Sims (1988) generated sample paths that seemed promising for explaining a Volcker-like stabilization prompted by the government's being able to learn a good enough version of the natural rate hypothesis. However, formal econometric attempts to implement the model by Chung (1990) and Sargent (1999) failed to fit the U.S. data well, mainly because the government's adaptive algorithm catches on to the adverse shifts in Phillips curve so quickly in the early 1970s that the Phelps problem tells the government to stabilize inflation much earlier than actually occurred. Sargent et al. (2006b) replaced the constant gain algorithm used in the earlier models with the Bayesian updating procedure implied by a drifting coefficients model with a covariance matrix V for the innovations in the drifts to the coefficients. By estimating V along with the parameters of nature's model by maximum likelihood, they reverse engineered a drifting set of government beliefs that, when put into the Phelps problem each period, produce a sequence of first period Phelps policy recommendations that do a good job of matching the actual inflation data. The estimated V makes the intercept in the Fed's model quite volatile and thus makes contact with descriptions of Arthur Burns's Fed, which according to Hetzel (1998), attributed much of the inflation of the 1970s to special factors akin to dummy variables that capture intercept drift in regressions. The maximum likelihood estimate of V is large and conveys the image of a government that expects coefficients to drift so much that it discounts past data heavily. The model's conjuring up a Fed that over fits its models to recent data is food for thought for Fed watchers. The synthesized government beliefs succeed in rationalizing inflation *ex post* as a response to these government beliefs, and the beliefs themselves do a good job of forecasting inflation, thus capturing what seems to have been a remarkably good record of inflation forecasting by the Fed (see Bernanke (2007)).³⁰

³⁰But relative to available alternatives, the imputed beliefs do a poor job of forecasting unemployment, a deficiency of the model that hints that the reverse-engineering exercise may be imputing unrealistic views about *joint* inflation-unemployment dynamics to the Phelps problem in order to rationalize observed inflation outcomes. Carboni and Ellison (2007) also note this deficiency in the model and make substantial progress in repairing it while also reducing the estimated V by conditioning estimates on greenbook forecasts.

7.1.2 Best of all possible worlds?

In this story, policy choices recurrently revisit ones that are optimal under the correct model, but they don't stay there because the mean dynamics attracting them to the suboptimal SCE are destined to reassert themselves. Thus, this story at best only temporarily supports the optimism expressed by Sims (1988) that the government's misspecified model can approximate the lack of an exploitable $U - \pi$ tradeoff well enough to induce the government to do what would be the right thing if it actually knew the true model. For the misspecified model to reveal the lack of a tradeoff, the government has to induce adequate variation in inflation, which it doesn't within an SCE. So the first story stops short of being one in which evolution converges to 'the best of all possible worlds.' However, a more optimistic outcome prevails in our next story, which hard wires things to make the government's misunderstandings of off-equilibrium-path choices eventually vanish.

7.2 A Keynesian account

The previous story is about how the troublesome possibility raised in subsection 4.2 plays out. The model of Primiceri (2006) envisions a world in which that possibility is off the table because $f(y^\infty, v^\infty | \rho) = f(y^\infty, v^\infty | \theta_o)$ for all v^∞ and an SCE equals an REE. All of the action in Primiceri's model comes from calibrating an initial $\hat{\theta}_0 \neq \theta_o$ that leads to a stochastic path that converges to an SCE presided over by Greenspan and whose transient dynamics mimic the post WWII U.S.

Primiceri's model has a time invariant true data generating model featuring (i) an expectations augmented Phillips curve, (ii) a Cagan (1956)-Friedman (1956) adaptive expectations scheme that describes how the public forms the expectations of inflation that appear in (i)³¹; and (iii) an aggregate demand equation that describes how the time t value of an

³¹Primiceri assumes that a fraction of agents form expectations this way, while the remainder have rational expectations. Primiceri's specification imposes that the sum of weights on lagged inflation equals unity. Lucas (1972a) and Sargent (1971) argued that, except in a special case, the sum of the weights on lagged inflation being one is not a valid characterization of the natural rate hypothesis. See King and Watson (1994) and Sargent (1999).

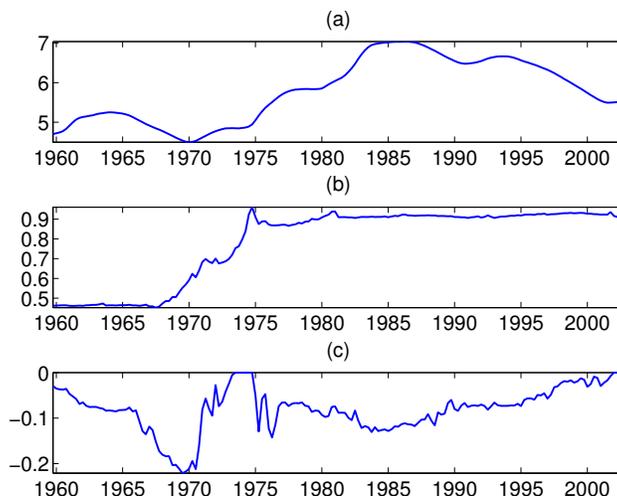


Figure 1: Evolution of policy-maker's beliefs about: (a) the natural rate of unemployment; (b) the persistence of inflation in the Phillips curve; and (c) the slope of the Phillips curve in King and Watson's Keynesian direction. (Primiceri 2006, p. 882)

uninterpreted government policy instrument v_t affects current and future gaps between the unemployment rate u_t and a natural rate of unemployment u_t^N .³² The model neatly allows the government's misperception of the natural rate to influence policy, as advocated by Orphanides (2002, 2003). It also allows other potentially important government misperceptions to influence policy. Primiceri shows that the lower is the sum of the weights on lagged inflation in the expectational Phillips curve, and therefore the less persistent is inflation under a passive government policy, the less counterinflationary is the policy that emerges from the Phelps problem. A lower estimated persistence of inflation indicates to the government that mean reverting inflation will evaporate soon enough on its own. Coefficients that measure the strength of the feedback from unemployment to inflation also influence how actively counterinflationary is the policy called for by the time t Phelps problem.

Primiceri calibrates initial government beliefs by using data between 1948 and 1960 to estimate the model's parameters.³³ These calibrated beliefs feature a level of persistence

³²Feature (ii) of Primiceri's model embraces a Keynesian spirit of assuming that the authority influences output directly through the aggregate demand function, then inflation indirectly through the expectations-augmented Phillips curve. Contrast this with the classical specification adopted by Sims (1988), Chung (1990), Sargent (1999), Cho et al. (2002), and Sargent et al. (2006b).

³³Primiceri calibrates initial beliefs for the government about the value of the natural rate of unemployment and the coefficients in a reduced-form Phillips curve. The reduced form is derived by substituting the adaptive

of inflation in the Phillips curve that is much lower than what prevails in the estimated model's self-confirming equilibrium. In addition to these initial conditions, Primiceri sets two constant gain parameters, one for the natural rate, another for all other coefficients in the government's beliefs. These calibrated objects, the data, and the parameters of the structural relations pin down the government's beliefs. Primiceri uses maximum likelihood to estimate parameters appearing in the government's objective function and the time-invariant structural equations.

Primiceri accounts for the acceleration of inflation in the 1960s and 1970s, then the fall in the 1980s in terms of the government's initial underestimates of the natural unemployment rate as well as a temporal pattern of underestimates of the persistence of inflation and overestimates of the costs of disinflation coming from its estimated inflation-unemployment tradeoff. Figure 1 reproduces Primiceri's figure II, which shows his estimates of the evolution of the Fed's estimates of the natural rate of unemployment, the persistence inflation, and the slope of the Phillips curve. The Phelps problem attributes the acceleration of inflation to the monetary authority's initial underestimates of both the natural rate and the persistence of inflation. After inflation had risen, the Phelps problem attributes the monetary authority's reluctance to deflate to its overestimation of the costs of disinflation as captured by the slope of the Phillips curve. We will return to this point in subsection 7.3, where we link it to the conceptual issues about direction of fit raised by King and Watson (1994).³⁴

Under-estimates of the natural unemployment rate and over-estimates of the sacrifice ratio are connected. When the Fed under-estimates the natural rate and over-estimates the unemployment gap, it over-predicts the amount of disinflation. That causes it to revise its estimate of the slope of the Phillips curve towards zero. Thus, Orphanides's story about the consequences of misestimating the natural rate of unemployment complements Primiceri's

expectations scheme into the expectations augmented Phillips curve.

