Essays on Securitization and the Resolution of Financial Distress

by

Manuel Adelino

B.S., Business Administration, Universidade Católica Portuguesa, 2003

Submitted to the Alfred P. Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2010

Author ....
Alfred P. Sloan School of Management
May 27, 2010

Certified by ....
Antoinette Schoar
Michael Koerner '49 Professor of Entrepreneurial Finance
Thesis Supervisor

Accepted by ....
Ezra W. Zuckerman
Director, Ph.D. Program, Sloan School of Management

Essays on Securitization and the Resolution of Financial Distress

by

Manuel Adelino

Submitted to the Alfred P. Sloan School of Management
on May 27, 2010, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This thesis consists of three empirical essays on securitization and the resolution of financial distress, examining the information content of prices of mortgage-backed securities, the role of securitization in preventing mortgage renegotiation and the effect of minimum bids in ascending price auctions of repossessed assets. In the first chapter, I present evidence that the prices at origination of residential mortgage backed securities contain information on the quality of the underlying asset pools above what is reflected in the ratings. Yield spreads at issuance predict both future downgrades and defaults even after the information contained in ratings is taken into account. This holds for all rating classes except triple-A. Yield spreads of the highest rated securities have no predictive power for future performance. This suggests that investors in triple-A were less informed about the quality of the securitized assets than investors in riskier, more information sensitive securities.

In the second chapter (co-authored with Kris Gerardi and Paul Willen), we show that securitization does not explain the reluctance among lenders to renegotiate home mortgages. We focus on seriously delinquent borrowers from 2005 through the third quarter of 2008 and show that servicers renegotiate similarly small fractions of securitized and portfolio loans. The results are robust to several different definitions of renegotiation and hold in subsamples where unobserved heterogeneity is likely to be small. We argue that information issues endemic to home mortgages where lenders negotiate with large numbers of borrowers lead to barriers to renegotiation fundamentally different from those present with other types of debt.

In the third chapter (co-authored with Antoinette Schoar), we study the effect of changing the reservation prices in online English auctions of repossessed motorcycles. The main intervention involves lowering the minimum bid initially set by the auctioneer by up to 10 percent in randomly selected vehicles. We find that lowering the minimum bids increases the probability of sale and reduces the final price conditional on a motorcycle being sold. This is in line with the predictions of standard auction theory and is evidence against auction fever being induced by lower initial minimum bids. The fact that bidders are regular participants in auctions and experienced buyers of the items for sale may help explain why they weren't subject to auction fever. Also, in our setting the number of bidders increased only slightly for discounted items, which is the opposite of what is necessary to induce an auction fever effect. For a subset of motorcycles that were bundled together for the auctions we a find lower probability of sale and lower prices conditional on sale.

Thesis Supervisor: Antoinette Schoar
Title: Michael Koerner '49 Professor of Entrepreneurial Finance

Acknowledgments

I am grateful to my advisors Antoinette Schoar, Stewart C. Myers and Nittai Bergman for their guidance, to my colleagues Jack Bao, António Sodré, Tara Bhandari, Apurv Jain, Alejandro Drexler, Jialan Wang, Mary Tian and Ngoc-Khanh Tran for their friendship and helpful comments. I also thank Paul Willen and Serdar Dinc for many invaluable discussions.

Finally, I want to acknowledge the love and support I received from all my family, in particular my grandparents Chinchá (my first teacher) and Álvaro, my parents Cristina and José Manuel and my brother João. I am most grateful to my wife, Teresa, and to my son, Francisco, for being a source of happiness every day.

Contents

List of Tables

Chapter 1

Do Investors Rely Only on Ratings? The Case of Mortgage-Backed Securities

1.1 Introduction

Rating agencies are believed to have played a prominent role in the development of the mortgage-backed securities market and have been at the center of the discussion on the roots of the 2007-2009 financial crisis. The agencies have been accused of performing insufficient screening of the underlying assets and of assisting issuers in designing deals to obtain the desired ratings, which ultimately resulted in misleading ratings for these securities (SEC, 2008; Bolton, Freixas and Shapiro, 2009; Mason and Rosner, 2007; Skreta Veldkamp, 2009).

Given the doubts cast on the quality of ratings and the subsequent poor performance of these assets, many have argued that investors in non-agency mortgage-backed securities (MBS) relied too heavily on the assessment made by the rating agencies and did not perform enough due diligence on the credit quality of the mortgage pools. Federal Reserve Chairman Alan Greenspan warned in 2005 that investors should "not rely solely on rating-agency assessments of credit risk."[1] In a May 2009 hearing of the House of Representatives, Rep. Hensarling argued that "Unfortunately, many investors, due to legal imperatives or practical necessity, relied exclusively on ratings from the three largest [credit rating agencies], without performing their own conservative due diligence."[2] This view was expressed by several other members of Congress in a variety of hearings in 2008 and 2009. A 2008 article in the New York Times summarized the conventional wisdom on this subject: "Who was evaluating these securities? Who was passing judgment on the quality of the mortgages[...]? Certainly not the investors. They relied on a credit rating."[3]

The view that investors did not perform any independent due diligence on the asset pools

[1]Greenspan, A., May 5, 2005, "Risk Transfer and Financial Stability". Remarks to the Federal Reserve Bank of Chicago's 41st Annual Conference on Bank Structure.

[2]Hensarling, J., May 19, 2009. Hearing before the Subcommittee on Capital Markets, Insurance, and Government Sponsored Enterprises of the Committee on Financial Services. U.S. House of Representatives, One Hundred Eleventh Congress, First Session.

[3]Lowenstein, R., April 27, 2008. "Triple-A Failure". New York Times.

is at odds with the more traditional view in finance that market prices reflect fundamentals and should be trusted to incorporate most if not all of the publicly available information. While ratings are expected to be part of the price setting process, they are only part of the information set available to investors.

This paper addresses two related questions about the information content of MBS yield spreads in the period leading up to the crisis. First, I address whether investors did in fact rely exclusively on ratings in setting the yield spreads of mortgage-backed securities. Specifically, I look at whether ratings are a sufficient statistic for the future credit performance of MBS or whether, alternatively, the issuance yield spreads contain additional information for predicting performance. Second, I consider whether the information content of yield spreads of these securities varies across different ratings. I analyze whether yield spreads at origination predict the likelihood of downgrade separately for securities of each rating, which implies that investors were able to (at least partially) identify the relative credit risk of different bonds with the same rating.

The empirical tests rely on the 2007-2009 shock to US house prices and the impact this shock had on the performance of mortgages and, consequently, mortgage-backed securities. There is widespread consensus that this shock was at the origin of the crisis in the financial sector and that it ultimately led to a recession in most of the industrialized countries (Greenlaw, Hatzius, Kashyap and Shin, 2008; Gorton, 2008; Brunnermeier, 2009). The analysis in this paper assumes that a shock of this nature should have been a priced risk in mortgage-backed securities' yields, so that securities with different sensitivities to nationwide house price drops would have different yields and also different outcomes in the crisis. Empirically, the advantage of using this event is that it produced significant variation in the credit outcomes at all rating levels, including triple-A.

The results show that investors did not rely exclusively on ratings when pricing the deals at origination. In fact, yield spreads have predictive power for both the probability of downgrade and of default after taking into account all the information contained in the ratings. The main specifications are logit models for each of the performance outcomes (downgrade and default) with rating fixed effects and controls for security characteristics and quarter of origination. The point estimates show that a move from the 25-th to the 75-th percentile of yield spreads leads to a 12-18 percent change in the probability of downgrade and to a 23-33 percent increase in the probability of default. The results hold in a variety of specifications and subsamples, including a linear probability model and if I use the magnitude of the rating change as the dependent variable.

The predictive power of yields comes, however, exclusively from ratings below triple-A. Yield spreads of the (lowest priority) triple-A classes have no statistically significant predictive power for future credit performance. Spreads of all lower rated classes predict both downgrades and defaults. In particular, while moving from the 25-th to the 75-th percentile in yield spreads of triple-A securities produces no statistically significant change in the probability of downgrade, a similar move in the distribution of yield spreads of all lower rated classes (AA through BB) leads to a significant 5-9 percent increase in the probability of downgrade. The difference in predictive power between triple-A and the lower ratings holds for fixed and floating rate classes and over most quarters of the downgrade wave of 2007:Q4 to 2008:Q3.

