DOCUMENTOS DE TRABAJO

Private Information and the Mortgage Market: Evidence and a Theory of Crises

Robert Shimer

N.º 716 Diciembre 2013

BANCO CENTRAL DE CHILE

DOCUMENTOS DE TRABAJO

Private Information and the Mortgage Market: Evidence and a Theory of Crises

Robert Shimer

N.º 716 Diciembre 2013

BANCO CENTRAL DE CHILE

BANCO CENTRAL DE CHILE

CENTRAL BANK OF CHILE

Documentos de Trabajo del Banco Central de Chile
Working Papers of the Central Bank of Chile
Agustinas 1180, Santiago, Chile
Teléfono: (56-2) 3882475; Fax: (56-2) 3882231

Documento de Trabajo
N° 716

Working Paper
N° 716

Private Information and the Mortgage Market:

Evidence and a Theory of Crises*

Robert Shimer
University of Chicago

Abstract
A growing body of evidence suggests that private information is important in the mortgage-backed securities market. Recent research reconsiders the theory of how investors trade in the presence of private information. This paper summarizes the evidence and uses the new theoretical approach to explain how trade in mortgage-backed securities can collapse during a crisis.

Resumen
Una creciente evidencia sugiere que la información privada es importante en el mercado de mortgage- backed securities. La investigación reciente reconsidera la teoría de cómo los inversionistas comercian en la presencia de información privada. Este artículo resume la evidencia y utiliza un enfoque teórico nuevo de manera de explicar cómo el comercio de mortgage- backed securities puede colapsar durante una crisis.

* Paper presented at the Central Bank of Chile 2012 Annual Conference, November 15-16, 2012. University of Chicago, Department of Economics, 1126 East 59th Street, Chicago, IL, 60637 (e-mail: robert.shimer@gmail.com). I am grateful for comments from Larry Christiano, Robert Hall, and conference participants on an earlier draft of this paper and to the National Science Foundation for research support.

1. INTRODUCTION

The securitization boom in the United States mortgage market from 2000 to 2005 was enormous (figure 1). According to the Securities Industry and Financial Markets Association (SIFMA), new issuance of securities backed by mortgages that were not insured by the U.S. government rose by a factor of twelve during that five year period, from $58 billion in 2000 to $726 billion in 2005. Issuance of securities backed by home equity loans also soared from $75 billion to $460 billion over the same five year span. The subsequent collapse was even faster. By 2008, issuance of these two types of securities had fallen to $36 billion, further declining to $8 billion by 2012. In contrast, the market for securities backed by insured mortgages has boomed since 2005, nearly doubling from $983 billion to $1.731 trillion by 2012 in the face of declining interest rates. This paper summarizes existing empirical evidence that private information was important in the uninsured mortgage market and then describes recent theoretical models that explain how the emergence of private information can lead to a decline in trade in these securities.

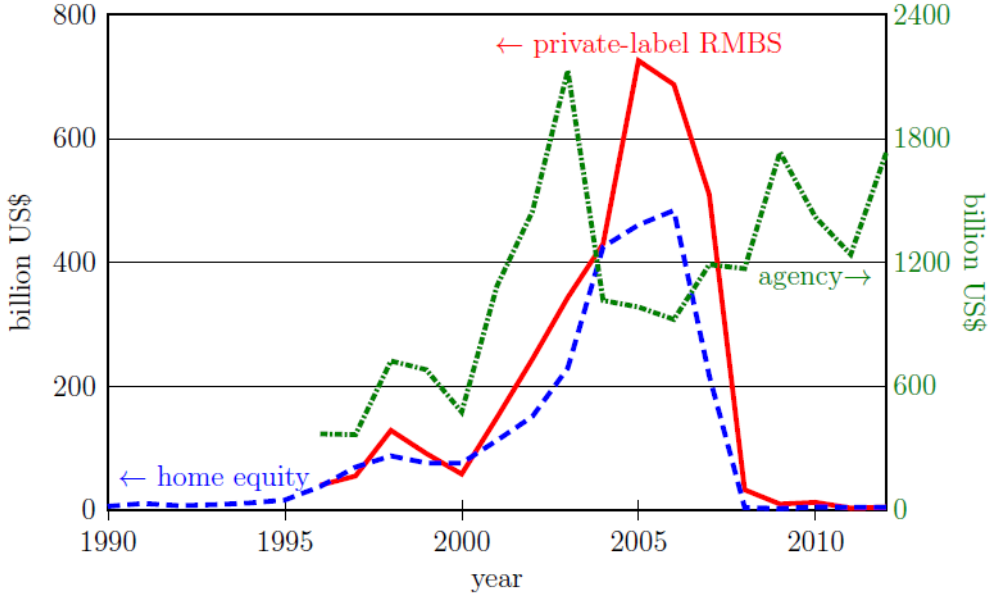

Figure 1: New issuance of private-label residential mortgage-backed securities (left axis), home equity loans (left axis), and agency mortgage-backed securities (right axis). Source: SIFMA, http://sifma.org/research/statistics.aspx, U.S. Mortgage-Related Issuance and Outstanding (private-label RMBS and agency MBS) and U.S. ABS Issuance and Outstanding (home equity loans).

The link between savers and borrowers typically involves a chain of intermediation; the mortgage market is no exception. This paper starts off in section 2 by describing that chain. This is important because private information may arise in one part of the chain but affect intermediation in another part. In particular, I argue in section 3 that mortgage originators obtained private information about the quality of their loans during the origination process. Private information was particularly acute for mortgages that were not insured by the U.S. government because of the very real risk of default on an uninsured mortgage; however, although private information was gathered during the origination process, its consequences were felt in the securitization process. The buyers of mortgage-backed securities (MBSs) were rightly concerned that mortgage securitizers were better informed about the MBSs' default risk.

It is not news that private information is an issue in the MBS industry, and indeed the industry has developed a number of techniques to moderate the amount of private information and mitigate its consequences (section 4). MBSs offer warranties, which were valued by independent specialists and traded as part of a long-term relationship between the buyer and seller. Each of these institutions should reduce buyers' concern that sellers will attempt to profit from their private information. Moreover, tranching was designed to create safe debt from risky mortgage pools, ideally eliminating the relevance of the seller's private information. Similarly, haircuts in repurchase agreements left repo sellers as the residual claimants on an income stream and created a safe, information-insensitive asset for repo buyers.

Unfortunately, all these safeguards were insufficient to prevent the emergence of private information and risk in the MBS industry once house prices began to decline in 2005. Section 5 discusses recent theoretical models that I have jointly developed with Veronica Guerrieri, which offer a framework for analyzing how buyers and sellers set prices in the presence of private information. The key insight is that an endogenous shortage of buyers at high prices allows sellers who have favorable information to separate themselves from those with unfavorable information. This is because sellers with favorable information are more willing to accept a reduction in the probability of trading in return for an increase in the price since they hold onto a better MBS if they fail to sell. I use this model to offer two related stories about the emergence of a crisis in the market for MBSs. In one, a crisis occurs because of a change in fundamentals. In the other, a crisis is a shift in equilibrium in the absence of any intrinsic change in the environment, perhaps caused by contagion from another market. Finally, I conclude in section 6 with a brief discussion of how a crisis in the market for MBS affects the ability of current homeowners to refinance their loans and potential homeowners to obtain loans.

2. PARTICIPANTS IN THE MBS MARKET

The MBS market effectively borrows money from large savers, such as insurance companies and pension funds, and lends it to homeowners and homebuyers. Because the large savers typically do not have any expertise in making loans, there are usually multiple intermediaries lying between the ultimate lenders and ultimate borrowers. In principle, since private information problems could arise at each stage of the intermediation process, this section briefly summarizes who those intermediaries are.

Broadly speaking, there are two main stages in the lending process: origination and securitization. Origination involves making a loan to an individual homeowner. Securitization involves bundling loans together and reselling them to the ultimate lenders. While this paper focuses more on securitization than origination, many of the information issues that arise in the securitization market start in the origination market; therefore, it is useful to think about both stages of this process together.

The main intermediary in the origination process is the mortgage broker. In the United States, a homeowner (or homebuyer) typically deals directly with a mortgage broker. The broker collects relevant information from the homeowner and then connects the homeowner with a mortgage originator who actually grants the loan. Some homeowners bypass this process, obtaining their loan directly from a retail lender rather than a broker. Smaller mortgage originators then typically resell their loans to wholesale aggregators, while larger originators may skip this

step.

At this stage, a securitizer bundles together a large number of mortgages and tranches them, creating a series of bonds with different promised coupon payments, maturities, and seniority. The coupon payments are supposed to be covered by the homeowners' principal and interest payments on the underlying loans. If cash flows ultimately turn out to be too small to support the payments, then the bonds go into default, with junior bonds defaulting before the more senior ones. Finally, the bonds are rated by one or more of the major credit rating agencies and then sold.