³⁴Among many interesting features of Primiceri's results are his estimate of k , a parameter in the government objective function that allows Primiceri to evaluate the government's temptation to deviate from the natural rate (he finds that the temptation is small) and the time series that he extracts for v_t , which tracks a real interest rate very well after 1980.

story about sacrifice ratio pessimism.

7.3 An eclectic account

The models in the previous two sections take stands on what both the true and the government's approximating models are. Cogley and Sargent (2005) performed an exercise that did not require them to specify a true data generating mechanism, it being enough for their purposes to consult the empirical distribution. But the government's views about policy choices not made play a key role. The government's model $f(y^\infty, v^\infty | \theta)$ is a mixture of three submodels and $\hat{\theta}_t$ includes Bayesian posterior probabilities that the government uses to mix the three submodels.

A government entertains three models that Cogley and Sargent chose to capture prominent specifications from the literature about U.S. unemployment-inflation dynamics described by King and Watson (1994). The models are (1) a Samuelson-Solow Phillips curve with King and Watson's Keynesian direction of fit, a model that implies a long-run exploitable trade-off between inflation and unemployment; (2) a Solow-Tobin model with a Keynesian direction of fit that features a short-run but no long-run trade-off between inflation and unemployment (albeit according to what Lucas (1972a) and Sargent (1971) claimed was a dodgy notion of long-run); and (3) a Lucas specification with a classical direction of fit that implies no exploitable trade-off between inflation and unemployment. If probability one is put on the Lucas model, the Phelps problem gives the trivial solution that the government should set the systematic part of inflation equal to zero. If probability one is put on either of the other models, the systematic part of inflation is a linear function of the state variables appearing in those exploitable dynamic Phillips curves. The government puts positive probability on all three models, so the Phelps problem brokers a compromise among the recommendations of the three models. But what kind of compromise? It depends on submodel probabilities times value functions.

The government starts with a prior with non-zero weights on all three models in 1960,

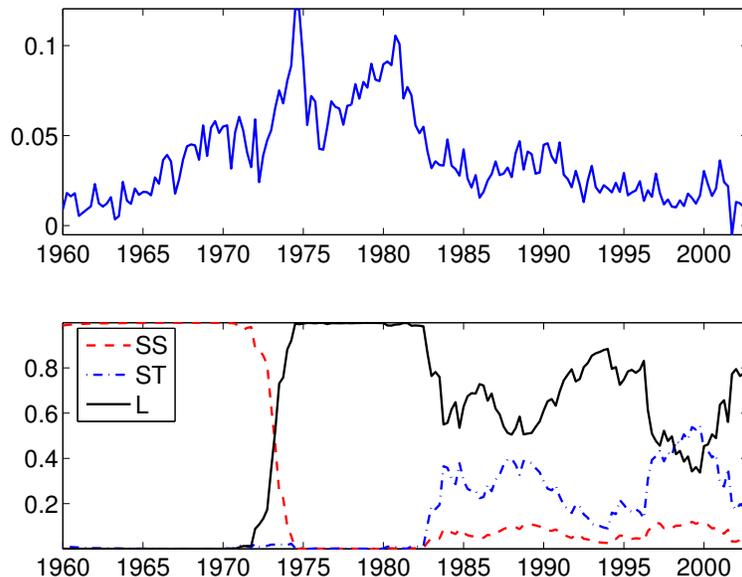


Figure 2: Top panel: CPI inflation. Bottom panel: Bayesian posterior model weights on the Samuelson-Solow (SS), Solow-Tobin (ST), and Lucas (L) models.

estimates each sub model using Bayesian methods, and updates its prior over the three sub models.³⁵ In each period, the government solves a Phelps problem that penalizes inflation and unemployment and that uses its time t submodel probabilities to average over its time t estimates of its three submodels. Cogley and Sargent put prior probabilities in 1960 of .98 on the Samuelson-Solow model and .01 each on the Solow-Tobin and the Lucas model. We put those low prior probabilities on the Lucas and Solow-Tobin models because only the Samuelson-Solow model existed in 1960.³⁶ Putting U.S. inflation-unemployment data into this machine, Cogley and Sargent computed time series of (1) the posterior model weights $\alpha_{i,t}$, and (2) the systematic part of the inflation rate set by the government in the Phelps problem.

Figures 2 and 3 taken from Cogley and Sargent (2005) frame the following puzzles. By

³⁵See Cogley et al. (2007b) for a related setup that has only two submodels, each of which has known coefficients, and in which the government designs purposeful experiments because it includes the submodel probabilities in the state vector. By way of contrast, the model being discussed in the text has the government making decisions as if its temporary mixture of models will prevail forever and therefore excludes purposeful experimentation. Cogley et al. (2007a) study purposeful experimentation when a government trusts neither its submodels nor its Bayesian posterior over submodels.

³⁶We have to put positive probabilities on the yet-to-be invented models in 1960 in order to launch our story. Foster and Young (2003) introduce new models randomly.

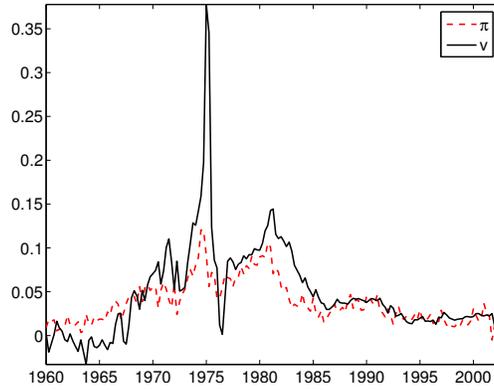


Figure 3: CPI inflation and recommendation from Phelps problem.

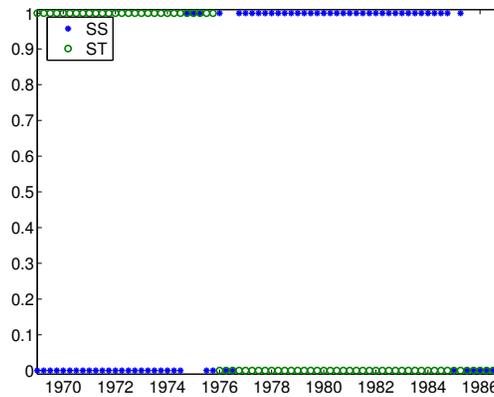


Figure 4: Loss from SS model (*) and ST model (o). When value is 1, it denotes infinite loss under Lucas zero inflation policy. When value is 0, it denotes finite loss.

the early 1970s, the data had moved the government's prior to put probability approaching 1 on the Lucas model (see figure 2). Since the Lucas model recommends zero inflation, why nevertheless was actual inflation so high and variable in the 1970s? And why was the systematic part of inflation that emerges from the Phelps problem (see figure 3) even higher and more variable? Why does the Phelps planner seem to disregard the recommendations of the Lucas model despite its high posterior probability?