The result that only the yield spreads of triple-A securities are not predictive of future performance suggests that investors in the non-agency MBS market may have been differentially informed about the securities they bought. Golub and Crum (2009) argue that investors in triple-A securities "thought that their investments were sufficiently "out-of-the-money" to any reasonable estimate of model risk so that this level of expertise was not required." This could be a reflection of a segmented MBS market where lower rated (more junior) classes were absorbed by investors who had information about the credit quality of the assets being securitized and triple-A securities were bought by passive or less informed investors. Another explanation is that the same investors performed different levels of due diligence when buying triple-A versus other securities.

Evidence from industry reports points to different types of investors in triple-A and lower rated securities. Among the investors that focused almost exclusively on triple-A were government-sponsored agencies (in particular Fannie Mae and Freddie Mac), some Federal Home Loan Banks and also asset-backed commercial paper conduits established by both US and European commercial banks, set up primarily to take advantage of "cheap" short term borrowing. The bank-related vehicles had the added incentive that triple-A securities were subject to the lowest capital requirements in case they had to be brought back onto the banks' balance sheets. The buyers of lower rated tranches included specialized mortgage securities funds and also CDOs (typically set up by investment banks) who re-packaged these securities into new structures.

The presence of differentially informed investors in senior and junior securities is part of many academic models of tranching and securitization. Boot and Thakor (1993) argue that in a setting with costly information acquisition the tranching of assets is a way for originators to mitigate a typical "lemons" problem. The authors predict that more senior securities are absorbed by liquidity (or noise) traders, whereas informed traders buy the junior securities. Other models with a similar separation between informed and uninformed investors include Riddiough (1997) and Pagano and Volpin (2008) and my results provide evidence consistent with this feature of these models.

To further examine the information content of yield spreads I exploit the fact that there are securities in the data where the ranking by yield spreads and by ratings disagrees. For example, some double-A classes have lower issuance yield spreads than A classes of a different deal issued at around the same time. I use the securities where the ranking by spreads and by ratings is different to determine whether it was the rating agencies or the investors who predicted future downgrades better in these cases. For each rating pair, I construct a sample where all the lower rated securities (i.e. judged of worse quality by the rating agencies) have lower yield spreads than all the higher rated securities.

For low ratings (below double-A) I find that the ranking of securities using yield spreads is more effective at predicting future downgrades than the ranking implied by ratings. This happens despite much of the variation in yield spreads coming from differences in types of securities and levels of prepayment protection. Ratings do better at predicting downgrades for double-A and triple-A securities (consistent with yield spreads containing less information at higher rating levels). However, if I focus on subsamples where the disagreement is more extreme, the predictive power of yield spreads improves for double-A and triple-A and in some specifications becomes superior to that of ratings.

The remainder of the paper is organized as follows. Section 2 provides some background on ratings and the pricing of debt securities. Section 3 describes the data and provides summary statistics. Section 4 details the methodology and presents the empirical results. Section 5 discusses the results. Section 6 concludes.

1.2 Background on Ratings and the Pricing of Bonds

The question I address in this paper is whether ratings assigned at the time of origination are a sufficient statistic for the future credit performance of MBS or whether yield spreads have some additional information about the asset pools. The relationship between ratings and yield spreads has traditionally been viewed from a different perspective and several authors have addressed the question of whether ratings add any new information above what is already reflected in the prices. In the context of corporate bonds, Weinstein (1977) finds that rating changes do not affect prices and that, in fact, all the information of downgrades and upgrades was reflected in prices well in advance of the actual change. Pinches and Singleton (1978) find that equity prices also do not react to rating changes and that all relevant information had already been impounded into the prices. Hand, Holthausen, and Leftwich (1992) consider only unexpected ratings changes and find that bond prices do react to downgrades and upgrades. In a study of the determinants of bond prices at issuance, Gabbi and Sironi (2002) argue that ratings are the most important determinant of yield spreads in the eurobond setting and that in fact the dependence of prices on ratings became stronger over the time period the authors analyze (1991-2001). Hull, Predescu, and White (2004) consider credit default swaps on corporate bonds and find that those instruments correctly anticipate ratings announcements. Bao (2009) studies the ability of existing structural models of default to explain the cross section of yield spreads of US corporate bonds and shows that ratings are strongly correlated with the component of yield spreads not captured by the models. Kisgen and Strahan (2009) use the certification of a fourth credit rating agency by the SEC to conclude that ratings-based regulations on bond investment matter for a firm's cost of capital.

Outside of the context of corporate bonds, Cavallo, Powell and Rigobón (2008) find that ratings changes matter for sovereign debt and that prices are not a sufficient statistic for the underlying credit quality of the securities. Ammer and Clinton (2004) find that rating downgrades are not fully anticipated by the asset-backed securities market (unlike in the corporate bond market) and that these changes do have a significant price impact. The authors conclude that in the ABS setting investors tend to rely more on ratings for negative news about credit risk than in the corporate bond market. In a study of European securitization transactions, Cuchra (2005) shows that ratings are the most important determinant of spreads at the time of issuance. The author also finds that ratings are, however, not a sufficient statistic for prices, as other factors such as placement characteristics and creditor rights also affect yield spreads after ratings are taken into account.

In order to study the correlation of issuance spreads with future outcomes I rely on rating changes, as I have no information on the time series of yield spreads of MBS[4]. While it would be interesting to analyze the pace at which information was impounded into prices over the period I analyze, the advantage in looking only at prices at origination is that

[4]There is no organized secondary market for MBS or mortgage-related CDOs, so reliable pricing data is very hard if not impossible to come by.

both investors and rating agencies could potentially acquire a similar information set for evaluating the securities (the publicly available information was the same). The question of whether ratings are a sufficient statistic for yield spreads is more meaningful when both ratings and spreads can rely on the same information set, rather than having one or the other vary over time and potentially incorporate information learned at different times.

One of the important features of the 2007-2009 crisis for this study is the large number and the magnitude of downgrades experienced in the mortgage-backed securities market, which generates variation in the outcomes for tranches of all ratings (including triple-A). Another option for performing the analysis in this paper would be to look at the corporate bond market. However, while there have been periods with large numbers of downgrades for corporate bonds, for the most part these occurred at lower ratings than triple-A. There are very few instances of downgrades or default of triple-A bonds (Moody's, 2005), which makes it virtually impossible to test whether the pricing of these bonds at issuance contained information for their future performance in terms of ratings changes. Another advantage of mortgage-backed securities relative to corporate bonds is that MBS issues had several classes and there was a heavy issuance of MBS during the period of analysis, which provides a large sample of bonds of each rating for comparing across issues with different underlying fundamentals.

Benmelech and Duglosz (2009) provide a summary of the severity of the crisis in the mortgage-related securities market, as well as a comparison to the outcomes in other asset-backed securities sectors. The authors report a large wave of downgrades of mortgage-backed securities, with some tranches being downgraded by several notches (consistent with Table 1.2, described in more detail in the next section). The biggest collapse in ratings happened, however, in ABS CDOs (collateralized debt obligations backed by asset-backed securities) and these securities were responsible for the largest write-downs of financial institutions worldwide. The triple-A securities of ABS CDOs suffered severe downgrades earlier and of much larger magnitude than the ones observed for triple-A mortgage-backed securities. Still, the widespread downgrades also in MBS are in stark contrast to the stability in ratings that was typical in this market in the past (Violi, 2004). Luo, Tang and Wang (2009) argue that the problem with ratings of CDOs was that rating agencies didn't have enough data and used the wrong models to estimate the probabilities of default of CDO assets.