Ultimate lenders, including insurance companies and pension funds, purchase some of these investment-grade bonds, which raise capital that the securitizer can reinvest into new mortgages. Foreign and domestic banks purchase other high-grade bonds, holding some on their balance sheet and offloading others into asset-backed commercial paper (ABCP) programs. Banks issue deposits and debt to fund their asset holdings, while ABCP programs typically sell very short-term debt to money market funds to finance their holdings. In both cases, the loan is ultimately funded by a lender who invests in the bank or ABCP program, thus completing the chain from borrower to lender. Finally, securitizers typically hold the junior bonds with junk ratings and the still riskier "equity" tranches on their balance sheet.[1]

3. Evidence of Private Information in MBS Markets

MBS in the United States are divided into two broad categories, agency and private-label, distinguished by the entity that issues the security. Agency MBSs are issued by a government-sponsored enterprise (GSE), especially the Federal National Mortgage Association (Fannie Mae) and the Federal Home Loan Mortgage Corporation (Freddie Mac), or directly by the U.S. government through the Government National Mortgage Association (Ginnie Mae). Agency MBS aggregate and tranche a large number of underlying residential mortgages, promising a coupon payment linked to the homeowners' principal and interest payments. The sponsor, eliminating an important source of risk, in turn guarantees those payments. Still, some residual sources of risk remain, particularly the risk that the mortgage is pre-paid when interest rates fall.

Private-label MBSs are issued by private financial institutions with no guarantee of principal or interest payments. The underlying mortgages are both commercial and residential, but this essay focuses on residential MBSs. These loans typically do not conform to the guidelines imposed by the GSEs because, for example, the loan is too large either in absolute terms or relative to the value of the collateral, or relative to the borrower's income. While pre-payment risk exists for private label securities, default is a bigger source of risk. When a borrower fails to make a payment on his mortgage, some portion of the MBS's promised coupon may be lost.

Data from the Securities Industry and Financial Markets Association (SIFMA) shows that the issuance of private-label MBS collapsed during the financial crisis, from a peak of $883 billion in 2005 to a trough of $18 billion in 2009. In contrast, the agency market performed very well as home owners refinanced to take advantage of

[1] Prior to the financial crisis, the junior tranches were frequently bundled and retranched as Collateralized Debt Obligations; however, that market has largely disappeared.

low interest rates; issuance increased sharply from $983 billion to $1.734 trillion over the same time period. If private information is likely to be important for understanding the financial crisis, the private-label MBS market is therefore a natural place to look for evidence of it. I start my review there. The underlying mortgages backing private-label MBSs are usually well documented. A prospectus describes many characteristics of both the loan and the borrower, such as the distribution of interest rates, maturities, and loan-to-value ratios, as well as the credit score, income, and owner occupancy status of the borrowers. A spreadsheet containing all this information is available to prospective buyers. Despite this, there are several reasons why a mortgage originator may have superior information to the MBS's ultimate buyer, which is often a money market fund, pension fund, or insurance company.

First, many private-label loans had low or no documentation (low-doc loans). In this case, the homeowner was either not asked about his income and assets, or his reports were not verified. Instead, the collateral was supposed to protect the mortgage originator against the borrower's inability to pay. In practice, however, originators had other "soft" information that they used when making these loans, information that could not easily be quantified and was not reported in the prospectus. To the extent that this information is useful for predicting future payments and default, it creates a natural and potentially quantitatively significant source of private information.

Second, it appears that borrowers sometimes provided incorrect answers to questions on loan applications in order to borrow at a better rate. Those misrepresentations naturally carried over to the prospectus and any other information available to MBS buyers. To the extent that mortgage originators could observe these misrepresentations, they were another source of private information. Again, there is some evidence that this is quantitatively significant.

Third, the mortgage originator specializes in evaluating the quality of loans and may therefore be better than the MBS buyer at valuing the fundamentals described in the prospectus, even if the two parties observe the same information. While I believe that this is a reasonable hypothesis, I am unaware of any direct information supporting it. Nevertheless, indirect information suggests that this too may be a relevant source of private information both in the agency and private-label market: mortgage issuers held onto better quality loans and securitized lower quality ones.

I turn next to a more detailed description of these three types of evidence.

3.1. Low Documentation Loans

Unquantified, soft information, such as the mortgage originator's expectation about the buyer's income stability, plays a key role in qualifying borrowers for low-doc mortgages. Since, by its nature, soft information cannot be reported in the MBS prospectus, low-doc loans offer the widest scope of private information. Indeed, the literature has found compelling evidence that supports this hypothesis.

Keys, Mukherjee, Seru and Vig (2010) show that when originators expect to retain (rather than resell) a low-doc mortgage, they screen the loan more carefully. This is consistent with the hypothesis that some aspects of low-doc mortgages are not properly priced in the MBS market because the information is unavailable to mortgage buyers.

The paper uses a regression discontinuity approach to analyze how the likelihood of retaining a mortgage affects

screening. Due to a historical anomaly, it was more difficult for a mortgage originator to securitize a loan if the borrower's credit (FICO) score was below 620, and so the originator was more likely to hold the loan to maturity. The FICO distribution is smooth by construction and it is nearly impossible for a borrower to precisely manipulate his FICO score. Despite this, Keys, Mukherjee, Seru and Vig (2010) show that mortgage originators made roughly twice as many loans to homeowners whose FICO score was slightly above 620 than to homeowners just below this threshold. This suggests that lower quality borrowers were screened more carefully than higher quality ones. In itself, that is not surprising, but the discontinuity in lending practices is most apparent at this critical threshold for securitization.

Next, Keys, Mukherjee, Seru and Vig (2010) document that the few loans made to borrowers with a FICO score just below the key threshold were much less likely to default than the loans made to borrowers just above the threshold. For example, a one-year-old low-doc loan to a borrower with a FICO score between 615 and 619 was about twenty percent (or two percentage points) less likely to default than a similar loan to a similar borrower with a slightly higher FICO score (between 620 and 624). To emphasize, the borrowers with a worse credit rating were less likely to default than those with a better credit rating.

Also, Keys, Mukherjee, Seru and Vig (2010) show that there was no difference in the lending terms to these two groups of borrowers. They paid the same mean interest rate, had the same loan-to-value ratio, and came from zip codes with the same median income. In other words, whatever information mortgage originators used to successfully screen out some low quality loans to the group of borrowers with a FICO score just below 620 was unavailable to the investors who purchased the MBS backed by loans to borrowers with a slightly higher credit score.

Finally, they document a similar discontinuity in the ease of securitizing full documentation loans, albeit at a slightly different threshold (a FICO score of 600). Once again, mortgage originators are much more likely to lend to borrowers just above the threshold than to those just below it. This reflects the fact that mortgage originators wish to avoid holding loans on their balance sheet due to the ensuing risk and capital requirements. Despite the difference in the amount of lending, they find that loans just above the threshold perform as well as those just below the threshold. The distinctive behavior of the low-doc market strongly suggests the importance of private information in that market, while there is a narrower scope for private information in the market for full documentation loans.

In a recent paper, Jiang, Nelson and Vytlacil (2011b) extend these results using proprietary data from a major, unidentified mortgage-origination bank. The particular bank specialized in broker-originated, low-doc, privately securitized lending. The data include all the information that the bank collected on all the loans it made from January 2004 to February 2008. They verify that the bank sold far fewer loans just below the "620" threshold (compared to just above it) with the density jumping about five-fold at the critical threshold. Moreover, the delinquency rate also jumps up at the threshold by about eight percentage points. Both of these results are consistent with the findings in Keys, Mukherjee, Seru and Vig (2010).

In contrast to the earlier paper, Jiang, Nelson and Vytlacil (2011b) are able to observe not just the ex-ante probability that the loan will be securitized, but also the ex-post outcome. Surprisingly, they find that securitized loans are actually less likely to default, even in a neighborhood at the critical "620" threshold. Their

interpretation lies in the timing. For example, a fraction of loans go delinquent immediately upon issue because the homeowner never makes a payment. The terms of the MBS do not allow such loans to be included in the security, and so the originator is left holding the loan. Similarly, investors may be able to select higher quality loans by using additional aggregate information that is revealed between the time of origination and securitization, such as the behavior of the local housing market. This indicates that symmetric lack of information at the time of origination works against the origination bank. Note, however, that the evidence in Jiang, Nelson and Vytlacil (2011b) does not speak against the possibility that the adverse selection problem remains in the market for MBS.

Demiroglu and James (2012) also look for evidence that mortgage originators have private information on low-documentation loans, but they use a different empirical approach. They look at how an originator's exposure to potential losses affects the quality of their loans. More precisely, some mortgage originators also securitize their loans to create MBSs. They then typically sell off the safest tranches to MBS buyers, but hold onto the riskiest (equity) tranches, exposing themselves to potential losses. Other originators sell their entire portfolio of loans to an unaffiliated mortgage securitizer, which performs the same function: it creates the MBS, sells off the safe tranches, and holds onto the riskiest ones.

The important difference is that when an originator has private information about loan quality but retains some exposure to losses through its affiliation with the securitizer, it may screen the loans more carefully. Demiroglu and James (2012) find evidence that these loans outperform loans where the originator and securitizer are unaffiliated. Low-documentation loans issued in affiliated deals, after conditioning on all the information available to the buyer of the MBS, are about twenty percent less likely to default than those issued in unaffiliated deals

Demiroglu and James (2012) offer another piece of evidence that private information is important in this market. They look exclusively at loans in which the originator and securitizer are unaffiliated, and so one would expect that the loans perform badly. They show that loan performance depends on whether the originator also services the loans, i.e. collecting and distributing the mortgage payments for a fee. A one standard deviation increase in the fraction of loans for which the originator is also the servicer implies about a 25 percent lower cumulative loss rate. Again, this suggests that originators screen loans more carefully if they expect to enjoy some of the benefits from the loans' performance.