The answer is to be found in what the Samuelson-Solow and Solow-Tobin models say would happen if the Lucas zero-target-inflation policy were to be adopted, as summarized in figure 4. The Phelps problem weights the submodel posterior probabilities against losses associated with various off-taken-path recommendations. In the early 1970s, their Keynesian direction of fit moved the coefficients in those submodels in ways that pointed to very high sacrifice ratios. Despite their low posterior probabilities, those models implied very high expected discounted losses if the Lucas policy recommendation were to be implemented immediately. In contrast, the high-probability Lucas model implied less adverse consequences if the recommendations of the Samuelson-Solow or Solow-Tobin models were allowed to prevail. So the Cogley and Sargent story is that the Lucas models policy recommendation did not prevail in the 1970s because there remained a low probability that it would be disastrous. In order for a low-inflation recommendation to emerge from the Phelps problem, it was necessary that the estimated coefficients in the Samuelson-Solow and Solow-Tobin models adjust in ways that would moderate the consequences of a low-inflation policy. That happened by the mid 1980s.³⁷

The direction-of-fit issue discussed by King and Watson (1994) helps understand how

³⁷The data also indicate that Bayes' law sponsors comebacks for the Samuelson-Solow and Solow-Tobin models in the 1980s and 1990s. One reaction that a true believer in the Lucas model might have is that Bayes' law is just too forgiving in still putting positive probability on those other models after the early 1970s data had come in, and that the inflation problem of the 1970s would have been solved by driving a stake through those other models. But no one has the authority to drive stakes, and models with operating characteristics much like those two survive today. The dispute between the fallacious (according to Friedman and Schwartz (1963, p. 191)) real bills doctrine and the quantity theory of money is mottled with repeated episodes having one of these doctrines being disposed of in favor of the other, then the other making a comeback. The real bills doctrine rides high in times like these when the Fed pegs a short term interest rate.

some of Primiceri's results relate to Cogley and Sargent's. Both models emphasize how monetary policy changed as the authorities updated their estimates, and Primiceri also attributes the inflation of the 1970s to the high perceived sacrifice ratio that Keynesian Phillips curve models presented to policy makers. But Primiceri assumes that the Fed relied exclusively on a version of the Solow-Tobin model and is silent about why the Fed disregarded the recommendations of the Lucas model. The central element of his story – the high perceived cost of disinflation or sacrifice ratio – is not a robust feature across the three submodels used by Cogley and Sargent because it depends critically on the direction of fit, as documented by Cogley and Sargent (2005, p. 546-547). The sacrifice ratios differ so much across submodels because of how the submodels interpret the diminished, near-zero contemporaneous covariance between inflation and unemployment that had emerged by the mid 1970s. In a Keynesian Phillips curve, this diminished covariance flattens the short-term tradeoff, making the authorities believe that a long spell of high unemployment would be needed to bring inflation down, prompting Keynesian modelers to be less inclined to disinflate. But for a classical Phillips curve, the shift toward a zero covariance steepens the short-term tradeoff, making the authorities believe that inflation could be reduced at less cost in terms of higher unemployment. Thus, a classically-oriented policy maker was more inclined to disinflate.

8 Concluding remarks

Sawhill (1995) also began by citing Keynes (1936) but disagreed with him by lamenting that politics, not good economics, drives decision making in Washington. In monetary economics today, the relationship between Washington policy makers and economists is so much easier than it was for Sawhill that I think it acceptable to summarize the case for adaptive models by quoting one “academic scribbler” who is now a monetary authority:

The traditional rational-expectations model of inflation and inflation expecta-

tions has been a useful workhorse for thinking about issues of credibility and institutional design, but, to my mind, it is less helpful for thinking about economies in which (1) the structure of the economy is constantly evolving in ways that are imperfectly understood by both the public and policymakers and (2) the policymakers' objective function is not fully known by private agents. In particular, together with the assumption that the central bank's objective function is fixed and known to the public, the traditional rational-expectations approach implies that the public has firm knowledge of the long-run equilibrium inflation rate; consequently, their long-run inflation expectations do not vary over time in response to new information.

Although variations in the extent to which inflation expectations are anchored are not easily handled in a traditional rational-expectations framework, they seem to fit quite naturally into the burgeoning literature on learning in macroeconomics. The premise of this literature is that people do not have full information about the economy or about the objectives of the central bank, but they instead must make statistical inferences about the unknown parameters governing the evolution of the economy. Bernanke (2007)

It is easy to agree with a warning by Sims (1980) that leaving the rational expectations equilibrium concept sends us into a “wilderness” because there is such a bewildering variety of ways to put discrepancies between objective and subjective distributions.³⁸ But the adaptive models described in this paper are very cautious departures from rational expectations theories and from rational expectations econometrics. The timidity of the departure from rational expectations reflects a desire to retain much of the discipline of rational expectations econometrics. I have focused some of the things that can happen when a government solves an intelligent design problem while using a misspecified model. I view the very simple statistical models in section 7 as a parable for the situation that we are always in, namely,

³⁸There is an infinite number of ways to be wrong, but only one way to be correct.

that our probability models are misspecified.³⁹ By stressing the possibility that learning has propelled us to a self-confirming equilibrium in which the government chooses an optimal policy based on a wrong model, the learning literature changes how we should think about promoting the novel policies that will allow misguided governments to break out of the lack-of-experimentation traps to which self-confirming equilibria confine them.

Milton Friedman was willing to use expected utility to explain why people simultaneously gamble and buy insurance ((Friedman and Savage (1948))), but he refrained from using it to form recommendations about macroeconomic policy. He also abstained from using rational expectations models, though early applications of them to the consumption function and the Phillips curve strengthened some of the policy recommendations that he had made. I do not know why he resisted applying expected utility and rational expectations, but Friedman's works on monetary history are about the evolution of monetary authorities' models, how wrong models sometimes led to bad policies, and how arguments and evidence sometimes improved models. Friedman said that evolution may not yet have taught everything we have to learn about running a fiat money system:

... the world is now engaged in a great experiment to see whether it can fashion a different anchor, one that depends on government restraint rather than on the cost of acquiring a physical commodity ... The verdict is far from in on whether fiat money will involve a lower cost than commodity money ... Friedman (1991, p. 42).

Nonetheless, it remains an open question whether the temptation to use fiat money as a source of revenue will lead to a situation that will ultimately force a return to a commodity standard ... The final answer will come only as history unfolds over the next decades. What that answer will be depends critically on our success in learning from historical episodes such as those that have been

³⁹This is the starting point of calibration in macroeconomics, i.e., the refusal to use maximum likelihood because the model builder views it as an approximation.

examined in this book. Such a learning process has been under way for centuries, ever since the first appearance of systematic analyses of money and monetary institutions. It has entered a new and urgent stage as the world ventures into hitherto unexplored terrain. Friedman (1991, pp. 259-260).

Appendixes

A Learning in games

In a game, a Nash equilibrium is the natural counterpart of a rational expectations equilibrium or a recursive competitive equilibrium. An extensive literature studies whether a system of adaptive players converges to a Nash equilibrium. A range of plausible adaptive algorithms have been proposed that are differentiated by how much foresight and theorizing they attribute to the players.⁴⁰ At one extreme are adaptive models that have naive players who ignore strategic interactions and either play against histograms of their opponents past actions (this is called fictitious play) or alter their moves in directions that *ex post* reduce their *regret* at not having taken other actions in the past, given their opponents' histories of actions. At the other extreme are models in which players construct statistical theories about their opponents' behavior, use them for a while to make forward-looking decisions, occasionally subject their theories to hypothesis tests, discard rejected ones and choose new specifications.

This literature has sought plausible and robust algorithms that converge to a Nash equilibrium. Hart and Mas-Colell tell us that this is a tall order:

It is notoriously difficult to formulate sensible adaptive dynamics that guarantee convergence to Nash equilibrium. In fact, short of variants of exhaustive search (deterministic or stochastic), there are no general results. Hart and Mas-Colell (2003, p. 1830)

Hart and Mas-Colell and Foster and Vohra (1999) show that the source of the difficulty is that most adaptive schemes specify that adjustments in a player's strategy do not depend on the payoff functions of other players, an uncoupling of the dynamics that in general prevents the system from converging to a Nash equilibrium. Many examples of the adaptive schemes in the literature are uncoupled. Because many game theorists find uncoupled schemes desirable, parts of the literature have lowered the bar by looking for convergence to something weaker than Nash equilibria, namely, correlated equilibria or coarse correlated equilibria. Hart and Mas-Colell (2003, p. 1834) observed that "It is thus interesting that Nash equilibrium, a notion that does not predicate coordinated behavior, cannot be guaranteed to be reached in an uncoupled way, while correlated equilibrium, a notion based on coordination, can."⁴¹

Hart and Mas-Colell (2000, 2001, 2003) study adaptive schemes that are backward looking. For example, some of the most interesting ones have a player construct counterfactual historical payoffs that he would have received had he played other strategies, then compute

⁴⁰For a critical survey of this literature, see Young (2004).

⁴¹Experimental economics has supplied data sets designed to check ideas from the literature on adaptive learning in games. It is remarkable that laboratory experiments using macroeconomics are rarer than those using microeconomics. See Duffy (2006) for an account of the existing experiments. I suspect that the main reason for fewer experiments in macro than in micro is that the choices confronting artificial agents within even one of the simpler recursive competitive equilibria used in macroeconomics are very complicated relative to the settings that experimentalists usually confront their subjects with.

a measure of regret, then adjust future play in directions that would have minimized regret. These schemes impute little or no theorizing and foresight to the players.