Becker and Milbourn (2008) show that the increased competition in the market for corporate bond ratings due to the entry of Fitch made ratings less informative. While there is limited direct evidence on the competition between the three agencies in the structured finance sector, the fact that Fitch has been active in this market segment from very early on and the large volume of security issuance points to at least a similar level of competitiveness as observed in the corporate bond market. The intense competition and the asset complexity in this market probably resulted in less informative ratings and this could have alerted investors to the need to perform increased due diligence on the quality of the assets. I explore the extent to which investors did or did not perform their own due diligence on the quality of the assets in the following sections.

Ashcraft, Goldsmith-Pinkham and Vickery (2009) show that ratings did not reflect all publicly available information at the time of issuance and that it is possible to construct a measure that predicts future performance better than ratings. This paper complements the work of Ashcraft et al (2009) by both confirming that ratings are not a sufficient statistic for future performance and by showing that investors actually incorporated some of the additional information into prices. Also, I consider the differential informational content by rating level and show that triple-A securities were different from lower rated tranches in the level of information contained in prices.

1.3 Data and Summary Statistics

1.3.1 Data Source and Variable Construction

The data used in this paper was obtained from the internal MBS database of JP Morgan and substantially complemented with hand collected data from Bloomberg. The data fields contain all the security identifiers (including CUSIP and ticker), the issuer name, the date of issuance, all original ratings (by the three major rating agencies), a description of the type of security (detailed in Appendix 1), the coupon at origination, the weighted average life and limited information on the underlying assets in each pool. The data covers approximately 80% of all RMBS issued in the US between 2003 and 2007. I also collected all ratings outcomes for each security up to the end of the third quarter of 2008 from Bloomberg. Treasury rates and swap rates were obtained from the Federal Reserve website and Datastream.

The sample obtained initially includes 87,017 securities. I drop from the sample all securities with missing coupon, amount issued or rating information (18,797 observations) as well as the following types of securities: Interest Only, Principal Only, Inverse Floater, Fixed to Variable and Descending Rate Bonds (127 bonds, many are dropped using the previous criterion)[5]. I also winsorize the data (using the coupon variable) at the 1% level to eliminate observations that are clear mistakes in data entry (this drops an extra 681 observations). The final sample contains 67,412 different securities from 5,712 RMBS issues.

The key variables for the analysis are the ratings and yield spreads. The ratings of each security were converted to a point scale where triple-A corresponds to 28 and then each rating notch corresponds to 1 point less (i.e., AA+ is 26 because AAA- does not exist, AA is 25, AA- is 24 and so on). This conversion of the rating scales to a numerical scale follows exactly Becker and Milbourn (2008). I also use the correspondence of the Moody's scale with that of S&P and Fitch that is common in the literature (AA of S&P and Fitch corresponds to Aa in Moody's, BBB of S&P and Fitch corresponds to Baa, etc.).

The numerical value of the ratings does not matter for most of the analysis because all specifications use ratings fixed effects as independent variables rather than just the point scale. Another alternative would be to use the scale directly as the independent variable. I use indicator variables for each rating to make sure that I capture all the information contained in the ratings. Using the rating scale directly would imply a number of assumptions that are not testable, including a linear relation between the rating scale and the outcome variables (or quadratic relation, if I included a rating squared term) which means that moving from AAA to AA+ is the same as moving from AA+ to AA- in terms of underlying credit risk.

Importantly, most (78%) of all securities I consider are rated by more than one rating agency (this is consistent with the findings in Benmelech and Dlugosz, 2009). As a summary

[5]These securities are dropped because the coupon cannot be directly compared to the other bonds.

variable for the ratings I choose to average the ratings in the point scale. For example, one AA rating (25 in the point scale) from one agency and one AA- from another agency (24 in the point scale) corresponds to 24.5. The regressions include dummies for all these cases to account for the possibility that this disagreement includes relevant information for future performance. In cases where there are three ratings and only one disagrees I approximate the rating to the closest half point. In specification 2 of Tables 1.4, 1.6 and 1.7 I perform the analysis using indicator variables for the rating classes instead of having one indicator variable for each rating as a robustness check (i.e., I only include dummies for AAA, AA, ... instead of doing this for each rating step AA+, AA, AA-, ...).

I define a downgrade as a negative transition from one rating class to a lower class (e.g., AA to A) and default is defined as a transition to CC or lower (or Ca or lower on Moody's scale). I don't observe actual default, but according to Moody's (2005) CC rated corporate bonds have a one-year default rate above 40% and this is likely to be much higher in the context of the 2007-2009 crisis and the deteriorating credit quality of the underlying assets. The dependent variable "downgrade" (used first in Table 1.4) is equal to 1 if the average rating of a bond is at least one class below its rating at origination at the time of analysis (either Q4 2007 or Q3 2008 in Table 1.4). The "default" variable used in Table 1.6 is constructed similarly, but the rating at the time of analysis must be CC or lower.

The yield spread for fixed rate securities is given by the difference between the coupon rate and the treasury constant maturity rate of similar maturity. I use the weighted average life[6] of each security as their maturity. Unlike most corporate bonds, the nominal maturity of mortgage-backed securities typically has no relation with the moment at which the principal is actually expected to be repaid and the weighted average life as reported by the issuer is the best approximation available for the expected maturity of each security[7].

In order to construct the yield spread for floating rate securities I use the swap rate of equivalent maturity (measured by the WAL as before), add the spread over the index rate[8] and subtract the treasury constant maturity rate of closest maturity. For example, a security that pays 1-month Libor + 50 basis points and has a weighted average maturity of 3 years will be converted using a 3 year swap rate. The equivalent coupon will be the 3-year swap rate plus 50 basis points. This assumes the spread over the index rate would be constant for all maturities of the index. Floating rate securities represent approximately 51% of the final sample by count, so I chose to convert the rates rather than focus on only one type of security for the main part of the analysis. The analysis is broken down by fixed and floating rate securities when I focus on triple-A securities only.

The above construction of yield spreads assumes that all securities are issued at par. I am not able to test this assumption because the price at issuance is not available in the data. There are, however, a few arguments that suggest that this assumption is unlikely to be problematic. First, the price formation process at origination (as described by industry sources) involves a competitive negotiated sale where the arranging bank adjusted the yield

[6]The weighted average life of a security is not the duration but rather the expected time to repayment of all capital. It is usually provided by the originator of the securities based on the expected repayment speed of the underlying mortgages.

[7]The most striking examples are the very senior triple-A tranches that usually have a weighted average life of one year or less and a nominal maturity of up to 30 years.

[8]The index rate for 97% of all the floating rate securities in the sample is the 1-month Libor.

spread as close as possible to the prevailing spread in the secondary market for deals of similar characteristics before marketing the deal. The deals are then placed with investors by first collecting orders and then delivering securities on a pro-rated basis (no auction takes place during the marketing of a new deal). The adjustment of the yield spread is done to avoid over or under-subscription of the new issue, both of which are costly to the issuer. Under-subscription usually implies negative deviations from par and a capital loss for the issuers and over-subscription is dealt with by investors receiving a smaller allocation of securities than they requested. Second, anecdotal evidence from July 2008 issues of the Global ABS/CDO Weekly Market Snapshot by JP Morgan indicates that the overwhelming majority of mortgage-backed securities are issued at par or very close to face value (within 1% of par). Finally, if the assumption that securities are issued at par does not hold this will introduce noise in the measurement of yield spreads and bias the results against finding a relationship between spreads and probability of downgrade. This could be a concern for triple-A securities (given that I find no relationship between yield spreads and future downgrades) and I address this issue in more detail in the discussion of Table 1.10 below.

1.3.2 Summary Statistics

Table 1.1 contains summary statistics for the frequency and total value of each type of security included in my sample. Panel A shows the breakdown by year. There was a significant increase in both the dollar amount and total number of securities issued during the period, with the peak activity in issuance happening in 2006. The total value of the RMBS securities included in the sample grows from $496 billion issued in 2003 to $1,080 billion issued in 2006 and the count of rated classes grows from 8,574 to 18,206 in the same period. This compares to approximately 9,000 rated RMBS tranches in 2003 and about 20,000 rated tranches in 2006 showed in Moody's (2008). The mean yield spread dropped between 2003 and 2006, although this fact could be driven either by a change in the composition of securities or by a change in market conditions over time. The deal complexity (measured by the number of rated tranches per deal) increased somewhat between 2003 and 2006, moving from about 10 tranches per deal to close to 13.