Finally, Demiroglu and James (2012) show that neither of these results carries over to the full documentation loan market. The likelihood that a loan defaults depends neither on whether the originator is affiliated with the sponsor, nor on whether the originator services the loan. Since there is a narrower scope for the private information problem, this is again consistent with private information driving the results in the low documentation market.

3.2. Misrepresentation

Homeowners sometimes misrepresent important loan characteristics in their application. To the extent that the mortgage originator is aware of the misrepresentation, this creates another potential source of asymmetric information between the originator and the MBS buyer.

In the popular press, low documentation loans are often called "liar's loans," reflecting the temptation for a borrower to lie about his income and assets when these are not verified. Using the same proprietary data from the same mortgage originator as Jiang, Nelson and Vytlacil (2011b), Jiang, Nelson and Vytlacil (2011a) uncover evidence suggesting that income misreporting was pervasive in low documentation loans.

Low documentation loans were 5 to 8 percentage point more likely to go delinquent than full documentation loans, even after conditioning on all observable characteristics of the loan. This is not necessarily evidence of misreporting, since borrowers who select low documentation loans are less desirable. But Jiang, Nelson and Vytlacil (2011a) uncover two other pieces of evidence that strongly suggest misreporting.

First, they show that borrower information is much better for predicting the performance of full documentation loans than low documentation loans. One would expect this result if there were systematic misreporting for low documentation loans. Second, they show that for full documentation loans, higher income reduces the likelihood of default, but the opposite is true for low documentation loans. An increase in a borrower's self-reported income raises the likelihood that he defaults on his loan, an unexpected correlation if income is truthfully reported. Unfortunately, Jiang, Nelson and Vytlacil (2011a) do not have any direct evidence on borrowers' true income and so cannot definitively establish that borrowers in fact lied about their income.

A recent paper by Piskorski, Seru and Witkin (2013) uncovers direct evidence that borrowers misrepresented another important characteristic of their loan: whether the property is owner-occupied. Moreover, their paper suggests that the mortgage originator knew about some portion of this misrepresentation.[2]

Owner occupancy status is an important predictor of future default risk. This may be because owners place greater value on living in their dwelling than the market rent, while investors simply compute the option timing of default given expectations about future prices and rents, or it may be because investors are more financially sophisticated. In any case, after conditioning on a large number of other controls, Piskorski, Seru and Witkin (2013) use loan-level data on mortgages originated between 2005 and 2007 to show that when a borrower truthfully reports that that he does not intend to occupy a property, he is about 3.5 percentage points more likely to default than a borrower who truthfully reports that he intends to occupy the property.

They then turn to evidence of misrepresentation. They match their loan-level data to subsequent mailing addresses reported to a major credit bureau. If a borrower reports that he will occupy a property but does not move to the appropriate zip code at any time during the subsequent year, Piskorski, Seru and Witkin (2013) record him as a misrepresented non-owner occupant. This group comprises about 6.4 percent of reported owner-occupied mortgages. A misrepresented non-owner occupant is about 9.5 percentage points more likely to default on his loan than an owner-occupant after conditioning on the same large set of controls. In other words, a misrepresented non-owner occupant is a far worse risk than a truthfully reported non-owner occupant. Perhaps the fact that an individual is willing to lie on a loan application is a signal of his financial sophistication.

Piskorski, Seru and Witkin (2013) also uncover evidence that mortgage originators knew about some of the misrepresenting. Borrowers who truthfully reported their non-owner occupancy paid an interest rate that was 35

[2] They also uncover evidence that borrowers misrepresented whether there was a second lien on the property; however, they do not find that mortgage originators were aware of this misrepresentation.

basis points higher than owner-occupants, reflecting the increased risk of default. Borrowers who misrepresented their non-owner occupancy status also paid a higher interest rate, but only 23 basis points higher than owner-occupants. With the benefit of hindsight, it is clear that neither of these surcharges compensated for the subsequent higher default rates, which presumably reflects the general underpricing of risk during their pre-crisis sample period. But the fact that misrepresented, non-owner occupants paid a premium indicates that banks were able to partially distinguish them from owner-occupants. The fact that the premium is smaller than the one for truthful non-owner occupants, despite the higher default risk, suggests that the distinction was imperfect.

Investors who purchased MBS backed by a high percentage of misrepresentations were not compensated for the resulting low quality of the security. For example, the safest tranches of the securities were not protected by a greater amount of subordinated debt. Potentially, however, these investors may be protected by the MBS's warranty, depending on the outcome of pending court cases. Originators may have a reasonable defense by arguing that they simply asked borrowers to state whether they intended to occupy the home. Once the loan closed, there was little the originator could do to force the homeowner to move in. In any case, I further discuss both tranching and warranties in section 4 below.

3.3. Superior Valuation Models

A third potential source of asymmetric information is that mortgage originators may simply be better at valuing mortgages and the securities backed by them than MBS buyers are. This seems plausible because of gains from specialization: mortgage originators were in the business of giving mortgages, while valuing MBSs were comparatively unimportant for money market funds, pension funds, and insurance companies.

To my knowledge, there is no direct test of this hypothesis, but it is possible to look for indirect evidence. In the presence of this type of asymmetric information, one might expect mortgage originators to retain the best mortgages and securitize the worst, anticipating that the unsophisticated buyers would underprice quality. In fact, this is exactly the pattern in the data.

Krainer and Laderman (2013) look at mortgages originated between 2000 and 2007 for properties in California. Their main empirical results rely on estimates of a proportional hazard model for the risk of a loan going into default. After conditioning on other loan characteristics, they find that adjustable-rate loans that were privately securitized defaulted at a 13 to 16 percent higher rate than comparable loans that the originator retained, a statistically and economically significant difference. Curiously, they find no robust difference for fixed-rate loans. On the other hand, they find that adjustable rate loans which were privately securitized charged about 50 basis points lower interest than similar loans which the originator retained. That is, originators retained loans with a high interest rate and a low default risk and sold off loans with the opposite characteristics. It is difficult to understand why this would happen unless MBS buyers did not understand how to value the loans.

Downing, Jaffee and Wallace (2009) look at the agency MBS market where the main risk lies in early prepayment. In particular, a mortgage performs badly from the lender's perspective if it is prepaid early in an environment with lower-than-expected interest rates and if it is prepaid late in the opposite environment. They

study all the MBS issued by Freddie Mac Gold Participation Certificates from 1991 to 2002, a period well before the housing crash. These MBS were constructed in two stages. The first stage simply pooled the mortgages without creating tranches. The second stage tranched the assets to create Real Estate Mortgage Investment Conduits (REMICs) that were then resold in the private market. Their main result is that MBS that were converted into REMICs performed worse from the lender's perspective than MBS that were not converted; however, the spread is small, about four to six basis points. The source of this spread appears to lie in the originator's superior model of prepayment risk.

4. How Markets Deal with Private Information

Sophisticated investors understand that private information can be a problem when trading securities. Well-functioning markets therefore develop techniques to mitigate the impact of private information. In some sense, these techniques are the "dog that didn't bark": they provide indirect evidence that adverse selection must be an issue in securities markets since it would be hard to understand why these techniques would be employed if buyers and sellers had the same information. This section reviews a number of those techniques and explains how they help to mitigate adverse selection.

4.1. Warranties

The prospectus for an MBS summarizes a number of characteristics of the underlying mortgages and warrants the buyer against defects. More precisely, an MBS is administered by an independent third party, the trustee. The trustee has a specified amount of time, typically 90 days after the execution of the MBS, to uncover any material defects in the underlying loans. If the trustee uncovers such defects, the securitizer must either purchase the loan by paying off the principal and interest, or it must replace the loan with a similar asset.

This type of warranty is useful in the presence of asymmetric information. It reduces the incentive of the securitizer to misrepresent the characteristics of the securitized assets and mitigates the need of the buyer to look for evidence of such misrepresentation. In addition, the 90-day window suggests that detecting such defects is time-consuming and difficult, which is important for understanding both why warranties are useful and the extent to which they are limited. A warranty enables an unsophisticated buyer to quickly purchase an MBS, despite being unable to evaluate the accuracy of the underlying documentation.

In light of the evidence in Piskorski, Seru and Witkin (2013), it is important to note that warranties are restrictive. In particular, the prospectus for an MBS would typically limit the securitizer's responsibility for a borrower misrepresenting his intent to occupy the property. A typical prospectus states, "The sole basis for a representation that a given percentage of the loans are secured by single family property that is owner-occupied will be either (i) the making of a representation by the mortgagor at origination of the loan, either that the underlying mortgaged property will be used by the mortgagor for a period of at least six months every year or that the mortgagor intends to use the mortgaged property as a primary residence, or (ii) a finding that the address of the underlying

mortgaged property is the mortgagor's mailing address as reflected in the servicer's records."[3] It seems that courts will have to determine whether the evidence in Piskorski, Seru and Witkin (2013) establishes that securitizers engaged in fraud or whether the language in this type of clause applies. Still, this example suggests that warranties protected a MBS buyer against certain risks, but still left considerable scope for private information. Some of this had a significant impact on buyers' realized returns.