For my present purposes, one of the most interesting contributions comes from part of the literature that attributes more sophistication to players, in particular, the work of Foster and Young (2003), which is also summarized in Young (2004, ch. 8).⁴² Their model has the following components: (1) each player has a large set of potential models that describe his opponents' strategies; (2) players use a random device to select a particular model; (3) after that model is selected, there is an 'act and collect data' period during which a player (incorrectly) assumes that he will believe his current model forever; during this period, each player chooses his actions via a smoothed best response to what his model tells him about opponents' actions (e.g., a quantal response function); (4) after a data collection period, a player compares the empirical pattern of his opponents' play with that predicted by his model. He performs an hypothesis test that compares the theoretical and empirical distributions. If he rejects his current model, he randomly draws a new model from his set of models, then returns to step 2. If he accepts the model, he returns to step 3, waits a random number of periods, and then begins another data collection period.

With suitable assumptions about the lengths of testing periods and the tolerances of the hypothesis tests, Foster and Young (2003) show that behaviors eventually emerge that are often close to Nash equilibria. Their notion of hypothesis tests is sufficiently broad to include many plausible procedures. Their convergence result seems to be an uncoupled multi-agent learning scheme that actually approaches Nash equilibria, not something weaker like the coarse correlated equilibrium that the entirely backward-looking schemes mentioned above can approach. They avoid the conundrum of Hart and Mas-Colell partly by weakening the notion of convergence.

B From commodity to fiat money

My theme is the data and theorizing that helped inspire David Ricardo's idea:

The introduction of the precious metals for the purposes of money may with truth be considered as one of the most important steps towards the improvement of commerce, and the arts of civilised life; but it is no less true that, with the advancement of knowledge and science, we discover that it would be another improvement to banish them again from the employment to which, during a less enlightened period, they had been so advantageously applied. Ricardo (1816, p. 65)

A long process with "much backing and filling and confusion about purpose and power" led to Ricardo's idea,⁴³ which led Keynes and others to propose a fiat currency.

⁴²For a distinct but related approach, see Jehiel (1995, 1998). The Foster and Young (2003) model seems to me to capture some of the flavor of the anticipated utility framework advocated by Kreps (1998). The classifier models in Marimon et al. (1990) have a similar flavor.

⁴³I borrowed the words in quotes from Friedman and Schwartz (1963, p.193), who used them to describe the evolution of the beliefs and policies within the Federal Reserve.

B.1 Learning to supplement a commodity currency with tokens

Redish (1990, 2000) and Sargent and Velde (2002) described how it took 800 years to understand and cope with two imperfections that marred an ideal self-regulating commodity money system in which coins of all denominations were meant to exchange at values proportional to silver (or gold) content. In the ideal system, a government instructed a mint to offer to sell coins of different denominations for silver at prices proportional to their weights in silver. The mint did not buy coins for silver, but citizens were free to melt silver coins to recover silver. If minting and melting were costless, this self-regulating system would automatically adjust the denomination structure of coins to suit coin holders' preferences by letting them melt coins of a denomination they wanted less of, then take the silver to the mint to buy coins of the denomination they wanted.⁴⁴ In the ideal system, a silver melt point equaled a silver mint point, denomination by denomination.

In practice, two imperfections hampered this system: (1) it was costly to produce coins; and (2) coins depreciated through wear and tear and sweating and clipping. The first imperfection gave rise to nonempty intervals between melt and mint points for gold or silver coins of each denomination – an upper point that indicated a melting point for that coin and a lower one that prompted minting. The proportionate spreads between minting and melting points differed because as a fraction of the value of the coin, it was cheaper to produce a large denomination coin than a small denomination coin. Unless the government were to subsidize the mint for producing low denomination coins, the spread between minting and melting points would be proportionately wider for low denomination coins. The second imperfection allowed underweight coins to circulate along side full weight coins.

A nonempty interval between melting and minting points allowed coins to circulate by *tale* (i.e., by what is written on the coin rather than by weight) at an exchange value that exceeded their value by weight. Indeed, as Adam Smith pointed out, in the presence of costs of producing coins, the money supply mechanism provided incentives for people to purchase new coins from the mint only when their value in exchange exceeded their value by weight by enough to cover the mint's brassage and seigniorage fees (Smith 1789, Book I, ch. 5).

Nonempty intervals with proportionately wider widths for lower denomination coins and a consequent exchange rate indeterminacy allowed the intervals to shift over time and eventually to become so misaligned that they recurrently provided incentives to melt small denomination coins. That created the recurring shortages of small coins documented by Cipolla (1956, 1982).⁴⁵

Cipolla (1956) described a temporary practical remedy for these shortages. The authorities debased small denomination coins, thereby shifting their mint-melt intervals in a direction that motivated citizens to purchase new coins from the mint. Monetary authorities throughout Europe used this method for hundreds of years. There were repeated debasements in small denomination silver coins and secular declines in rates of exchange of small denomination for large denomination coins.

Many experiments, some inadvertent, others purposeful, were performed, and numerous

⁴⁴Sargent and Velde (2002, p. 95) cited Bernardo Davanzati, who in 1588 wrote that “metal should be worth as much in bullion as in coin, and be able to change from metal to money and money to metal without loss, like an amphibious animal.”

⁴⁵This multi-interval commodity money system in which coins circulate by tale is taken for granted by Smith (1789, book I, ch. 5).

theoretical tracts were written and disputed before what Cipolla (1956) called the ‘standard formula’ for issuing token small denomination coins was put into practice in the mid 19th century.⁴⁶ It solved the problem of misaligned mint-melt intervals for coins of different denominations by, first, having only one large denomination full weight coin that the mint sold for a precious metal, and, second, having the government issue difficult-to-counterfeit small denomination token coins that it promised to convert on demand into the large denomination coin. This required a technology for manufacturing coins that were difficult to counterfeit.⁴⁷

As examples of inadvertent experiments, token monies were occasionally issued inside besieged cities and sometimes they worked. A document that prefigured later arguments of John Law, Adam Smith, and David Ricardo sparked a purposeful experiment. It advised King Ferdinand II of Spain that he could issue token copper coins that Spanish residents would voluntarily accept from the government in exchange for full bodied silver coins. It described how this fiscal boon to the Spanish treasury could be attained in a noninflationary way.⁴⁸ Three successive Spanish Kings tried this experiment, which had all of the ingredients of the 19th century standard formula except convertibility. For 25 years, the experiment worked well, yielding the government substantial revenues without inflation. But eventually excessive issues of copper coins caused inflation, in the aftermath of which the Spanish monetary authorities pursued a fascinating sequence of experiments. They restamped copper coins and manipulated the unit of account in order either to adjust the price level or raise revenues for the Spanish government.

The quantity theory can operate only in the limited interval between the mint and melt points for the precious metal, so a commodity money system conceals the quantity theory. When the Spanish broke through those restrictions, they gave the British statistician Sir William Petty data that he used to discover a quantity theory of money (see Hull (1899)). Other episodes created more data to substantiate the quantity theory of money, for example, the construction and collapse of John Law’s system (see Velde (2007)) and the overissuing of French assignats after the sales of the church lands that had initially backed them were suspended after war broke out in 1792 (see Sargent and Velde (1995)). But the episodes that lent vivid empirical support to a quantity theory also brought evidence that government monetary authorities could not be trusted to administer a pure fiat standard in ways that stabilized prices.⁴⁹

In 1660, the master of the British mint, Henry Slingsby, added an element missing from the Spanish experiment, namely, convertibility of token coins, and recommended what in the 19th century became the standard formula.⁵⁰ But perhaps because the inflation accompanying the Spanish and similar experiments had given token coins a bad name, the British

⁴⁶This process of shuttling through experiments, reformulations of theories, and further experiments reminds me of the hypothesis-testing learning models of Foster and Young (2003) and Cho and Kasa (2006), but I might be imagining things.

⁴⁷See Redish (1990, 2000) and Selgin (2003).

⁴⁸See the document cited in Sargent and Velde (2002, pp. 231-232).