In terms of the breakdown by rating (Panel B), a majority of the securities in my sample are rated triple-A at origination (almost 35,000 tranches representing about 54% of all securities). In terms of amounts issued, on average about 90% of the value of all rated classes is triple-A (although this number excludes the non-rated component of these deals). The total amount issued is decreasing as we move to lower ratings and the total amount issued below investment grade only amounts to 10.3 million dollars in the 5-year period.

Panel B of Table 1.1 also shows the average spreads and the standard deviation of yield spreads by rating. Triple-A securities have a higher mean spread and higher standard deviation of yield spreads than AA securities. The main reason for this is the wide variety of types of securities issued in each deal and also variation between issues and in the time of issuance. I address this issue in more detail in Panel C. The last column of Panel B shows the average number of classes of each rating in each deal, conditional on a deal having at least one security with that rating. This column illustrates a widespread feature of MBS deals, namely that there are usually several triple-A classes, both with different seniority levels and also different characteristics along other dimensions (typically different prepayment risk or interest rate risk). Below triple-A, though, most deals have one or two classes of each rating. I will take this feature into account in the estimations below.

Panel C shows the results of a regression of yield spreads on indicator variables for each rating class and characteristics of each security. This is intended to address the issue that mean triple-A spreads are higher than those of double-A securities. The coefficients show that once security types are included, triple-A classes are associated with the lowest yield spreads and all other ratings sort as we'd expect *a priori* (with the exception of B). Including quarter of origination fixed effects and weighted average life as dependent variables doesn't change the coefficients significantly. This table also shows that almost 50% of the variation in yield spreads is left unexplained by these controls.

Table 1.2 shows the percentage of the securities downgraded and in default (defined as a transition to CC or lower) by cohort and by rating. Panel A shows the well known fact that the quality of the deals was decreasing over time, with the 2006 cohort showing the worst performance of the 5-year period. The percentage of tranches downgraded by the third quarter of 2008 is over 60% for the deals closed in 2006, compared to about 30% for the 2005 deals. The other striking feature of both Panel A and Panel B is the rapid deterioration of the credit quality of MBS securities that we observe. Before the onset of the crisis (shown in the column labeled June 2007), only a few of the lowest rated bonds had been downgraded. By September of 2007 there had already been at least one large wave of downgrades by each of the rating agencies. Then, between the end of 2007 and the third quarter of 2008, the percentage of securities downgraded was multiplied by about three for all cohorts, likely driven by the significant drop in house prices and the worsening economic conditions in the US during the same period.

The rating by rating analysis reveals that the percentage of tranches downgraded is decreasing with the rating level (triple-A being the highest) and the same applies to the percentage of securities that default. The one exception to this are B classes and the reason for this is that the most fragile deals at origination did not issue B tranches at all (they were left unrated for these issues). This means that once the crisis hit the B tranches actually were downgraded less often and defaulted less than BB securities. While this explains the pattern we observe, the fact that B classes perform better than BB is an indication that rating agencies may have missed the true quality of some of the securities being issued in this period.

One striking feature of Table 1.2 is that almost all triple-A securities had no change in ratings by the end of 2007, but in the subsequent 9 months a full 16% of triple-A tranches had been downgraded. The increase in downgrades for double-A securities is also surprising, with a total of 42% of all tranches being downgraded between January and September of 2008.

The level of downgrades and defaults reported in Panel B is far from what is typical for both securitized transactions and for corporate bonds. If we compare the level of downgrades in Panel B to the figures reported for structured finance bonds by Violi (2004), we can see that the typical 1-year probability of downgrade for triple-A structured finance bonds is around 1%, which compares to a total of 16% of all securities issued between 2003 and 2007 that got downgraded up to the third quarter of 2008.

Table 1.3 shows the transition matrix up to the end of 2008 for all securities issued in 2005 and in 2006 that are part of my sample. This can be compared to the transition matrices in other studies on ratings and also included in most reports put together by ratings agencies. The difference in quality between the 2005 and 2006 issues was already visible in Panel A of Table 1.2 and it is confirmed here for all rating levels. These matrices also show that there are very few upgrades of securities across the whole ratings spectrum.

1.4 Informativeness of Yield Spreads

1.4.1 Methodology

The main analysis in this paper addresses whether ratings are a sufficient statistic (in terms of predicting future credit performance) for yield spreads at origination. I use ratings downgrades and upgrades as a proxy for credit performance and make use of the fact that the crisis caused a large number of securities (including triple-A bonds) to be downgraded to identify more and less fragile securities within each rating.

The changes in ratings (almost all of which happen after the onset of the crisis in the housing market as shown in Table 1.2) is the only outcome I observe for each bond, so an important assumption in all the regressions is that these ratings changes are positively correlated with the changes in the credit quality and value of the underlying assets (i.e., the securities that were downgraded more also suffered a larger deterioration in credit quality). Brennan, Hein and Poon (2008) point out that the ratings of S&P and Fitch aim to reflect the default probability of the securities and that the model used by Moody's tries to predict expected losses due to default. Both these measures of credit quality are adversely affected by increased delinquencies in the underlying asset pool, so it is likely that this assumption holds for most cases.

In order to tease out the information content of yield spreads not contained in ratings I use indicator variables (or "dummies") for each rating as right hand side variables in all regressions. One attractive feature of this approach is that it is applicable irrespective of the level of informativeness of ratings. If ratings are completely uninformative for the future performance of securities, the ratings dummies should be statistically insignificant and yield spreads have a good chance of capturing the variation in credit quality of different securities. If, on the other hand, ratings are formed using all available information at the time of origination, it is still possible that prices predict future performance given the relative coarseness of ratings relative to prices. In fact, ratings can be seen as a coarse summary measure of an underlying continuous variable (e.g. probability of default), so there may still be variation in this underlying variable that is left uncaptured by ratings. Clearly, the more information incorporated by ratings, the harder it is to find predictability of future performance on the part of yield spreads.

The outcome variables in most of the analysis are indicators for whether a security has been downgraded or defaulted, so I use a logit model in the estimation. The equation I estimate is of the form:

$$Outcome_i = f(\alpha + \beta * YieldSpread_i + \lambda_{Rating_i} + \gamma * Characteristics_i + \epsilon_i)$$

The independent variable of interest is the yield spread (constructed as described in Section 3.1). The controls are fixed effects for each rating and other characteristics of each security that include timing of origination, maturity and security type.

The timing of origination is taken into account using quarter of origination fixed effects. These are intended to account for both the market conditions at the time of issuance and for the fact that most downgrades happen in 2007-2008 and securities issued at different times were at different points in the underlying pool's life cycle when the crisis hit (in terms of principal repayment, rate resets, etc.). I account for maturity by including both weighted average life (WAL) and the square of WAL as dependent variables. Finally, all regressions include security type dummies. This partially bypasses the need to model all the cash flows of each security, which I don't have enough information to do. Appendix 1 lists the variety of types of securities issued in RMBS deals. Most of the types refer to either protection from interest rate risk (fixed vs floating), protection from prepayment risk (e.g. planned amortization classes have fixed WAL and duration within a pre-set window of prepayment speeds) or the priority in repayment of principal[9]. In order to give yield spreads the best chance to reflect information on the pool fundamentals, the regressions include no information on the asset pools (such as average credit scores, prime vs. subprime, location of the mortgages, lien type, etc.).

The standard errors in all tests are clustered at the issue level to take into account the fact that the outcomes of different classes of the same deal are not independent observations (it is typically the case that several classes in the same deal are downgraded at the same time).