4.2. Credit Rating Agencies

MBS issued in the U.S. are typically rated by one of the three big credit ratings agencies (CRAs): Moody's, Standard & Poor's, and Fitch. These ratings serve at least two important roles. The first is regulatory: regulators forbid certain financial entities, such as money market funds, from holding any asset that does not have the highest credit rating; furthermore, under the Basel Accords, banks are required to hold more capital against assets with lower ratings. The second role is informational: CRAs specialize in offering independent, objective, and reliable assessments of an asset's quality. Of course, the two roles are linked. Financial regulators rely on CRAs because determining an asset's quality is difficult.

Since CRAs specialize in evaluating the quality of a security, it seems possible that they will actually have superior information to the buyer and seller of a MBS, thus entirely eliminating the adverse selection problem. But even if they do not have superior information, they can mitigate the extent of the private information problem, helping buyers distinguish between securities that might otherwise appear indistinguishable. As I discuss further in the theoretical section of this essay, anything that mitigates asymmetric information may help to facilitate trade in a security. Thus the fact that MBS are graded by the CRAs, and that regulators place weight on these ratings, suggests that private information may be pervasive in these markets.

During the financial crisis from 2007 to 2009, there was a widespread perception that the CRAs had not been providing independent, objective, and reliable assessments of assets' quality. For example, early in the financial crisis, investors started to realize substantial losses on AAA-rated Collateralized Debt Obligations (CDOs). CDOs are created by combining junior tranches of MBSs. The models employed by the CRAs assumed that the losses suffered by the MBS would be uncorrelated, which meant that the senior tranches of CDOs would be nearly risk-free. This assumption turned out to be incorrect, resulting in massive losses. To date, nearly 40 percent of AAA-rated CDOs have suffered some losses. By contrast, Moody's idealized expected loss rate over a 5-year period for an AAA security is 0.0016 percent.

Whether the CRAs provided independent, objective, and reliable assessments prior to the financial crisis remains in dispute.[4] What is indisputable is that their reputation was damaged by the crisis, and that no third party could

[3] This quote comes from Bear Stearns Asset Backed Securities Trust 2006-1, form 424(B)(5). Similar language was used in many other prospectuses.

[4] See Foote, Gerardi and Willen (2012), especially their fact 12. They point out that among all the residential MBS and home equity loans that Moody's had originally rated AAA, about 15 percent were impaired —suffered losses or had been downgraded to junk status— by the end of 2011 (Moody's Investors Services, 2012). While 15 percent impairment is much higher than one would normally expect from a AAA security, these losses occurred during the deepest recession since the Great Depression. Indeed, it seems that conditional on the size of the national house price decline, the models used by the

immediately step in to provide their traditional services. The loss of these key players at a critical juncture exacerbated private information problems and contributed to the collapse in financial intermediation during this period.

4.3. Reputation

In markets in which a seller's private information is revealed slowly over time, a seller may obtain a reputation for truthfully revealing the quality of its product. Moreover, if a particular buyer frequently purchases assets from a particular seller, the value of future interactions may guarantee that the seller wishes to ensure the buyer's survival, relaxing any incentive constraints.

The market for MBSs was certainly a small market. Relatively few large banks provided most of the mortgages to a small number of securitizers, which in turn sold many of the MBS to large institutional investors. For example, in 2006 the three largest "Alt A" lenders accounted for over 45 percent of the market,[5] while the six largest underwriters accounted for 53 percent of private-label mortgage purchases and the six largest MBS issuers accounted for 43 percent of the private-label market.[6] This means that the opportunities for sellers to obtain and maintain a reputation for honesty were abundant.

Unfortunately, important forces conspired against sellers obtaining a good reputation before the financial crisis and maintaining one during it. Prior to the crisis, mortgage originators and securitizers may have correctly perceived that the boom in mortgage issuance would be temporary. A good reputation is hard to sustain when the short-run profits from exploiting buyers are large relative to the long-run loss from a bad reputation, exactly the situation during a temporary boom. And then during the financial crisis, many originators and securitizers experienced bankruptcy or a forced sale to a competitor, while all surviving financial intermediaries were concerned that these undesirable outcomes were possible. Again, intermediaries may be tempted to boost short-run cash-flow if that significantly increases their survival probability, even if doing so significantly reduces their long-run value through a loss in reputation and the ensuing bankruptcy of their clients.

As was the case with CRAs, the usefulness of financial intermediaries' reputation as a device that ensured truthful revelation of information collapsed at a key juncture during the financial crisis. This meant that buyers had to be more aware than ever of sellers taking advantage of short-run profit opportunities. The MBS market was therefore rife with private information.

4.4. Tranching

A securitizer originates or purchases a large number of loans, combines them into a single pool, and then issues

---

CRAs correctly forecast the losses suffered by MBS, although they severely underestimated the size of the house price decline.

[5] See Inside Mortgage Finance (2011a) p.161.

[6] See Inside Mortgage Finance (2011b) p. 39.

MBSs backed by the revenue stream coming into the pool. The securitizer creates several different securities from a single pool of mortgages distinguished, in large part, by their seniority. For a private-label bond, about 80 percent of the notional stream of interest and principal payments are typically promised to the most senior tranche. This means that if 40 percent of the homeowners default on their mortgage payments and the mortgage holder is able to recover half the value of those loans by selling the collateral, the holders of the most senior tranche will still receive the full promised coupon payment. Indeed the size of the most senior tranche is typically set so as to ensure that it receives the highest (AAA) rating from the CRAs, with enough collateral and buffer from junior tranches to protect bondholders against losses in any likely scenario.[7] The next few percent of the notional income stream is then promised to a more junior investment-grade bond. That bond experiences losses before the AAA-rated tranches, but is still buffered against losses by lower-grade securities. Finally, the securitizer typically retains the rights to the marginal income streams, often called the equity tranche.

Widespread tranching is evidence that private information is an issue in this market. Tranching divides a stochastic stream of payments into several assets, ranging from risk free debt to levered equity. The buyers of the risk-free debt, in turn, are protected against the need to understand the stochastic process of the underlying payment stream. In the terminology of Gorton and Pennacchi (1990), risk-free debt is "information insensitive." As long as a buyer is (almost) certain that a MBS will pay off at face value, he does not need to understand the risks to the mortgage pool. In contrast, equity is information sensitive since it absorbs all the variation in the payment stream. This means that if a securitizer has superior information about the quality of a mortgage pool, tranching allows the securitizer to sell much of the payment stream without encountering private information problems. The most senior tranches are safe and, hence, private information does not distort their sale. The securitizer then retains the junior tranches with the associated risk and information problems.

Still, the ability of tranching to mitigate private information problems is limited by the underlying amount of risk. If there is a chance that the stream of payments to a mortgage pool will dry up completely, then it is impossible to create any risk-free debt from the promised revenue stream. Risk in itself is of course not a problem for financial markets. The problem is that a seller may have superior information about the stochastic process for the revenue stream. Indeed, as noted by Dang, Gorton and Holmström (2012), once debt is risky, investors have an incentive to acquire information about the nature of the risk. "The crisis is not just the bad shock about fundamentals that back debts. Instead, the crisis is a bad enough shock to cause information-insensitive debt to become information acquisition sensitive."[8] The crisis is the emergence, or threat of emergence, of private information in a market that was previously immune from this problem.

In summary, tranching represents an attempt to turn a risky stream of income into a safe bond, in part, to suppress private information problems. The prevalence of tranching in the MBS market therefore suggests that

[7] In the aftermath of the financial crisis, these buffers were insufficient and some AAA-rated MBS did default on payments (see footnote 4). Still, AAA-securities performed significantly better than lower-rated securities, with 58 percent of AA-rated MBS impaired by 2011 (Moody's Investors Services, 2012)

[8] See Dang, Gorton and Holmström (2012) p. 32.

private information may be relevant to those markets. Moreover, a worsening of the left tail of the income stream lowers the maximum amount of risk-free debt that can be created. As the supply of risk-free debt disappears, either securitizers must stop creating MBS or the MBS market must deal with the existence of private information.

4.5. Repurchase Agreement Haircuts

A repurchase agreement (repo) consists of the sale of a security together with the promise to buy it back at a specified date and price. In other words, in its simplest form, a repo is a collateralized loan. For example, a repo seller gives a repo buyer an MBS in return for some cash, then the contract specifies that the seller must repay the cash with interest at a later date in return for the securities. The haircut in a repo contract is defined as one hundred percent, minus the ratio of the cash lent by the repo buyer for the market price of the securities lent by the repo seller. In other words, if the repo seller receives $70 in cash in return for $100 in securities, the haircut is 30 percent.

For my purpose, the relevant aspect of a repo agreement is how it treats a default. Repos are governed by a Global Master Repurchase Agreement, which contains detailed rules following a default. To be concrete, suppose the repo seller does not repay the repo buyer on the specified date. At this point, the repo buyer is permitted to sell or retain enough of the MBS so as to compensate himself for the lost payment, returning the rest to the repo seller.

It is in this instance that the repo haircut is important. If the haircut is sufficiently large (i.e. the loan is sufficiently over-collateralized), then the repo buyer does not need to worry about the value of the collateral. In the event of default, the repo buyer simply sells the collateral to make up for the lost payment, ensuring that the loan is risk-free. This means that if the repo buyer is concerned that the repo seller has superior information about the value of the collateral, he can simply demand a larger haircut to protect himself against the risk of loss.