⁴⁹I suspect that is why later advocates for replacing the gold standard with ‘more scientific’ systems of managed currencies including Adam Smith and Ricardo to Keynes purposefully omitted references to some of the historical experiments that generated the data that were sources for the quantity theory of money. For example, Smith (1789) did not cite John Law’s theoretical writings as among the sources for his monetary recommendations.

⁵⁰See Sargent and Velde (2002, pp. 268-269).

government ignored Slingsby's recommendations. Many experts, including Locke (1691), continued to insist that token coins of any denomination were dangerous and that a good faith commodity money system required that coins of all denominations be full bodied. For a long time, that sentiment convinced national governments not to issue tokens, but that did not stop other entities from creating them. In seventeenth and eighteenth century Britain, hundreds of private firms and municipalities issued small denomination tokens that formed a substantial part of the country's coinage. Between 1816 and 1836, the British government implemented the standard formula by nationalizing a token coin industry that had long existed.

B.2 Ricardo's proposal

It required 156 years to take the short logical step from Slingsby's 1660 standard formula for issuing convertible token subsidiary coins to David Ricardo's 1816 recommendation. Ricardo proposed that a country's domestic money supply should ideally consist of paper notes that the government promises to exchange at a pegged price for gold bullion bars, but that no gold coins should actually be minted. A variant of Ricardo's scheme in which a government promises to redeem domestic notes for gold, but only for foreign residents, came to be practiced around 1900. This arrangement, by which "a cheap local currency [is] artificially maintained at par with the international standard of value," (Keynes 1913, p. 25) was called the "gold exchange standard." Keynes described how by 1913 this system had come to prevail in India through a sequence of haphazard administrative decisions that eventually produced a coherent system that no one had planned but that Keynes applauded. Keynes (1913, p. 25) predicted that Ricardo's scheme would be an essential part of "the ideal currency system of the future."⁵¹

The standard formula eliminates the gold or silver points for all coins except one standard coin and uses the mint and melt points for that coin to regulate the total quantity of money, while it uses its promise freely to convert tokens into that standard coin to produce the correct denomination composition. It was one more step from the standard formula or Ricardo's proposal to the recommendation of Fisher (1920), Keynes, and others that well intentioned government officials should administer a fiat currency in ways that stabilize the price level. Doing that would allow them to remove the mint and melt points for the one standard coin too. Discovering the quantity theory of money was an essential step in learning the conditions under which a fiat money system could be managed to provide greater price level stability than could be achieved with a gold standard.

As Keynes wanted, in the twentieth century governments throughout the world carried out the historically unprecedented experiment of managing currencies completely cut off

⁵¹Speaking of how a change in Indians' preferences for holding gold could cause world-wide inflation in prices:

The time may not be far distant when Europe, having perfected her mechanism of exchange on the basis of a gold standard, will find it possible to regulate her standard of value on a more rational and stable basis. It is not likely that we shall leave permanently the most intimate adjustments of our economic organism at the mercy of a lucky prospector, a new chemical process, or a change of ideas [preferences for holding gold] in Asia. (Keynes 1913, p. 71)

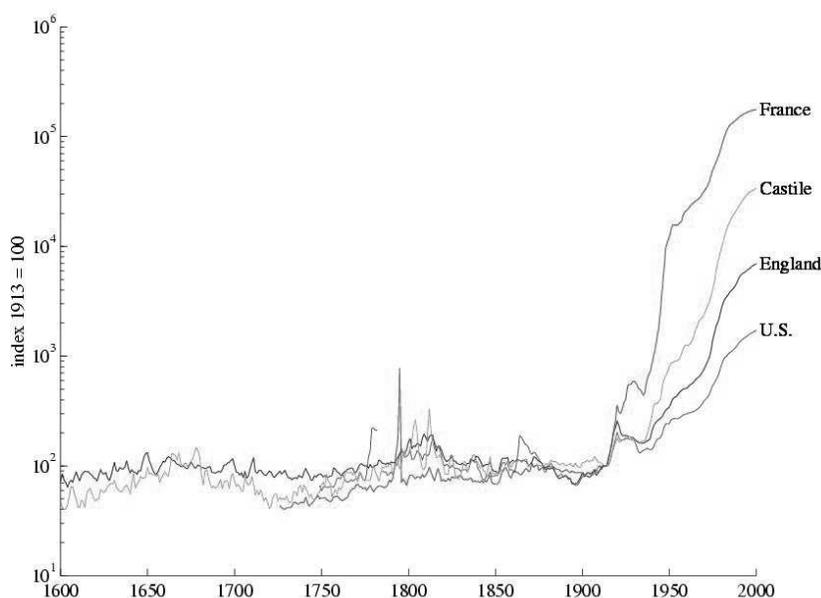


Figure 5: Indices of prices in terms of unit of account in England, the United States, France, and Spain. Sargent and Velde (2002, p. 35)

from gold backing (see Friedman (1991, p. 245)). Figure 5 documents that, at least until very recently, the monetary authorities in four hard-currency countries failed to deliver the kind of price stability that a commodity standard had achieved. Figures 6 and 7 show price indexes for Istanbul and Argentina, places with softer currencies (compare the vertical scales).

C A monetary policy rules literature

The adaptive models described in section 7 explain the rise and fall of post WWII U.S. inflation in terms of monetary policy rules that drifted over time in response to drifts over time in the monetary authorities' models of the economy. All three models embed very crude descriptions of the monetary policy rules and sidestep interesting questions about monetary policy transmission mechanisms. It is appropriate to say a few words about a related literature that uses time series data to infer the structure of post WWII U.S. monetary policy rules and how they have changed over time. The bottom line is that this literature has mixed evidence about whether monetary policy rules shifted enough to validate stories along the lines of our three adaptive models.⁵²

Bernanke and Mihov (1998) developed an SVAR methodology for measuring innovations in monetary policy and their macroeconomic effects. They compared alternative ways of measuring monetary policy shocks and derived a new measure of policy innovations based

⁵²This mixed news partly reflects the fact that it is statistically difficult to detect drifts or shifts in the systematic part of a vector autoregression and much easier to detect changes in volatilities.

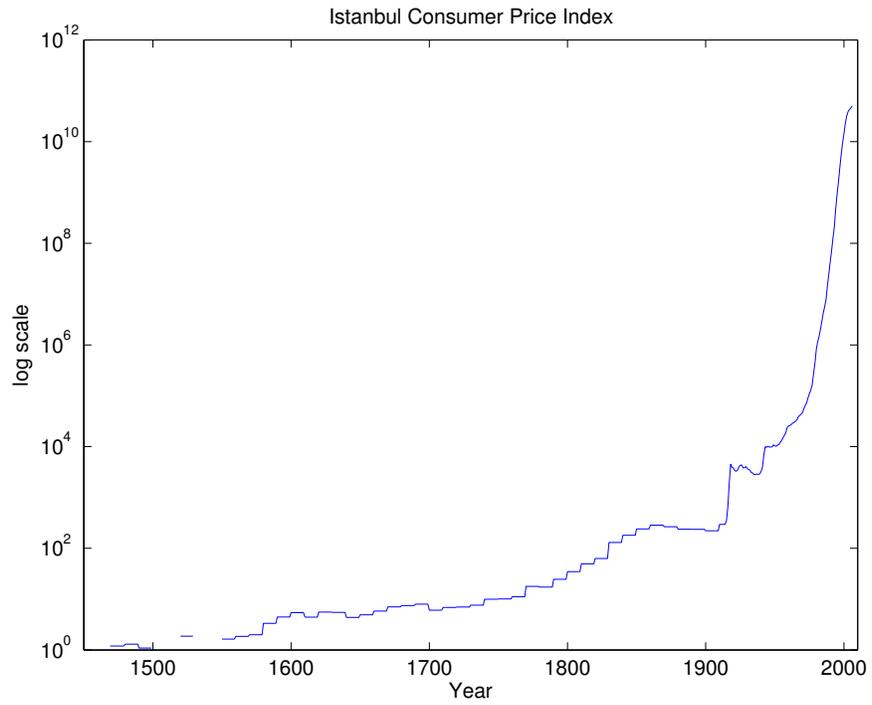


Figure 6: Indices of prices in Istanbul.