The difference between specification 1 and specification 3 in tables 1.4 to 1.7 is that specification 3 only includes the lowest priority triple-A class in each asset pool along with all the lower rated classes[10]. As discussed in the previous section (and shown in Panel B of Table 1.1) the typical MBS deal has about 6 different triple-A securities with different levels of seniority. If we ignore differences in security types[11], the yield spreads naturally rank the different triple-A securities by their seniority and the results in the regressions that follow could mechanically be driven by this (uninteresting) within-issue variation (higher spread triple-As are more junior and more likely to be downgraded). To be clear, the relevant test is whether investors can distinguish more fragile securities across different issues, not whether the prices rank securities within each issue (which follows trivially from the contract terms). Removing the higher seniority triple-A classes imposes a more demanding test on yield spreads, but one which answers the more relevant question of whether yield spreads have information about the fragility of each deal.

Limiting the analysis to the lowest triple-A classes should work in favor of finding a relationship between yield spreads and downgrades if the issuers engaged in ratings shopping

[9]Most securities are of more than one type (e.g. many floating rate bonds are also step rate bonds). This introduces noise in the measurement of yield spreads and could lead to finding no information in yield spreads. I address this concern in the rating by rating analysis below.

[10]Most RMBS deals include more than one asset pool (e.g. conforming and non-conforming) and the priority rules are usually designed so that the triple-A classes have cash flow rights over a specific pool, while the lower rated classes receive the remaining cash flows from both pools. I remove the higher priority triple-A classes in each pool whenever I have pool information (about 93% of the sample), so that most deals will have more than one triple-A class in the regressions.

[11]Usually there is significant variation in the types of triple-A securities issued in the same deal due primarily to differences in prepayment risk and interest rate risk (fixed vs floating, step rate bonds, planned amortization, etc)

as suggested by many industry sources, as well as Benmelech and Dlugosz (2009) and Bongaerts, Cremers and Goetzmann (2009). In fact, ratings shopping should make ratings less informative and (if the existence of ratings shopping was public information) should lead to more due diligence on the low priority triple-A tranches, the most likely to be the object of ratings shopping.

1.4.2 Results for the Pooled Sample

In this section I discuss the results for the whole sample, i.e. without splitting it by rating class. In Table 1.4 the dependent variable is an indicator that takes the value 1 if the security was downgraded by at least one rating class, as described in Section 3.1 (e.g., from AA to A). A positive coefficient on the yield spread variable means that securities with higher spreads are downgraded more often after taking into account the information for future downgrades contained in each rating. Table 1.4 reports results for all downgrades that happened up to the end of 2007 (columns labeled Q407) and the results using all the data (labeled Q308). This separation is intended to show the extent to which investors anticipated a relatively mild wave of downgrades that initially hit only the worst deals (the first columns) and a more widespread shock to the asset class (the full sample). All the reported coefficients are marginal effects from logit regressions.

The results show that investors were able to rank securities by their fragility and distinguish between bonds of the same rating but of different underlying asset quality. The regression with only quarter of origination fixed effects as controls (Panel 1 of Table 1.4) shows very significantly positive coefficients: a 10 basis points increase in yield spreads is associated with a 34-52 basis points increase in the probability of downgrade up to the end of 2007 and a 46-52 basis points increase in the probability of downgrade up to the third quarter of 2008. These estimates become 22-31 and 56-62 basis points respectively when I add the controls of weighted average life (WAL), WAL squared and security type dummies (Panel B of Table 1.4). By including the indicator variables for security types the effect of yield spreads is estimated only within security types, which alleviates the concern that spreads may have very different meanings when coupled with certain prepayment protection or other contract characteristics. The results hold also in a linear probability model (Table 1.5) rather than a logit and also if I use an ordered logit model with the size of downgrade as a dependent variable (unreported regressions).

The results are significant when we consider just the first wave of downgrades up to the end of 2007 and also when we take into account the full ratings history in the sample (and in fact become stronger in general). This is consistent with investors having ranked not only the *a priori* most fragile deals (downgraded early) but also the more resilient ones that were hurt when the crisis became more widespread.

The coefficients on the yield spread variable don't vary much across the three specifications. From specification 1 to specification 2 the sample is unchanged, but specification 2 includes only rating class dummies (i.e. one dummy for AAA, another for A, and so on), rather than one dummy for each rating (e.g. one for A+, another for A and one more for A-). Interestingly, while the coefficient on yield spreads does become larger, the change is small, suggesting that the ratings refinements add little information about the credit quality of a security. The sample used for specification 3 drops the higher priority triple-A classes, as described in the previous section. Yield spreads seem more informative once we do that, which suggests that the most senior triple-A classes added mostly noise to the estimation. The concern that the priority structure of the deals might be driving the result on yield spreads (as discussed in the previous section) is thus put to rest.

To make the results more clear from an economic perspective, the last two rows of both Panel A and Panel B have the mean probability of downgrade for the subsample used in each regression and also the effect of moving from the 25-th to the 75-th percentile in the dependent variable (yield spread)[12] on the probability of downgrade. The effects are large in economic terms, ranging from 3 to 6 percentage points change in the probability of downgrade in Panel A and 2-6 percentage points in Panel B, despite the fact that yield spreads are measured with significant noise (stemming from the construction of the yield spreads and the different types of securities in the sample). These quantities represent approximately a 10-25% change in the probability of downgrade across the several specifications.

As we include all the downgrades up to the end of the third quarter of 2008 the effect of yield spreads on the probability of downgrade becomes stronger but the share of tranches that were downgraded increases even more. This means that the percent change in the probability of downgrade implied by changing the yield spreads becomes smaller. This suggests that the yield spreads work better at identifying weaker securities and perform less well at distinguishing securities that are downgraded later.

The analysis of defaults is presented in Table 1.6 and while the estimated marginal effects are smaller in magnitude than those for downgrades, the mean probability of default is also much lower than that of downgrade. The coefficient on yield spreads is statistically significant at the 1% level across all specifications. The results point to a 7 basis points increase in probability of default up to the end of 2007 for each 10 basis points increase in yield spreads and a 16-22 basis points increase in the probability of default until the end of the third quarter of 2008 for the same increase in spreads. As in the previous table, the results are very stable across the three specifications.

In terms of the economic magnitude of the results reported in Table 1.6, moving from the 25-th to the 75-th percentile of the overall distribution of yield spreads (a move of around 105 basis points for the observations used in these regressions) changes the probability of downgrade by 20-30% (or 1 to 2 percentage points). The significant reduction in the number of observations relative to the previous tables is due to the fact that some of the "dummy" variables included for origination quarters and security types perfectly predict the outcome variable (the absence of default in most cases).

As a robustness check, I also run an ordinary least squares regression of the change in the rating (measured on a numeric scale as described in Section 3) as a dependent variable instead of an indicator for downgrade and default as in the previous tables. I present the results in Table 1.7. A one percentage point increase in yield spreads is correlated with a 0.2-0.3 points larger downgrade up to the end of 2007 and approximately a 0.5 points larger downgrade up to the end of the third quarter of 2008. A move from the 25-th to the 75-th percentile in yield spreads corresponds to a change in the size of the downgrade of similar magnitude (around 0.25 point larger downgrade up to 2007 and 0.5 points up to the third quarter of 2008).

[12]This move represents approximately 95 basis points.

The results by cohort (year-by-year) are presented in Table 1.8. The results are qualitatively similar to those presented in Tables 1.4 and 1.6 and are statistically significant for the large majority of subsamples considered. This contradicts claims that investors / originators became more reliant on ratings over the recent boom. In fact, the predictive power of yield spreads for future downgrades is not weaker for issues made later in the sample when compared to those made earlier. Also, the results become stronger when I consider all downgrades and defaults up to the third quarter of 2008 when the large wave of downgrades was in full swing (see Table 1.3).

1.4.3 Rating by Rating analysis

While it seems clear that yield spreads were informative for the future credit performance of MBS, it is informative to look at the results broken down by each rating (we can think of the previous results as a weighted average of the rating-by-rating coefficients on yield spreads). The coefficients shown in Table 1.9 reveal that triple-A rated classes are different from all the lower rated tranches. Indeed, yield spreads on triple-A classes are uncorrelated with both downgrades and with future rating changes. The coefficients have the right sign, but are statistically insignificant at conventional levels even though the standard errors are of the same order of magnitude as those of all other ratings. The result for defaults is in line with the other two columns but it is estimated using only a very small number of defaults of triple-A bonds.