These arguments imply that repo haircuts are effectively equivalent to tranching.[9] A repo seller has private information about the quality of a risky asset that it owns. By demanding the entire asset as collateral against a relatively small loan, a repo buyer is protected against the seller's default, even if he has little information about the collateral's quality. The repo buyer therefore effectively purchases the senior tranche of the asset, while the repo seller is left holding the residual (i.e. the equity tranche). In short, a repo haircut effectively circumvents the seller's private information, allowing the repo seller to sell a safe stream of income secured by a risky asset.

Conversely, it is difficult to understand why repo haircuts would exist in an environment with symmetric information. For example, if MBS are risky but the buyer and seller have the same information about their value, the market price of the MBS would compensate for the risk. A haircut would shift the risk to the seller, but since there is no general reason to believe that repo sellers are better at bearing risk than repo buyers, this cannot give a satisfactory theory of repo haircuts.

Gorton and Metrick (2010) show that haircuts increased as the financial crisis worsened. In the first half of 2007, there was no haircut on subprime-related structured products. By the time Lehman Brothers collapsed in September

[9] The notion that a repo is equivalent to tranching comes from Gorton and Metrick (2010). That paper stresses a different type of private information that the repo buyer may be less well informed about than a potential trading partner in the secondary market. For the purpose of this paper, it is enough to note that private information was likely prevalent in the MBS market.

2008, the haircut had increased to 100 percent (i.e. the products were worthless as collateral). This suggests again that, although private information problems were suppressed prior to the crisis, they emerged in the private label MBS market as the underlying income streams dried up.

5. MODELING MARKETS WITH PRIVATE INFORMATION

The previous sections offered a subtle and nuanced view of the sources of private information in the MBS market. This section describes a stark and abstract framework for analyzing how private information can lead to illiquidity in financial markets. The analysis here is based on Guerrieri and Shimer (2013a) and Guerrieri and Shimer (2013b), and I refer the reader to those papers for a formal treatment of these ideas.

The basic idea is that illiquidity acts as a costly signal of an asset's quality. The notion of "costly signals" dates back to Spence (1973) in the context of school enrollment. Gale (1996) first proposed that illiquidity could serve as a costly signal. Illiquidity is costly because there are gains from trade, and so the failure to trade imposes a cost on the seller. Illiquidity can serve as a signal because the costs depend on the asset's quality. If an asset does not sell, the seller is left holding it, which is more costly to the seller if the asset is of lower quality. This means that any observable action that a seller takes to make an asset illiquid can serve as a useful signal to potential buyers about the asset's quality.

DeMarzo and Duffie (1999) propose that a seller may commit to hold onto a fraction of a stream of payments to signal its quality. Even if the senior tranche remains risky, the willingness to retain an equity tranche is a useful signal about the quality of a mortgage pool. While this security design literature offers important insights into how markets cope with private information, it leaves one important question unanswered: how can a seller credibly commit to hold onto the equity tranche? After selling the senior tranche, there are still gains from selling the equity tranche. But if the initial buyer knows that the seller can do this, then holding onto the equity tranche is no longer a costly signal. Put differently, retaining a portion of an asset as a signal of its quality leads to a classic time-inconsistency problem.

I propose that in addition to tranching, markets use a different costly signal: the seller's ask price for a security. I construct a market economy in which sellers can potentially sell a security at a range of different prices, but are more likely to sell it at a lower price. Demanding a high price therefore incurs the potential cost of illiquidity. Sellers who value an asset less, either because of the characteristics of the seller or the characteristics of the asset, will set a lower price and sell their security with a higher probability. Buyers understand this and are willing to pay a higher price only to the extent that they expect that this will give them a higher quality security.

This insight leads to a theory of how illiquidity arises naturally in a market economy, even without any commitment. If it were too easy to sell a security at a high price, then the holders of low quality securities would offer them at a high price, driving away the buyers. The equilibrium I construct has the minimum amount of illiquidity required to induce sellers to offer different prices depending on their value of holding onto their security. I turn next to a detailed description of the environment to flesh this idea out.

5.1. Model

There are many investors. The number of investors is large in the sense that each of them believes that they cannot, acting individually, change the nature of the illiquidity problem. There are two assets in the economy: "cash" and "mortgage-backed securities." Initially each investor holds some cash and some MBSs. Cash is homogeneous, but MBSs are of heterogeneous quality or payoff. The private information problem is that only the initial owner of the MBS knows its quality. For example, the initial owner may be the securitizer, with all the informational advantages described in section 3. I denote the quality of an MBS by $\delta \in (\underline{\delta}, \overline{\delta})$.

Investors can trade cash for an MBS and receive a payoff that depends on their final cash and quality-adjusted MBS holdings. All investors are risk-neutral and I normalize the marginal utility of quality-adjusted MBS holdings to 1 for each investor. I allow different investors to have a different marginal utility (or value) of cash, which I denote by $\alpha \in (\underline{\alpha}, \overline{\alpha})$. For example, pension funds may have an ample cash flow with few direct investment possibilities, while securitizers can use the cash to purchase more MBSs (in some un-modeled market). In this case, I would expect α to be low for a pension fund and high for a securitizer.

The difference in the value of cash is critical because it implies that there are gains from trade. If all investors had the same value of cash and rational expectations, then there would be no gains from trade and hence, no trade (Milgrom and Stokey, 1982). The only reason one investor would be willing to sell a security at a particular price is if he believed it to be worth less than that price. But every other investor should understand this and so be unwilling to purchase the security at that price. All trade would cease. The realistic assumption that there are some intrinsic gains from trade means that it may be possible to construct an equilibrium with trade. The interesting question is how private information affects the amount of trade.

I look at two versions of the model. In the first, an investor's value of cash α is observable (Guerrieri and Shimer, 2013a), but the quality of the investor's MBS is not. For example, it is possible to distinguish a pension fund from an investment bank. In the second, both the value of cash and the quality of MBSs are the investor's private information (Guerrieri and Shimer, 2013b). For example, some investment banks may have better investment opportunities than others at any point in time. While outcomes are fairly similar in the two environments, I highlight some important differences in sections 5.2 and 5.3 below. For expositional purposes, my description of the environment focuses on the (notationally simpler) case in which both an investor's value of cash and the quality of its MBS is its private information.

I assume that investors can simultaneously sell their MBSs for cash, and use that cash to buy other MBSs. They do this by setting an ask price for each MBS, and a bid price to purchase other MBSs for cash. When they do this, two considerations are paramount: First, an investor anticipates that he can raise the probability of finding a buyer by reducing his ask price. Let $\Theta(p)$ denote the probability that a seller finds a buyer if he sets an ask equal to p per unit of MBS. Since it is a probability, impose $0 \leq \Theta(p) \leq 1$ for all p.[10] Second, an investor anticipates that he can

[10] In the full model, buyers also anticipate that they can raise the probability of finding a seller by raising their bid. In equilibrium, however, buyers' bids are always satisfied, and so I ignore that issue for expositional simplicity.

raise the quality of the MBS that he buys by raising his bid price. Let $D(p)$ denote the expected payoff from a unit of MBS purchased at a price p. The functions Θ and D are not arbitrary but must be consistent with investors' optimization, with rational expectations, and with markets clearing, as I explain in the following paragraphs.

First consider the optimal ask price for an investor who values cash at α and who has a MBS with value δ. This satisfies

$$P_s(\alpha,\delta) = \arg\max_p \Theta(p)(\alpha p - \delta) \tag{1}$$

If he succeeds in selling the MBS, he gets cash p that he values at α. Otherwise he retains the MBS, which he values at δ. Therefore, $\alpha p - \delta$ represents the gain from selling, which occurs with probability $\Theta(p)$.

Similarly, consider the optimal bid price for an investor who values cash at some α. This satisfies

$$P_b(\alpha) = \arg\max_p \left(\frac{D(p)}{p} - \alpha \right) \tag{2}$$

He values each MBS that he buys at price p at $D(p)$ and a unit of cash allows him to buy 1/ p of this security. His opportunity cost of each unit of cash is α. The bid price must solve this maximization problem given buyers' beliefs encapsulated in $D(p)$. If, however, α> $D(p)/p$ for all p, then an investor with cash value α will not bid for MBSs.

A word on notation is in order. The notation in the previous two paragraphs imposes that all investors with the same cash value and same quality asset set the same ask price, and all investors with the same cash value set the same bid price. I could easily relax this assumption and allow identical investors to sell identical assets at different prices, and similarly for buying. For example, I could allow an investor to offer one MBS at a high price and an identical MBS at a low price. Although such an environment would be notationally cumbersome, one can verify that relaxing this assumption would not affect the equilibrium.

In equilibrium, the functions Θ and D must be consistent with rational expectations. Broadly speaking, there are two types of prices p, those that are an ask price for some seller (formally, there is a (α,δ) such that $P_s(\alpha,\delta) = p$) and those that are not. If there are sellers with ask price p, then $D(p)$ is the average quality of the MBS offered by those sellers and $\Theta(p)$ reflects any shortfall of buyers bidding that price; if the amount of cash that buyers use to purchase MBS at price p exceeds the amount of MBS offered for sale at that price times the price, then all sellers are satisfied and $\Theta(p)$ =1; otherwise $\Theta(p)$ is the ratio of the buyers' total bids to the cost of the sellers' total asks at price p.