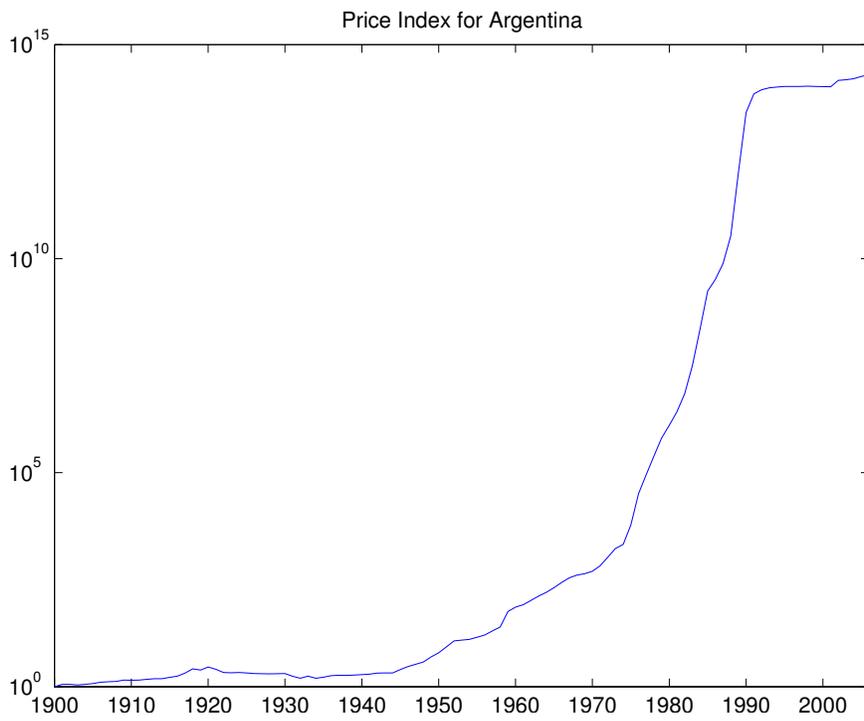


Figure 7: Price index for Argentina.

on possibly time-varying estimates of the Fed's operating procedures. They presented a measure of the overall stance of policy (see Bernanke and Mihov (1998, Fig. III, p. 899)) that is striking in how the distribution of tight and loose policies seems not to have changed much in the periods before and after 1980.

But Clarida et al. (2000) estimated a forward-looking monetary policy reaction function for the postwar United States economy before and after Volcker's appointment as Fed Chairman in 1979 and found substantial differences in the estimated rules across periods. They found that interest rate policy in the Volcker-Greenspan period has been much more sensitive to changes in expected inflation than in the pre-Volcker period. They then extracted implications of the estimated rules for the equilibrium properties of inflation and output in a new Keynesian DSGE model and found that the Volcker-Greenspan rule is stabilizing, but that the earlier rule was not. Lubik and Schorfheide (2004) estimated a new Keynesian model like Clarida et al.'s in which the equilibrium is undetermined if monetary policy is passive and constructed posterior weights for the determinacy and indeterminacy region of the parameter space as well as estimates for the propagation of fundamental and sunspot shocks. They found that U.S. monetary policy post-1982 was consistent with determinacy but that the pre-Volcker policy was not, and also that before 1979 indeterminacy substantially altered the propagation of shocks.

In contrast, working in terms of less fully interpreted models, Sims and Zha (2006) estimated a multivariate regime-switching model for monetary policy and found that the best fit allows time variation in disturbance variances only. When they permitted the systematic VAR coefficients to change, the best fit was with change only in the monetary policy rule. They estimated three regimes that correspond to periods across which the folk-wisdom states that monetary policy differed. But they found that those differences among regimes were not large enough to account for the rise and decline of inflation of the 1970s and 1980s. Likewise, by estimating a time-varying VAR with stochastic volatility, Primiceri (2005) found that both the systematic and non-systematic components of monetary policy had changed. In particular, he found that the systematic responses of the interest rate to inflation and unemployment exhibited a trend toward a more aggressive behavior, while also having sizeable high frequency oscillations. But Primiceri concluded that those had small effects on the rest of the economy and that exogenous non-policy shocks were more important than interest rate policy in explaining the U.S. inflation and unemployment episodes of the 1970's, thus coming down more on the 'bad luck' than the 'bad policies' side of the argument. I hope that conclusion is too pessimistic because we have learned to do better.

References

- Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1986. Optimal Cartel Equilibria with Imperfect Monitoring. *Journal of Economic Theory* 39:251–269.
- . 1990. Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica* 58 (5):1041–1063.
- Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini. 2006. Learning and Stock Market Volatility. Computing in economics and finance 2006, Society for Computational Economics.
- Aiyagari, S. Rao, Albert Marcet, Thomas J. Sargent, and Juha Seppala. 2002. Optimal Taxation without State-Contingent Debt. *Journal of Political Economy* 110 (6):1220–1254.
- Albanesi, Stefania, V.V. Chari, and Lawrence J. Christiano. 2002. Expectation Traps and Monetary Policy. Nber working papers, National Bureau of Economic Research, Inc.
- Barro, Robert J. 1979. On the Determination of Public Debt. *Journal of Political Economy* 87:940–971.
- Bernanke, Ben S. 2007. Inflation Expectations and Inflation Forecasting. Remarks at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts, available on the Federal Reserve Board web site.
- Bernanke, Ben S. and Ilian Mihov. 1998. Measuring Monetary Policy. *The Quarterly Journal of Economics* 113 (3):869–902.
- Bernanke, Ben S., Thomas Laubach, Frederic S. Mishkin, and Adam S. Posen. 2001. *Inflation Targeting: Lessons from the International Experience*. Princeton, New Jersey: Princeton University Press.
- Blinder, Alan S. 1998. *Central Banking in Theory and Practice (Lionel Robbins Lectures)*. Cambridge, Massachusetts: MIT Press.
- Blume, Lawrence and David Easley. 2006. If You’re so Smart, why Aren’t You Rich? Belief Selection in Complete and Incomplete Markets. *Econometrica* 74 (4):929–966.
- Bossaerts, Peter. 2002. *The Paradox of Asset Pricing*. Princeton, New Jersey: Princeton University Press.
- . 2004. Filtering Returns for Unspecified Biases in Priors when Testing Asset Pricing Theory. *Review of Economic Studies* 71:6386.
- Branch, William A. and George W. Evans. 2005. Model Uncertainty and Endogenous Volatility. Computing in economics and finance 2005, Society for Computational Economics.
- . 2006. Intrinsic heterogeneity in expectation formation. *Journal of Economic Theory* 127 (1):264–295.

- Bray, Margaret M. and David M. Kreps. 1987. Rational Learning and Rational Expectations. In *Arrow and the Ascent of Modern Economic Theory*, edited by George R. Feiwel, 597–625. New York University Press.
- Brock, William A. and Patrick de Fontnouvelle. 2000. Expectational diversity in monetary economies¹. *Journal of Economic Dynamics and Control* 24 (5-7):725–759.
- Brock, William A. and Cars H. Hommes. 1997. A Rational Route to Randomness. *Econometrica* 65 (5):1059–1096.
- Bullard, James and Kaushik Mitra. 2007. Determinacy, Learnability, and Monetary Policy Inertia. *Journal of Money, Credit and Banking* 39 (5):1177–1212.
- Cagan, Phillip. 1956. *The Monetary Dynamics of Hyperinflation*, 25–117. Chicago: University of Chicago Press.
- Calvo, Guillermo A. 1978. On the Time Consistency of Optimal Policy in a Monetary Economy. *Econometrica* 46:1411–1428.
- Campbell, John Y. and John H. Cochrane. 1999. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* 107:205–251.
- Carboni, Giacomo and Martin Ellison. 2007. The Great Inflation and the Greenbook. European Central Bank and University of Warwick.
- Chang, Roberto. 1998. Credible monetary policy with long-lived agents: recursive approaches. *Journal of Economic Theory* xxx–xxx.
- Chari, V. V. and Patrick J. Kehoe. 1993a. Sustainable Plans and Debt. *Journal of Economic Theory* 61 (2):230–261.
- Chari, V V and Patrick J Kehoe. 1993b. Sustainable Plans and Mutual Default. *Review of Economic Studies* 60 (1):175–95.
- Chari, V. V., Lawrence J. Christiano, and Martin Eichenbaum. 1998. Expectation Traps and Discretion. *Journal of Economic Theory* 81 (2):462–492.
- Cho, In-Koo and Kenneth Kasa. 2006. Learning and Model Validation. University of Illinois.
- Cho, In-Koo, Noah Williams, and Thomas J. Sargent. 2002. Escaping Nash Inflation. *Review of Economic Studies* 69:1–40.
- Chung, Heetaik. 1990. *Did Policy Makers Really Believe in the Phillips Curve?: An Econometric Test*. Ph.D. thesis, University of Minnesota.
- Cipolla, Carlo. 1956. *Money, Prices, and Civilization in the Mediterranean World, Fifth to Seventeenth Century*. New York: Gordian Press.
- . 1982. *The Monetary Policy of Fourteenth-Century Florence*. Berkeley: University of California Press.