The regressions for ratings below triple-A across all three specifications show a statistically significant relationship between spreads and the dependent variables in the direction that one would expect. The results are stronger for defaults than they are for downgrades. A move from the 25-th to the 75-th percentile in yield spreads changes the probability of downgrade by 5-10% for ratings between double-A and double-B, while the default probability moves by as much as 15-60%.

While there is a large number of triple-A securities and downgrades in the sample, yield spreads of tripe-A securities might be measured with more noise than those of other ratings. This could lead to an insignificant coefficient for triple-A bonds, so I address this issue in Table 1.10. It is important to note that while the percentage of downgraded securities is lower than for the other ratings (close to 20%) it still provides significant variation in the outcomes of securities and shouldn't affect the predictive power of yield spreads much. Also, the previous results for all ratings suggest that the informativeness of yield spreads exists even early on in the crisis, so yield spreads should be able to identify the weakest securities of each rating.

In order to reduce the noise in yield spreads of triple-A bonds I to split the sample into fixed and floating rate securities. By doing this I can use the yield spread variable as constructed before for the fixed rate subsample and the spread over the index rate (the 1-month LIBOR for 92% of the floating rate tranches) for the floating subsample. The conversion explained in Section 3 and used in all other regressions implicitly assumes that investors who bought floating rate mortgage-backed securities would be indifferent between receiving those securities and a fixed rate security of the same maturity at the prevalent rate in the swap market (plus the spread over the index rate). Directly using the spread over the index rate as reported by the issuer for floating rate bonds bypasses this assumption altogether.

In Table 1.10 I also consider the downgrades along the wave of downgrades. One concern about the results for triple-A versus all the other tranches is that the difference in coefficients could be present only for Q3:2008, so I repeat the regressions using downgrade data for all quarters between 2007:Q4 and 2008:Q3.

The results of the breakdown of triple-A securities into fixed and floating rate as well as a comparison to all tranches below triple-A are shown in Table 1.10. The first panel repeats the regression run in the first rows of Table 1.9 and also shows the results using downgrade information for the other quarters. The coefficients on the yield spread variable varies over the four specifications and actually has the wrong sign when I use the information up to Q1 and Q2 of 2008. The coefficient on the regression using rating changes is either insignificant or of the wrong sign. The coefficient for non-AAA securities is of the right sign and statistically significant for both downgrades and rating changes across all specifications. The percentage of downgraded tranches for triple-A bonds is lower than for non-AAA, but it is already at 5% by the end of the first quarter of 2008. Given that for non-AAA the ability of yield spreads to rank securities seems to be stable across the downgrade waves, the lower share of downgraded securities shouldn't be an important reason why triple-A yield spreads are uncorrelated with the future performance.

In Panel B of Table 1.10 I restrict the analysis to floating rate securities. As explained above, instead of the yield spread variable as a dependent variable I use the spread over the index rate in order to reduce measurement error[13]. The marginal effect of yield spreads on the probability of downgrade are insignificant or of the wrong sign for all quarters after 2007:Q4 and the coefficients from the rating change OLS regressions are also either statistically insignificant or the wrong sign. The results for tranches below triple-A are all of the expected sign (as in Table 1.10).

Fixed rate classes globally exhibit much lower shares of downgraded classes, which potentially makes the estimation of the coefficient on triple-A yield spreads harder. It is still the case though that one would expect yield spreads to identify the weakest securities. The coefficients of the first three specifications are statistically insignificant or of the wrong sign, but the last specification does have "correct" and statistically significant coefficients. The non-AAA classes always exhibit the right sign irrespective of the quarter we consider even though the share of downgraded securities grows from just 12% to over 40% over the four quarters.

Overall this table provides strong support for the differential ability of the yield spreads of triple-A versus non-AAA classes to predict the future credit performance and to pick out the most fragile securities. The results for non-AAA are very stable and survive all specifications, whereas the coefficients for triple-A securities are mostly either statistically insignificant or of the wrong sign, with the exception of fixed rate securities in the third quarter of 2008.

When I split the sample by the size of the MBS issue[14] the results (shown in Table 1.11) are broadly unchanged except for the largest quartile. In fact, the triple-A spreads of the

[13]The results are the same if I use the yield spread measure of the earlier regressions instead, suggesting the conversion from floating to fixed did not introduce a large amount of error into the estimation.

[14]The size of an issue is measured as the sum of the dollar amount of all securities included in an issue.

largest issue quartile predict future downgrades while yield spreads of smaller issues do not. This is consistent with investors having to incur a cost to acquire information about the underlying assets of MBS issues and preferring to do so for larger issues where it is possible to obtain a larger dollar amount allocation of securities. In unreported regressions, I find no difference between the explanatory power of yield spreads of issues arranged by the six most active investment banks in MBS deals (that accounted for around half of the securities in the sample) compared to all other deals.

Splitting the triple-A classes by cohort (Table 1.12) does not provide much additional information. Yield spreads do not consistently predict downgrades for any of the five cohorts considered and no clear pattern emerges in terms of the evolution of the informativeness of triple-A yield spreads over time.

1.4.4 Further Test on the Informativeness of Yield Spreads

The previous sections provided evidence that yield spreads of mortgage-backed securities contain information about the credit quality of the underlying assets. The question I address in this section is whether there are instances where yield spreads can do better than ratings, in particular whether spreads can provide a better guide to future performance of the securities in cases where there is disagreement between the two. It is clear from the summary statistics (Table 1.1) that lower ratings are subject to more downgrades, but also that there is substantial overlap in the yield spreads of adjacent ratings. Much of this overlap in yield spreads is due to differences in the types of securities and also to measurement error, so I test whether there is some information in this overlap that reflects disagreement in credit quality. This test is stronger than the previous analysis in terms of the informational requirements put on investors. In order to be right more often about the future performance of each security, investors need to use superior information relative to rating agencies or, at least, they need to incorporate the available information better.

In order to assess the relative performance of ratings and yield spreads in predicting downgrades I group securities by adjacent ratings and analyze ratings two by two. I sort all securities of both ratings by their yield spread (for example, I take AA and A securities originated in a quarter and sort them all by yield spread). In general, there will be higher rated securities with higher spreads than some of the lower rated securities. I repeat this procedure for each pair of ratings and each quarter of origination (so that the sort does not pick up time variation in yield spreads). I then select all securities of the higher rating that have a yield spread that is above the median spread[15] and all securities of the lower rating with a yield spread below the median spread for the two ratings. I then make the two groups of the same dimension by dropping observations from the larger group of the two (I drop the observations that are closest to the median spread). I repeat this exercise for other cutoffs. Instead of always using the median yield spread as the cutoff, I also try keeping the higher rated securities with spreads above the 60th percentile and lower rated securities with spreads below the 40th percentile. I repeat this for the 70th, 80th and 90th percentile cutoffs.

For all the securities selected by the procedure described above, yield spreads and ratings

[15]Within the rating pair and issuing quarter.

"disagree" about which is riskier. In fact, all selected higher rated securities have a yield spread that is higher than all lower rated securities, so either this is just "noise" due to differences in the securities themselves (prepayment risk, interest rate risk, expected maturity, etc.) or investors thought that the higher rated securities were more risky than the lower rated ones. The fact that yield spreads are noisy favors finding better predictability on the part of the ratings, because yield spreads do not pick up only credit risk, whereas ratings should do just that[16].

The results of this analysis are presented in Table 1.13. I run both unconditional OLS regressions (essentially a t-test) in the first column of each ratings pair and also an OLS regression with controls in the second column of each ratings pair. I find that yield spreads do better than ratings for low rated securities, which is more evidence that investors took credit risk into account when pricing RMBS. When I compare A to triple-B rated tranches, the A-rated group has an 11 percentage point higher downgrade frequency (4 p.p. when I include controls). This percentage rises to 16-26 percentage points when I pick the bonds with the largest differences in yield spreads. The results are similar when I compare triple-B and double-B rated securities, where the difference ranges from 13 to 36 percentage points.