Next consider prices p that are not an ask price for any seller. In this case, there are two possibilities. First, it may be the case that, even if $\Theta(p)$ =1, no seller would find it optimal to offer this price. In this case, buyers believe they cannot buy at this price, or equivalently they would only be able to buy a worthless asset, $D(p)$ = 0. Second, it may be the case that there is some $0 \le \Theta(p) \le 1$ at which one or more seller is indifferent about this price or his ask price, while all other sellers prefer their ask price. In this case, sellers believe that if they set this ask price, this is the sale probability; buyers believe that if they set this bid price, they would purchase some combination of the assets

offered by the sellers who are indifferent about offering this price.

The previous paragraph describes restrictions on beliefs that seem reasonable in this environment. If no seller sets ask price p, then a buyer should think about which sellers would be most willing to be rationed at that price. He should then anticipate that if he bids that price, only these types of sellers would set that ask price. But if no seller is willing to accept that price, even if there is no rationing as may be true at a very low price, then buyers should anticipate that they would be unable to buy anything at that price.

These restrictions on beliefs reduce the set of possible equilibria and so lead to strong predictions about the nature of equilibrium. Such restrictions are necessary because this is a signaling game, and signaling games typically have multiple equilibria. Different equilibria impose different assumptions on how one economic agent interprets the signals sent by another. Here, prices are signals and buyers must interpret which seller asks which price. For prices that are asked in equilibrium ($p = P_s(\alpha, \delta)$ for some (α, δ)), are pinned down by rational expectations. For other prices, these beliefs are potentially arbitrary. For example, buyers might choose not to bid a particular high price because they believe that only low-quality assets are offered at that price; sellers do not ask that price because it is impossible to find a buyer at that price. Therefore, the buyers' beliefs are never invalidated. I preclude this particular belief through the assumption that buyers believe that each price would be asked by the seller who is most willing to be rationed at that price. This helps to discipline what can happen in equilibrium in a reasonable way.

In closing, I discuss the critical assumption in this pricing game: an ask price represents a commitment to buy at that price. This means that an investor cannot ask two prices for one security, selling it at the higher price if he manages to find a buyer, or otherwise, selling it at the lower price.[11] If he succeeds in selling at the higher price, he would not be able to fulfill his commitment to sell it at the lower price. The question is what real world institution this model captures. MBSs are sold over-the-counter (OTC) and the price is determined by bargaining, something that is absent from this model. Still, the basic economic forces in this model are likely to be present in an OTC market. The model has bilateral asymmetric information since the seller knows both his cash value and the quality of the security, while the buyer knows his cash value. In such an environment, a seller may correctly perceive that a high ask price signals to a buyer that she is not too motivated to sell. This may reduce the chance that the seller ultimately transacts with the buyer but raise the price if trade does take place. I do not model such a bargaining game because there is generally no agreement on how to model bargaining with bilateral asymmetric information. Instead, I expect that the economic forces that I identify in this competitive framework are also relevant in decentralized markets.

[11] As previously noted, I can allow the investor to ask a high price for one tranche of his security and a low price for the remainder, without affecting the equilibrium allocation. In addition, it is straightforward to extend the model to allow for multiple rounds of trading; however, with no opportunity cost of delay, all trading must take place in the final round and so the equilibrium is unchanged.

5.2. Observable Value of Cash

I start by describing the equilibrium of the model in which sellers' value of cash α is observable (Guerrieri and Shimer, 2013a).[12] There exists a unique equilibrium outcome in this environment. In it, investors are endogenously partitioned into two groups, sellers and buyers, at a critical threshold for the value of cash, $\alpha = \hat{\alpha}$. Sellers are investors with a high value of cash, $\alpha > \hat{\alpha}$. They attempt to sell all of their MBSs and do not buy MBSs (or equivalently, set a bid price at which MBSs are unavailable). Their ask price is proportional to the quality of their MBS, while the sale probability depends on both the seller's value of cash and the quality of the seller's MBS. Buyers are investors with a low value of cash, $\alpha < \hat{\alpha}$. They are unable to sell their MBS but use all their cash to buy MBSs. All buyers are indifferent about purchasing from any seller at any bid price.

In equilibrium, a seller with the worst quality MBS, $\delta = \underline{\delta}$, is able to sell his security with probability 1, while other sellers are rationed. If the lower bound of the quality distribution is positive, $\underline{\delta} > 0$, rationed sellers trade with a strictly positive probability but not with certainty. The probability of sale is lower when the seller knows that his security has a higher quality (δ is higher), when the worst quality MBS is lower ($\underline{\delta}$ is lower), or when the investor values cash less ($\alpha > \hat{\alpha}$ is smaller) but $\delta > \underline{\delta}$. This reflects the fact that the sale probability must prevent the sellers of lower quality MBSs from misrepresenting them as being of higher quality. Illiquidity is more effective as a separating device when investors value cash more, and so less illiquidity is required in that case.

Figure 2 illustrates these points by showing the sale probability as a function of price for two sellers with a different value of cash. They coincide at the price of the lowest quality MBS, but at higher prices, the seller with a higher value of cash sells with higher probability. This reflects his relatively flat indifference curve. Even a small reduction in the sale probability compensates for an increase in the price when a seller has a high value of cash. Indeed, the sale probability is the lower envelope of the indifference curves of all sellers with that value of cash but different quality MBSs, and therefore is also flatter when the seller has a higher value of cash.

---

[12] The model setup is slightly different because the two key equilibrium objects Θ and D are functions of both the price and the seller's value of cash. The ask price of a seller (α, δ) satisfies

$$P_s(\alpha, \delta) = \arg\max_p \Theta(p)(\alpha p - \delta) \quad , \tag{1}$$

recognizing that the sale probability depends on both its ask price and its type. This is an immediate extension of equation (1). Buyers choose both a bid price, p, and the type of seller that they buy from, α_s, to solve

$$P_b(\alpha) = \arg\max_{p,\alpha_s} \left(\frac{D(p, \alpha_s)}{p} - \alpha \right),$$

where $D(p, \alpha_s)$ is the average quality asset sold at price p by a seller with cash value αs. Here $P_b(\alpha)$ denotes the price that a buyer with cash value α pays and $A_b(\alpha)$ denotes the seller's type. The remainder of the setup is common across the two models.

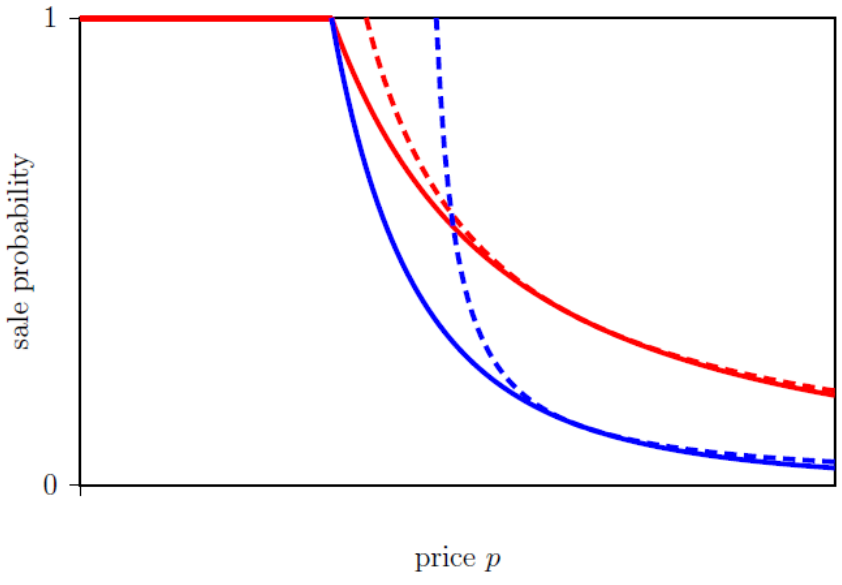

Figure 2: The solid curves illustrate the sale probability for two different investors. The red (flat) one has a higher value of cash than the blue (steep) one. The dashed curves illustrate the indifference curves of two different investors. Both have the same quality asset, but the red (flat) one has a higher value of cash than the blue (steep) one. Both investors set the same price, but the investor with the higher value of cash sells with a higher probability.

The equilibrium allocation is sensitive to the support of the MBS quality distribution. A reduction in $\underline{\delta}$ reduces the sale probability for all MBSs conditional on the marginal investor $\hat{\alpha}$. In particular, if $\underline{\delta} = 0$, there is no trade in any MBS with positive quality, $\Theta(p) = 0$ for all $p > 0$. This means that the equilibrium allocation depends on the exact specification of the model. For example, suppose there is a negligible probability that $\delta < \hat{\delta}$ for some $\hat{\delta} > \underline{\delta}$. Markets that disregard this possibility will have much more trade than markets that recognize the small chance that $\delta < \hat{\delta}$.

It is worth noting that all sellers with an MBS $\delta > \underline{\delta}$ wish they could mislead other investors into believing that they have a higher value of cash α. This would raise their sale probability without affecting the price. The assumption that buyers can observe sellers' value of cash therefore matters for the structure of equilibrium. I turn next to the other case.