- Clarida, Richard, Jordi Gali, and Mark Gertler. 1999. The Science of Monetary Policy: A New Keynesian Perspective. *Journal of Economic Literature* 37 (4):1661–1707.
- Clarida, Richard, Jordi Gal, and Mark Gertler. 2000. Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. *The Quarterly Journal of Economics* 115 (1):147–180.
- Cogley, Timothy. 2005. Changing Beliefs and the Term Structure of Interest Rates: Cross-Equation Restrictions with Drifting Parameters. *Review of Economic Dynamics* 8 (2):420–451.
- Cogley, Timothy and Thomas J. Sargent. 2005. The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty. *Review of Economic Dynamics* 8:528–563.
- Cogley, Timothy, Riccardo Colacito, Lars Hansen, and Thomas Sargent. 2007a. Robustness and U.S. Monetary Policy Experimentation. Unpublished.
- Cogley, Timothy, Riccardo Colacito, and Thomas J. Sargent. 2007b. Benefits from U.S. Monetary Policy Experimentation in the Days of Samuelson and Solow and Lucas. *Journal of Money, Credit and Banking* 39 (s1):67–99.
- Dagli, Ates C. and John B. Taylor. 1984. Estimation and solution of linear rational expectations models using a polynomial matrix factorization. *Journal of Economic Dynamics and Control* 8 (3):341–348.
- Duffy, John. 2006. Experimental Macroeconomics. Prepared for the New Palgrave Dictionary of Economics.
- Epstein, L. and S. Zin. 1989. Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica* 57:937–969.
- Evans, George W. and Seppo Honkapohja. 1999. Learning Dynamics. In *Handbook of Macroeconomics*, edited by John B. Taylor and Michael Woodford. Elsevier.
- . 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press.
- . 2003. Expectations and the Stability Problem for Optimal Monetary Policies. *Review of Economic Studies* 70 (4):807–824.
- Fetter, Frank Whitson. 1978. *Development of British Monetary Orthodoxy, 1797-1875*. Fairfield, New Jersey: Augustus M. Kelley. First published 1965 by Harvard University Press, Cambridge, Ma.
- Fisher, Irving. 1920. *Stabilizing the Dollar: A Plan to Stabilize the General Price Level without Fixing Individual Prices*. New York: Macmillan.
- . 1926. *The Purchasing Power of Money: Its Determination and Relation to Credit, Interest and Crises*. New York: Macmillan. New and revised edition, originally published in 1911.

- Foster, Dean P. and Rakesh Vohra. 1999. Regret in the On-Line Decision Problem. *Games and Economic Behavior* 29 (1-2):7–35.
- Foster, Dean P. and H. Peyton Young. 2003. Learning, hypothesis testing, and Nash equilibrium. *Games and Economic Behavior* 45 (1):73–96.
- Friedman, M. and A.J. Schwartz. 1963. *A Monetary History of the United States, 1857-1960*. Princeton, New Jersey: Princeton University Press.
- Friedman, Milton. 1956. *A Theory of the Consumption Function*. Princeton, N.J.: Princeton University Press.
- . 1991. *Money Mischief: Episodes in Monetary History*. New York, San Diego, and London: Harcourt Brace Jovanovich.
- Friedman, Milton and L. J. Savage. 1948. The Utility Analysis of Choices Involving Risk. *Journal of Political Economy* 56.
- Fudenberg, Drew and David K. Levine. 1998. *Learning in Games*. Cambridge, MA.: M.I.T. Press.
- . 2007. Self-Confirming and the Lucas Critique. Unpublished manuscript, Harvard University and Washington University.
- Gallant, A Ronald and George Tauchen. 1996. Which Moments to Match? *Econometric Theory* 12 (4):657–81.
- Golosov, Mikhail, Narayana Kocherlakota, and Aleh Tsyvinski. 2003. Optimal Indirect and Capital Taxation. *Review of Economic Studies* 70 (3):569–587.
- Golosov, Mikhail, Aleh Tsyvinski, and Ivan Werning. 2007. New Dynamic Public Finance: A User’s Guide. *NBER Macroeconomics Annual 2006* 317–364.
- Gourieroux, C., A. Monfort, and E. Renault. 1993. Indirect Inference. *Journal of Applied Econometrics* 8:S85S118.
- Grossman, Sanford J and Robert J Shiller. 1981. The Determinants of the Variability of Stock Market Prices. *American Economic Review* 71 (2):222–27.
- Hansen, Lars Peter. 2007. Beliefs, Doubts and Learning: Valuing Economic Risk. *American Economic Review, Papers and Proceedings* 97 (2):1–30.
- Hansen, Lars Peter and Ravi Jagannathan. 1991. Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy* 99 (2):225–62.
- Hansen, Lars Peter and Scott F Richard. 1987. The Role of Conditioning Information in Deducing Testable Restrictions. *Econometrica* 55 (3):587–613.
- Hansen, Lars Peter and Thomas J. Sargent. 1980. Formulating and estimating dynamic linear rational expectations models. *Journal of Economic Dynamics and Control* 2 (2):7–46.

- . 2006. Fragile beliefs and the price of model uncertainty. Unpublished.
- Hansen, Lars Peter and Kenneth J Singleton. 1983. Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns. *Journal of Political Economy* 91 (2):249–65.
- Harrison, J. Michael and David M. Kreps. 1979. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* 20 (3):381–408.
- Hart, Sergiu and Andreu Mas-Colell. 2000. A Simple Adaptive Algorithm Leading to Correlated Equilibrium. *Econometrica* 68:1127–50.
- . 2001. A General Class of Adaptive Strategies. *Journal of Economic Theory* 98 (1):26–54.
- . 2003. Uncoupled Dynamics Do Not Lead to Nash Equilibrium. *The American Economic Review* 93 (5):1830–1836.
- Hetzl, Robert. 1998. Arthur Burns. *Federal Reserve Bank of Richmond Economic Quarterly* 84 (1):21–44.
- Hull, Charles Henry, ed. 1899. *The Economic Writings of Sir William Petty*. Cambridge: Cambridge University Press.
- Jehiel, Philippe. 1995. Limited Horizon Forecast in Repeated Alternate Games. *Journal of Economic Theory* 67 (2):497–519.
- . 1998. Learning to Play Limited Forecast Equilibria. *Games and Economic Behavior* 22 (2):274–298.
- Kalai, Ehud and Ehud Lehrer. 1993. Rational Learning Leads to Nash Equilibrium. *Econometrica* 61 (5):1019–45.
- . 1994. Weak and strong merging of opinions. *Journal of Mathematical Economics* 23 (1):73–86.
- Keynes, John Maynard. 1913. *Indian Currency and Finance*. Cambridge, UK: University Printing House. The Royal Economic Society edition, 1971.
- . 1936. *The General Theory of Employment Interest and Money*. New York: Harcourt, Brace and Company.
- King, Robert G. and Mark W. Watson. 1994. The post-war U.S. phillips curve: a revisionist econometric history. *Carnegie-Rochester Conference Series on Public Policy* 41:157–219.
- Kocherlakota, Narayana R. 2005. Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation. *Econometrica* 73 (5):1587–1621.
- Kreps, David M. 1998. Anticipated Utility and Dynamic Choice. In *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997*. Cambridge University Press.