The results are not as strongly in favor of yield spreads for high rated securities. When I use the median cutoff (Panel A) ratings perform well when I compare both triple-A and double-A and also for double-A compared to single-A. As we move from Panel A to Panels B through E the results for the higher ratings change. As I pick the tranches with more disagreement in the yield spreads the ratings perform progressively worse. The result for triple-A vs double-A becomes weaker as the difference in the yield spreads between the two ratings becomes larger and actually at the 90th percentile cutoff it is actually the case that triple-A securities get downgraded more than double-A. A similar pattern emerges from the double-A vs A pair, where AA only has significantly less downgrades than A in the first two panels and then the coefficient becomes positive (and statistically significant in one specification).

Overall, this analysis confirms that yield spreads contain information about the underlying asset pools and that this information is more visible for lower ratings than it is higher in the rating scale. This does not mean that ratings didn't provide useful information - this analysis was set up to capture the information content of yield spreads by picking the securities where prices disagreed from ratings the most. A similar analysis for the ratings would involve picking securities within a certain band of yield spreads and seeing whether ratings predicted downgrades better than prices. Given the noise with which yield spreads are measured, it would not be surprising to find ratings performing very well in such a setting.

1.5 Discussion of the Results

If investors included all available information when determining yield spreads, the analysis I perform in Tables 1.9 and 1.10 should reveal that higher spread securities were more likely to be downgraded irrespective of the rating, as long as a couple of assumptions hold. The

[16] S&P indicates in their "Guide to Credit Ratings" that ratings are in essence opinions about credit risk, not affected by other dimensions of a security (http://www2.standardandpoors.com/aboutcreditratings/)

first important assumption is that the shock we observe to house prices in the US starting in 2006 was ex ante (at the time of pricing) important enough, i.e. the shock we observe must be a negative realization of a relevant risk factor for these securities. I rely entirely on the 2007-2009 shock to generate dispersion in the outcomes for securities in different pools and I cannot test rigorously whether the effects of a bad shock to house prices is something that should have been taken into account ex ante in pricing MBS. Several studies and industry reports from before the crisis do indicate, however, that participants in the market recognized the central importance of house prices to the performance of MBS (see Gerardi, Lehnert, Sherland and Willen, 2008 for a more complete discussion) and I confirm that yield spreads of tranches below triple-A priced in this risk.

The second important assumption in order for all securities to be ranked appropriately is that they have ex ante different sensitivities to the common risk component (i.e. there must be triple-A securities that ex ante were more and less sensitive to the drop in house prices). This is a function of both the underlying asset pool (e.g. prime vs. subprime vs. Alt-A) and the structure chosen for financing the pool (e.g. the level of over-collateralization for the triple-A security). It is likely that the dispersion in sensitivity of triple-A securities to the performance of the underlying mortgages is lower than that of non-AAA rated securities. In fact, given that the more junior tranches suffer the losses first, triple-A tranches are by definition less sensitive to shocks to the mortgage pool. However, we do see dispersion in outcomes, which suggests that the triple-A securities did have ex ante different exposures to bad realizations in US house prices. Also, the fact that I consider only the lowest priority triple-A classes (the ones that were closest to double-A level) should dismiss the concern that the securities were insensitive to the quality of the underlying assets. The above regressions test whether the yield spreads at origination detected the differential levels of exposure. If the two assumptions hold, then yield spreads that reflected all available information should rank securities by their outcome in the crisis.

One potential explanation for the lack of predictive power of triple-A classes is that the lower sensitivity of triple-A classes to shocks to the underlying assets leads to lower power in the estimation of the coefficient of triple-A securities. Put another way, it could be that the reason we see no predictive power from triple-A securities is that the standard errors are too large due to lack of variation in the prices, so it could just be an issue of lack of power in the estimation. The intuition behind this can be drawn from an OLS setting, where the standard errors of the coefficients are simply given by the well-known formula $\sigma^2(X'X)^{-1}$. If there is no variation in the regressors X, the standard errors will be large and the regression coefficients statistically insignificant. Looking at the standard errors reported in Table 1.9 this does not seem like the correct explanation for the results. The coefficients on triple-A in any of the three regressions is estimated with as much precision as those of any other rating. Also, from Table 1.1 we can see that there is a lot of variation in triple-A yield spreads, possibly even more than we would expect ex ante. In Table 1.10 one can draw a similar conclusion and, in fact, the problem is that the coefficient often has the wrong sign, not that the estimate is not precise.

A second potential explanation for the triple-A result is that the lack of predictive power of triple-A securities is due to measurement error in the yield spreads. I addressed this issue in Table 1.10, by splitting the sample into fixed and floating rate classes. As I argued before, the yield spreads of floating rate classes should suffer less from measurement error than the overall sample. As a further test of whether the results in Table 1.10 are driven by errors in variables problems, I repeat the analysis using a linear probability model and then instrument the yield spread of triple-A classes using the spreads of double-A, single-A and triple-B securities of the same deal. This will solve the errors in variables problem as long as the error in triple-A yield spreads is not correlated with the error in lower rated yields (any deal level measurement error is not eliminated). The results are presented in Table 1.14 and are largely consistent with the previous table, indicating that a "classical" errors in variables problem is unlikely to be driving the results for triple-A classes.

As an extra robustness check I run a duration (Cox) model where each loan is "at risk" between mid-2007 and September 2008 (given that I observe almost no downgrades before this). "Failure" is defined as a downgrade and the regression includes the same controls as all other previous tables. The results are shown in Table 1.15 and are in line with the rest of the analysis. Yield Spreads of triple-A classes are not correlated with downgrades (the hazard ratio is statistically indistinguishable from 1), while those of the lower rated classes predict downgrades. One advantage of this specification is that it takes into account the full downgrade wave (shown in Table 1.10) and summarizes the effect of yield spreads in just one coefficient.

1.5.1 Models of Securitization and Differentially Informed Investors

This section addresses one potential explanation for the pattern identified for the predictive power of yield spreads of triple-A and lower rated securities, namely that the MBS market may be segmented and that the investors in triple-A securities are differentially informed from those in the lower rated (more information sensitive) bonds.

Residential mortgage-backed securities (RMBS) are more complex structures than most other asset backed securities (Mason and Rosner, 2007), and the pricing of these securities requires some degree of sophistication on the part of investors. Most RMBS deals include many tranches (including many triple-A bonds issued on the same pool of assets, as shown in Table 1.1), as well as a wide variety of security types with different levels of prepayment protection and interest rate risk. Many have suggested that although (or maybe because) the securities in question are complex, some investors may have avoided performing (costly) due diligence on the RMBS they bought.

The fact that we find a pattern where the information content of yield spreads is restricted to ratings below triple-A suggests that the investors in triple-A bonds may be less informed than those in lower rated classes. One model of tranching that has this implication is that of Boot and Thakor (1993). In this model, information insensitive classes are bought by uninformed investors and the informed investors preferably buy more information sensitive securities. There are good and bad issuers and tranching the pools of assets is a way around the "lemons" problem faced by the good issuers. The model includes a fixed cost of acquiring information that is only worth incurring for trading in the information sensitive securities.

One limitation of the model of Boot and Thakor (1993) is that the information sensitive tranches are assumed to be risk-free. One can, however, extend their result to risky securities that are more or less information sensitive depending on the quality of the underlying assets if we make some simplifying assumptions. The intuition behind this is that investors only benefit from distinguishing between senior tranches in states of the world where "bad" issuers default on their senior securities. On the other hand, they benefit from distinguishing between junior classes in all other states of the world. If it is unlikely that the bad issuers default on the senior debt (i.e. if it is unlikely that any triple-A defaults) then risk-neutral investors prefer to invest in distinguishing between junior securities.