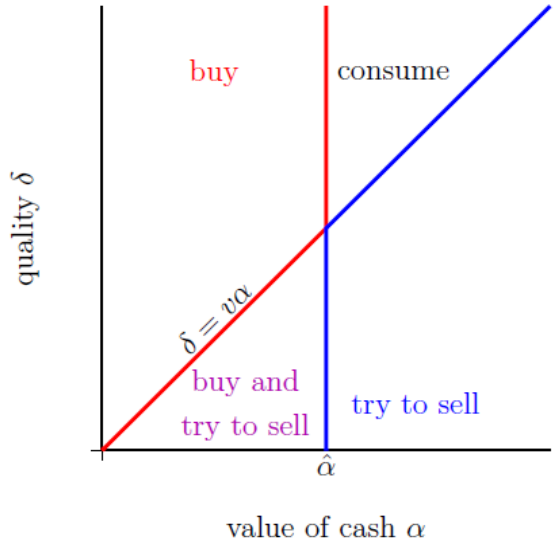

Figure 3: Behavior in partial pooling equilibrium.

5.3. Unobservable Value of Cash

When investors' value of cash is private information, the structure of equilibrium is necessarily different (Guerrieri and Shimer, 2013b). In this case, a continuum of equilibria may exist. In any equilibrium, there is again a critical threshold $\hat{\alpha}$ such that any investor with a lower value of cash, $\alpha < \hat{\alpha}$, uses all his cash to buy any MBS, and any investor with a higher value of cash does not purchase MBSs. But this threshold no longer determines whether an investor sells his MBS. Instead, this depends on the ratio of the quality of his MBS to his value of cash, δ / α. When this ratio is low, an investor sells his MBS at a low price with a high probability. As this ratio rises, the price rises and the sale probability falls. Finally, at some critical value $\delta / \alpha = v$, the sale probability falls to zero and the investor no longer attempts to sell his MBS. Figure 3 illustrates the equilibria outcomes. Investors with a low quality MBS and a low value of cash use their cash to buy MBSs while simultaneously attempting to sell their MBSs. If their MBS quality is higher, they do not attempt to sell, while if their cash value is higher they do not buy. Finally, investors with both a good quality MBS and a high value of cash simply consume their endowment.

The equilibrium has "partial pooling" in the sense that all investors with the same value of δ / α set the same price for their MBS and sell with the same probability. This follows from equation (1), which indicates that if one investor (α, δ) finds ask price p optimal, then any other investor $(\lambda\alpha, \lambda\delta)$ with $\lambda > 0$ finds the same ask price weakly optimal. On the other hand, investors with a higher value of δ/α are more willing to accept a reduction in the sale probability in return for an increase in the price and so, send the noisy signal of a high price in equilibrium.

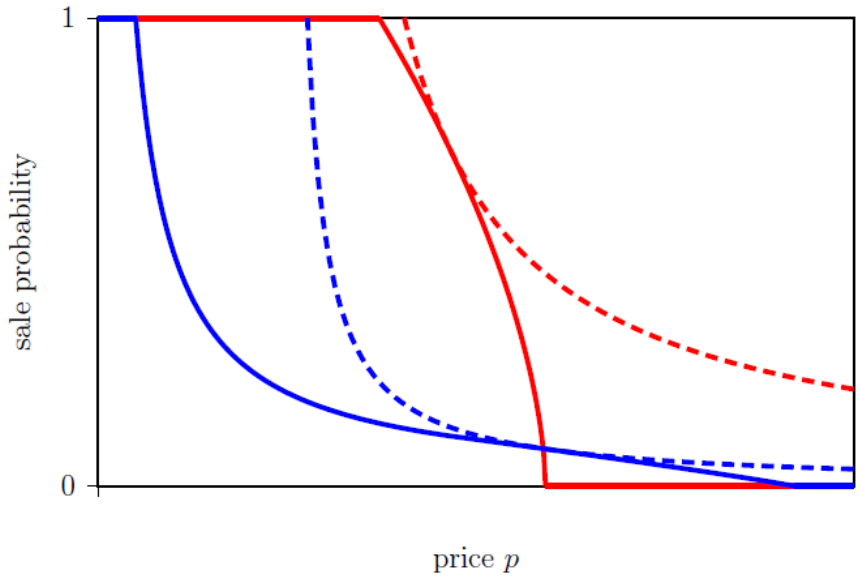

Figure 4: The solid curves illustrate the sale probability in two different equilibria. In the red (steep) equilibrium, the value of cash for the marginal buyer is higher than in the blue (flat) equilibrium, and so there are more buyers. The solid curves indicate the indifference curve of an investor with a particular value of δ/α. The investor chooses a higher price and has a lower sale probability in the blue (flat) equilibrium where are buyers are more scarce.

The solid lines in figure 4 illustrate the equilibrium sale probability as a function of the price in two different equilibria. The dashed lines indicate the indifference curve of one particular seller. The seller is on a lower indifference curve (in one equilibrium) than the other because of a shortage of buyers. Once again, in each equilibrium, the sale probability is the lower envelope of all sellers' indifference curves.

The figure does not illustrate buyers' indifference curves. Buyers' behavior is similar to the model with observable cash values. In any equilibrium, higher expected quality exactly compensates for higher price; therefore, buyers are willing to purchase at any price. The only subtle issue is that different types of sellers set a common price, so buyers do not know exactly what type of MBS they will purchase at each price. Still, under the risk-neutrality assumption, only the expected quality matters.

The set of equilibria now depends in an intricate way on the entire joint distribution of cash values and MBS quality. To be concrete, suppose that the distribution of the value of cash in the population is *Pareto* with parameter $a > 0$, so a fraction $1-\alpha^{-a}$ of the population value cash less than any level α> 1. Also suppose that the fraction of MBSs with quality less than $\delta \in [0,1]$ is δ^d for some parameter $d > 0$. Finally, suppose that the quality of an MBS held by an investor is independent of the investors' value of cash.

For all parameters a and d, Guerrieri and Shimer (2013b) prove that there exists an equilibrium with no trade: $\Theta(p) = 0$ for all $p > 0$. Buyers believe that any seller who is willing to sell at any positive price has an MBS with quality less than p. Since every investor values cash more than MBSs, no one is willing to purchase MBSs.

In addition, if $d > a$, so there are few low quality MBSs relative to the number of investors with a high value of cash, there is a continuum of equilibria with trade. In any equilibrium, the relative illiquidity for two investors

depends not just on their value of cash and MBS quality, but also on the distributional parameters a and d. These parameters affect buyers' perception of the average quality of MBSs available for sale at any price and affect buyers' willingness to pay a marginal increase in the price. In contrast, in the model with observable seller characteristics, illiquidity for any particular seller depended only on the minimum MBS quality $\underline{\delta}$.

Another important difference between the two models is that with observable seller characteristics and $\underline{\delta} = 0$, all trade breaks down. With unobservable characteristics, trade may continue to take place, even in this case, as the concrete example above shows. This seems like an attractive feature. Trade can occur even in markets in which a seller can have arbitrarily bad information about the quality of his MBS. A necessary condition for trade is that extremely motivated sellers are more likely, in some sense, to have an extremely high value of cash, rather than an extremely low quality MBS.

Finally, in the model with observable investor characteristics, there is a neat partition between buyers ($\alpha < \hat{\alpha}$) and sellers ($\alpha > \hat{\alpha}$). With unobservable characteristics, the set of buyers is qualitatively unchanged (although of course the threshold $\hat{\alpha}$ will, in general, be different). However, the set of sellers changes so that any investor with δ / α below a critical threshold v attempts to sell his MBS. This implies in particular that some investors are engaged only in buying, some only in selling, some do not participate in markets, and still others both buy and sell assets, as shown in figure 3. Such trade may be inefficient in the sense that an investor with a low value of cash may sell to an investor with a somewhat higher value of cash, reducing aggregate welfare.

5.4. Two Theories of Breakdown in Financial Markets

The model offers two mechanisms through which trade in financial markets can break down. The first is a change, or the perception of a change, in the joint distribution of the model's fundamentals: investors' cash value α and securities' quality δ. The second is a change in equilibrium for a given joint distribution. I describe how each of these crises might look in turn.

I start with a shift in the joint distribution of fundamentals. In the pre-crisis environment, the AAA-rated tranche of a mortgage pool is perceived to be riskless. All investors believe that the promised coupon will be paid with say, $\delta = 1$ certainty. Whether this is exactly correct or not is unimportant. Sellers do not pay attention to trivial risks in their MBS portfolio, so buyers can neglect this potential issue as well. This means that a mortgage securitizer can quickly purchase a pool of mortgages, tranche them, and sell off most of the earnings, giving himself enough cash to repeat the process.

The crisis begins with a decline in house prices. This has two effects. First, homeowners start to default at higher rates. Previously safe assets become risky and so information-insensitive debt becomes information *sensitive*. Buyers become aware that sellers may have private information about the quality of securities and scale back their demand appropriately. Since fewer securities are sold, and the securities that are sold, sell at lower prices reflecting the default risk, there must be less cash in the market. That is, some investors stop buying and the marginal buyer has a lower cash value. And since the marginal buyer prices the securities, the reduction in the marginal buyer's cash value implies that there is offsetting upward pressure on MBS prices.