- Kydland, Finn E and Edward C Prescott. 1977. Rules Rather Than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy* 85 (3):473–91.
- Ljungqvist, Lars and Thomas J. Sargent. 2004. *Recursive Macroeconomic Theory, Second Edition*. Cambridge, Mass.: MIT Press.
- Locke, John. 1691. *Some Considerations of the Consequences of the Lowering of Interest and the Raising of the Value of Money*. London: Awnsham and John Churchill.
- Lubik, Thomas A. and Frank Schorfheide. 2004. Testing for Indeterminacy: An Application to U.S. Monetary Policy. *American Economic Review* 94 (1):190–217.
- Lucas, Jr, Robert E. 1973. Some International Evidence on Output-Inflation Tradeoffs. *American Economic Review* 63 (3):326–34.
- Lucas, Robert E., Jr. 1972a. Econometric Testing of the Natural Rate Hypothesis. In *The Econometrics of Price Determination: Conference, October 30-31*, edited by Otto Eckstein. Washington, D.C.: Board of Governors, Federal Reserve System.
- . 1972b. Expectations and the neutrality of money. *Journal of Economic Theory* 4 (2):103–124.
- . 1976. Econometric Policy Evaluation: A Critique. In *The Phillips Curve and Labor Markets*, edited by K. Brunner and A. H. Meltzer, 19–46. Amsterdam: North-Holland.
- Lucas, Robert E., Jr. and Nancy L. Stokey. 1983. Optimal fiscal and monetary policy in an economy without capital. *Journal of Monetary Economics* 12 (1):55–93.
- Madison, James. 1987. *Notes of Debates in the Federal Convention of 1787 Reported by James Madison*. London: W.W. Norton & Company. Bicentennial Edition.
- Marcet, Albert and Juan P. Nicolini. 2003. Recurrent Hyperinflations and Learning. *American Economic Review* 93 (5):1476–1498.
- Marcet, Albert and Thomas J. Sargent. 1989a. Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models. *Journal of Economic Theory* 48:337–368.
- . 1989b. Least-Squares Learning and the Dynamics of Hyperinflation. In *International Symposia in Economic Theory and Econometrics*, edited by William Barnett, John Geweke, and Karl Shell, 119–137. Cambridge, England: Cambridge University Press.
- Marimon, Ramon. 1997. Learning from Learning in Economics. In *Advances in Economics and Econometrics: Theory and Applications*, edited by David Kreps and Kenneth Wallis. Cambridge University Press.
- Marimon, Ramon, Ellen McGrattan, and Thomas J. Sargent. 1990. Money as a medium of exchange in an economy with artificially intelligent agents. *Journal of Economic Dynamics and Control* 14 (2):329–373.

- Mehra, Rajnish and Edward C. Prescott. 1985. The equity premium: A puzzle. *Journal of Monetary Economics* 15 (2):145–161.
- Meyn, S. P. 2007. Control Techniques for Complex Networks. To appear, Cambridge University Press.
- Mishkin, Frederic S. 2007. Inflation Dynamics. Remarks at the Annual Macro Conference, Federal Reserve Bank of San Francisco, available on the Federal Reserve Board web site.
- Orphanides, Athanasios. 2002. Monetary-Policy Rules and the Great Inflation. *American Economic Review* 92 (2):115–120.
- . 2003. The quest for prosperity without inflation. *Journal of Monetary Economics* 50 (3):633–663.
- Orphanides, Athanasios and John C. Williams. 2005. The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations. *Journal of Economic Dynamics and Control* 29 (11):1927–1950.
- . 2007. Inflation targeting under imperfect knowledge. *Economic Review* 1–23.
- Phelan, Christopher and Ennio Stacchetti. 2001. Sequential Equilibria in a Ramsey Tax Model. *Econometrica* 69 (6):1491–1518.
- Phelps, Edmund S. 1967. Phillips Curves, Expectations of Inflation and Optimal Unemployment over Time. *Economica* 2 (3):22–44.
- Piazzesi, Monika and Martin Schneider. 2007. Equilibrium Yield Curves. *NBER Macroeconomics Annual 2006* 389–442.
- Primiceri, Giorgio. 2005. Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies* 72:821–852.
- . 2006. Why Inflation Rose and Fell: Policymakers’ Beliefs and US Postwar Stabilization Policy. *Quarterly Journal of Economics* 121 (3):867–901.
- Redish, Angela. 1990. The Evolution of the Gold Standard in England. *Journal of Economic History* 50 (4):789–805.
- . 2000. *Bimetallism: An Economic and Historical Analysis*. Cambridge: Cambridge University Press.
- Ricardo, David. 1816. *Proposals for an Economical and Secure Currency*. London: John Murray, Albermarle-Street. Reprinted in Vol. IV, of *The Works and Correspondence of David Ricardo*, edited by Piero Sraffa, with the collaboration of M.H. Dobb, Liberty Fund, Indianapolis, 2004.
- Sandroni, Alvaro. 2000. Do Markets Favor Agents Able to Make Accurate Predictions? *Econometrica* 68:1303–1342.

- Sargent, Thomas J. 1971. A Note on the 'Accelerationist' Controversy. *Journal of Money, Credit and Banking* 3 (3):721–25.
- Sargent, Thomas J. 1977. The Demand for Money During Hyperinflations under Rational Expectations: I. *International Economic Review* 18 (1):59–82.
- Sargent, Thomas J. 1979. A note on maximum likelihood estimation of the rational expectations model of the term structure. *Journal of Monetary Economics* 5 (1):133–143.
- Sargent, Thomas J. 1984. Autoregressions, Expectations, and Advice. *American Economic Review* 74 (2):408–15.
- Sargent, Thomas J. 1999. *The Conquest of American Inflation*. Princeton University Press.
- Sargent, Thomas J. and François Velde. 2002. *The Big Problem of Small Change*. Princeton, New Jersey: Princeton University Press.
- Sargent, Thomas J. and François R. Velde. 1995. Macroeconomic Features of the French Revolution. *Journal of Political Economy* 103 (3):474–518.
- Sargent, Thomas J. and Noah Williams. 2005. Impacts of Priors on Convergence and Escapes from Nash Inflation. *Review of Economic Dynamics* 8 (2):360–391.
- Sargent, Thomas J., Noah Williams, and Tao Zha. 2006a. The Conquest of South American Inflation. Nber working papers, National Bureau of Economic Research, Inc.
- . 2006b. Shocks and Government Beliefs: The Rise and Fall of American Inflation. *American Economic Review* 96 (4):1193–1224.
- Sawhill, Isabel V. 1995. Distinguished Lecture on Economics in Government: The Economist vs. Madmen in Authority. *Journal of Economic Perspectives* 9 (3):3–13.
- Selgin, George. 2003. Steam, Hot Air, and Small Change: Matthew Boulton and the Reform of Britain's Coinage. *Economic History Review* 56:478–509.
- Sims, Christopher A. 1980. Macroeconomics and Reality. *Econometrica* 48 (1):1–48.
- . 1982. Policy Analysis with Econometric Models. *Brookings Papers on Economic Activity* (1):107–152.
- . 1988. Projecting Policy Effects with Statistical Models. *Revista de Analisis Economico* 3:3–20.
- Sims, Christopher A. and Tao Zha. 2006. Were There Regime Switches in U.S. Monetary Policy? *American Economic Review* 96 (1):54–81.
- Smith, Adam. 1789. *An Inquiry into the Nature and Causes of the Wealth of Nations*. New York: The Modern Library. The Modern library edition is copied from the fifth edition published in 1789, the last life-time edition.

- Smith, Jr, A A. 1993. Estimating Nonlinear Time-Series Models Using Simulated Vector Autoregressions. *Journal of Applied Econometrics* 8 (S):S63–84.
- Stokey, Nancy L. 1989. Reputation and Time Consistency. *American Economic Review* 79:134–139.
- . 1991. Credible Public Policy. *Journal of Economic Dynamics and Control* 15.
- Taylor, John B. 1980. Aggregate Dynamics and Staggered Contracts. *Journal of Political Economy* 88 (1):1–23.
- Velde, François. 2007. John Law’s System. *American Economic Review Papers and Proceedings* 97 (2):276–270.
- Williams, Noah. 2004. Escape Dynamics in Learning Models Princeton University manuscript.
- Woodford, Michael D. 1990. Learning to Believe in Sunspots. *Econometrica* 58:277–307.
- . 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, New Jersey: Princeton University Press.
- Young, H. Peyton. 2004. *Strategic Learning and its Limits*. Oxford and New York: Oxford University Press. The Arne Ryde Memorial Lectures Series.