The extent of the increase in MBS prices is dictated by the distribution of buyers' cash value and by the ability of buyers to substitute by purchasing other securities. In the model described here, buyers can only invest in MBSs, but in reality, buyers can easily purchase other assets, such as treasury bonds. In this case, the crisis will also cause a flight to quality, with buyers' excess cash driving up the price of treasury bonds.

The second effect of the reduction in house prices is to reduce the value of investing in new mortgage pools, and hence the value of cash to mortgage securitizers. If the value of cash is observable, this directly reduces the liquidity of the securities market, as shown in figure 2. If it is unobservable, the thinning of the right tail of the distribution of α means that buyers anticipate getting a lower quality asset conditional on the value of δ / α and therefore willing to pay less. The simultaneous reduction in the demand and supply of MBS has an ambiguous impact on prices but, to the extent that prices rise, further accelerates the flight to quality.

If under the new distribution $\underline{\delta} = 0$, then all trade may break down in a crisis. More generally, some trade may continue to take place, but at depressed prices and liquidity relative to the pre-crisis environment. Securitizers retain a greater share of the mortgage pools and ultimately the availability of new loans dries up.

The second type of crisis can arise only when the value of cash is unobservable. There is no change in fundamentals, just a reduction in the number of investors who use their cash to purchase securities. This means that the value of cash to the marginal buyer is lower, driving up MBS prices. If buyers are able to purchase other securities, they will do so. This means that the crisis can again generate a flight to quality. This second type of crisis is accompanied by a shift in the illiquidity function Θ, as shown in figure 2. This illiquidity induces sellers to willingly charge higher prices, despite the collapse in the probability of trade, because the percentage reduction in the sale probability for an increase in price is smaller.

This type of crisis looks like a buyers' strike. Nothing fundamental has changed in the market, but buyers have disappeared. Still, sellers are unwilling to cut prices because this has little impact on their sale probability. Indeed, they take advantage of the inelastic demand to raise prices. Sellers of course would be willing to charge the old price if they could sell with the old probability, but that is simply not possible any more.

What can cause this second type of crisis? A full answer goes beyond the scope of the model. One possibility, however, is that the buyers' strike is caused by a collapse in another related market. For example, a shift in fundamentals may cause a collapse in the market for private-label MBS. A shift in equilibrium may then cause a similar collapse in the agency MBS market or the market for securities backed by automobile loans.

6. CONCLUSIONS

If a crisis in the MBS market simply slowed the sale of MBSs from securitizers to pension funds, it would not matter for ordinary individuals who are not engaged in financial intermediation. But the crisis matters for them because when the intermediation chain breaks down, lending breaks down as well. Securitizers will not buy mortgage pools if they anticipate it taking too long to sell the MBSs. Originators will be reluctant to make loans if they anticipate that they will have to hold the loans in their portfolios since the capital requirements on these undiversified risks are high.

And so in the end, the perception that securitizers have private information hurts potential homeowners who wish to buy their first home, and existing homeowners who wish to refinance or move to a bigger home. The reduced demand for new housing hurts construction workers who cannot find jobs as residential investment collapses. A crisis in the MBS market causes a widespread decline in individuals' well being.

REFERENCES

Dang, T.V., G. Gorton and B. Holmström. 2012. "Ignorance, Debt and Financial Crises." Unpublished manuscript, Columbia University.

DeMarzo, P. and D. Duffie. 1999. "A Liquidity-Based Model of Security Design." *Econometrica* 67(1): 65–99.

Demiroglu, C. and C. James. 2012. "How Important is Having Skin in the Game? Originator-Sponsor Affiliation and Losses on Mortgage-backed Securities." *Review of Financial Studies* 25(11): 3217–58.

Downing, C., D. Jaffee and N. Wallace. 2009. "Is the Market for Mortgage-backed Securities a Market for Lemons?" *Review of Financial Studies* 22(7): 2457–94.

Foote, C.L., K.S. Gerardi and P.S. Willen. 2012. "Why Did So Many People Make So Many Ex Post Bad Decisions? The Causes of the Foreclosure Crisis." Public Policy Discussion Papers 12-2, Federal Reserve Bank of Boston.

Gale, D. 1996. "Equilibria and Pareto Optima of Markets with Adverse Selection." *Economic Theory* 7(2): 207–35.

Gorton, G. and A. Metrick. 2010. "Haircuts." Unpublished manuscript, Yale University.

Gorton, G. and G. Pennacchi. 1990. "Financial Intermediaries and Liquidity Creation." *Journal of Finance* 45(1): 49–71.

Guerrieri, V. and R. Shimer. 2013a. "Dynamic Adverse Selection: A Theory of Illiquidity, Fire Sales, and Flight to Quality." Unpublished manuscript, University of Chicago.

_____. 2013b. "Markets with Multidimensional Private Information." Unpublished manuscript, University of Chicago.

Inside Mortgage Finance, Mortgage Market Statistical Annual. 2011a. Volume I: The Primary Market.

_____. 2011b. Volumen II: The Secondary Market.

Jiang, W., A.A. Nelson and E. Vytlacil. 2011a. "Liars Loan? Effects of Origination Channel and Information Falsification on Mortgage Delinquency." Unpublished manuscript, Columbia Business School.

_____. 2011b. "Securitization and Loan Performance: A Contrast of Ex Ante and Ex Post Relations in the Mortgage Market." Unpublished manuscript, Columbia University Business School.

Keys, B.J., T. Mukherjee, A. Seru, and V. Vig. 2010. "Did Securitization Lead to Lax Screening? Evidence from Subprime Loans." *Quarterly Journal of Economics* 125(1): 306–62.

Krainer, J. and E. Laderman. 2011. "Mortgage Loan Securitization and Relative Loan Performance." Working Paper Series 2009-22. Federal Reserve Bank of San Francisco.

Milgrom, P. and N. Stokey. 1982. "Information, Trade and Common Knowledge." *Journal of Economic Theory* 26(1): 17–27.

Moody's Investors Services. 2012. Default & Loss Rates of Structured Finance Securities: 1993-2011. 16 November.

Piskorski, T., A. Seru and J. Witkin. 2013. "Asset Quality Misrepresentation by Financial Intermediaries: Evidence from RMBS Market." Research Paper 13-7, Columbia University Business School.

Spence, M. 1973. "Job Market Signaling." *Quarterly Journal of Economics* 87(3): 355–74.

Documentos de Trabajo Banco Central de Chile	Working Papers Central Bank of Chile
NÚMEROS ANTERIORES	PAST ISSUES
La serie de Documentos de Trabajo en versión PDF puede obtenerse gratis en la dirección electrónica: www.bcentral.cl/esp/estpub/estudios/dtbc. Existe la posibilidad de solicitar una copia impresa con un costo de Ch$500 si es dentro de Chile y US$12 si es fuera de Chile. Las solicitudes se pueden hacer por fax: +56 2 26702231 o a través del correo electrónico: bcch@bcentral.cl.	Working Papers in PDF format can be downloaded free of charge from: www.bcentral.cl/eng/stdpub/studies/workingpaper. Printed versions can be ordered individually for US$12 per copy (for order inside Chile the charge is Ch$500.) Orders can be placed by fax: +56 2 26702231 or by email: bcch@bcentral.cl.

DTBC – 715

Macro-Prudential Policy and the Conduct of Monetary Policy
Denis Beau, Christophe Cahn, Laurent Clerc y Benoît Mojon

DTBC – 714

Central Banking after the Crisis
Frederick S. Mishkin

DTBC – 713

Monetary Policy and Macro-Prudential Regulation: The Risk-Sharing Paradigm
Atif Mian

DTBC – 712

Monetary Policy at the Zero Lower Bound: The Chilean Experience
Luis F. Céspedes, Javier García-Cicco y Diego Saravia

DTBC – 711

Jobless Recoveries during Financial Crises: Is Inflation the Way Out?
Guillermo Calvo, Fabrizio Coricelli y Pablo Ottonello

DTBC – 710

Does Expansionary Monetary Policy Cause Asset Price Booms? Some Historical and Empirical Evidence
Michel Bordo y John Lando-Lane

DTBC – 709
Investment Horizons and Asset Prices under Asymmetric Information
Elias Albagli

DTBC – 708
Comparison of Model for Pricing Volatility
Néstor Romero

DTBC – 707
A PCA Approach to Common Risk Exposures in the Chilean Banking System
Diego Avanzini y Alejandro Jara

DTBC – 706
Coordination of Expectations and the Informational Role of Policy
Yang Lu y Ernesto Pastén

DTBC – 705
Wicksell versus Taylor: A Quest for Determinacy and the (Ir)relevance of the Taylor Principle
Sofía Bauducco y Rodrigo Caputo

DTBC – 704
The Changing Roles of Education and Ability in Wage Determination
Gonzalo Castex y Evgenia Dechter

DTBC – 703
Proceso de Endeudamiento y Sobre Endeudamiento de los Hogares en Chile
Jaime Ruiz-Tagle, Leidy García y Álvaro Miranda

DTBC – 702
The Wealth Distribution in Developing Economies: Comparing the United States to Chile
Sofía Bauducco y Gonzalo Castex

DTBC – 701
Impacto de Sorpresas Macroeconómicas y Anuncios en Factores de la Estructura de Tasas de Chile
Luis Ceballos

BANCO CENTRAL
DE CHILE

DOCUMENTOS DE TRABAJO • Diciembre • 2013