We can consider a simple setting with an equal quantity of "good" and "bad" issuer that investors cannot distinguish and that there is a fixed cost of acquiring information about the true quality of the issuers (M). Investors can choose to become informed if they incur this cost. For simplicity assume additionally that (i) both types of issuers choose to tranche the pools of assets such that there is a junior and a senior tranche[17]; (ii) the (one period) cash flows available to the issuers are of the form $X_i = \theta_i + \epsilon$ where $i \in \{g, b\}$, $\theta_g > \theta_b \geq S$, $\epsilon\tilde{\ } f_\epsilon(0, \sigma^2)$ and $\epsilon > -\theta_g$ such that the good and the bad issuer differ in the expected amount available to investors (as in Boot and Thakor) and are subject to the same underlying risk factor; (iii) good and bad issuers choose to issue the same amount of each type of security (if they didn't, bad issuers would reveal themselves to investors as being of the bad type) and (iv) there is a way for informed investors to hide their information, i.e. the prices do not fully reveal the information they possess (one way to have this at origination is for investors to submit private bids for the assets). Each investor has exogenous demand for one security[18] and is risk neutral.

We set the amount promised by the senior security to be S for both types of issuers and have the junior security absorb the remaining cash flows. In this setting, the payoff to the senior tranche is given by $Min\{S, \theta_i + \epsilon\}$[19] and the payoff to the junior tranche is $Max\{0, \theta_i + \epsilon - S\}$. The difference in expected payoff between the good and bad senior tranches is given by

$\pi_{Senior} = \int_{-\theta_g}^{-(\theta_g - \theta_b)} (\theta_g - \theta_b) f_\epsilon \, d\epsilon + \int_{-(\theta_g - \theta_b)}^{-(\theta_b - S)} (-\epsilon) f_\epsilon \, d\epsilon + \int_{-(\theta_b - S)}^{+\infty} 0 * f_\epsilon \, d\epsilon$

and the difference in expected payoff between good and bad junior tranches is given by

$\pi_{Junior} = \int_{-\theta_g}^{-(\theta_g - \theta_b)} 0 * f_\epsilon \, d\epsilon + \int_{-(\theta_g - \theta_b)}^{-(\theta_b - S)} (\theta_g - \theta_b + \epsilon) f_\epsilon \, d\epsilon + \int_{-(\theta_b - S)}^{+\infty} (\theta_g - \theta_b) * f_\epsilon \, d\epsilon$

Whether an investor chooses to become informed about the issuer type depends on the size of M. If the (risk-neutral) investors choose not to become informed, they submit a bid equal to the expected value of the payoff of each tranche. Because $\pi_{Senior} < \pi_{Junior}$, there is a range of M, the cost of acquiring information, for which it pays to become informed when buying the junior tranche, but not if the investor invests in the senior class. This range increases with the difference between θ_g and θ_b and also with the distance between S and θ_b (i.e. the safer the senior tranche, the more likely it is that investors prefer to incur in M when they are buying the junior tranche instead of the senior one). Note that if investors are not risk neutral, then the states of the world where the senior tranche defaults may be very valuable (i.e. they my have very high state prices) and the relationship between π_{Senior} and π_{Junior} becomes less clear.

As the above exercise shows, the fact that investors failed to perform due diligence on triple-A securities does not imply that they were not behaving optimally from their point of

[17]Boot and Thakor (1993) don't assume that issuers tranche the cash flows, given that their main interest is to derive tranching as an optimal strategy on the part of the issuers.

[18]The (wealth constraint) assumption that each investor buys only one security is important, otherwise investors can re-use the information acquired about the deal to price both securities correctly.

[19]Note that the senior tranche is no longer risk-free.

view. This may, in fact, reflect a rational trade-off between the costs of acquiring information and the losses in the states of the world where the senior securities default. There are, however, policy implications from the fact that some investors choose to remain uninformed. In particular, just as the ratings agencies act as a certifier for the quality of the securities in each issue, there may be room for an external party that distinguishes between higher and lower quality senior securities or, alternatively, monitors institutional exposure to this type of security. It is unclear whether it would be optimal for the participants in the market to set up such an mechanism or whether this would require the intervention from regulators.

There are several other authors in the literature that have assumed that investors were not fully informed about the quality of the underlying assets in the structured finance market. Bolton, Freixas and Shapiro (2009) argue that agencies have a larger incentive to inflate ratings when the fraction of naive investors in the market is large. In a related paper, Skreta and Veldkamp (2009) point out that the greater complexity of the pool of assets being securitized may create incentives for issuers to shop for ratings when some investors are not fully informed. Also, in the model of Pagano and Volpin (2008) the tranching of assets is one way to increase transparency of ratings and this is achieved by selling the more information sensitive tranches to sophisticated investors.

There is recent empirical evidence that investors in triple-A securities may have mispriced the risk they were taking. Coval, Jurek and Stafford (2009) consider the pricing of senior structured finance bonds and compare it to the prices of similarly rated corporate bonds. They argue that structured finance bonds are set up as catastrophe bonds (i.e. only default in the worst states of the world) and that investors only included the probability of default in the prices, rather than also taking into account the states of the world in which default happens. This meant that structured finance bonds were too expensive relative to corporate bonds. In related work, Brennan, Hein and Poon (2008) model the gains to issuing securities that are priced only as a function of the probabilities of default and these authors assume that "some investors are not able to assess the value of the securities themselves, but must rely instead on the bond ratings provided by third parties."

1.5.2 Anecdotal Evidence of Segmentation

It is hard to find systematic data on who actually held triple-A mortgage-backed securities and the MBS of lower ratings. There is, however, some anecdotal evidence that the investors in triple-A may have been different from those in the other classes.

Fannie Mae and Freddie Mac (two large government-sponsored entities whose mandate is to finance mortgages) were among the largest buyers of private-label MBS in the 2004-2006 period. After in 2004 the U.S Department of Housing and Urban Development increased the target for affordable loans that Fannie Mae and Freddie Mac were required to finance each year, the two agencies started to aggressively purchase private-label MBS. By some accounts (e.g. Leonnig, 2008) the two agencies bought up to 424 billion dollars worth of subprime MBS between 2004 and 2006.[20] According to the 2007 third quarter 10-Q of Fannie Mae, 99% of the securities purchased by this institution were triple-A.

Another group of investors who held private-label mortgage-backed securities were the

[20]Leonnig, C., June 10, 2008. How HUD Mortgage Policy Fed The Crisis. Washington Post.

Federal Home Loan Banks (FHLB). According to the 2008 annual reports of the San Francisco FHLB and the Seattle FHLB, both these institutions had large amounts of private-label MBS in portfolio (about 7 billion dollars and 2 billion dollars, respectively) and both banks invested almost exclusively in triple-A securities.

One of the important drivers of the demand for triple-A paper may have been the regulatory framework of commercial banks. In fact, the capital requirements of many key buyers of MBS, including commercial banks, broker-dealers and insurance companies, are a function of the quality of the assets held on the balance sheet (where the quality of debt instruments is often measured by their rating). Triple-A mortgage- and asset-backed securities provided a higher yielding alternative to traditional triple-A corporate bonds or treasuries and, since 2002, these securities receive the same risk weighting as equally rated debt of government sponsored entities for calculating the risk weighted assets of US depository institutions[21]. Also, the more recent Basel II rules distinguish the capital requirements for holding triple-A mortgage-backed securities from all lower rated securities, as shown in Appendix 2 (no longer binning together triple-A and double-A securities as was the case under Basel I).

Bank-sponsored securities arbitrage asset-backed commercial paper (ABCP) programs were one group of investors that was notorious for buying mostly triple-A rated mortgage-backed (and CDO) securities. These programs were set up to benefit from low short term funding rates (ABCP has a typical maturity of one year or less) by investing the proceeds in longer dated (higher yielding) assets. According to Moody's (2007) the securities arbitrage ABCP programs sponsored by US banks held approximately 94% of its assets in triple-A securities (CDOs, commercial and residential mortgage-backed securities). The securities arbitrage ABCP programs of European banks held 19% of their assets in US RMBS and of those RMBS assets a full 98% were rated triple-A. Of the CDO securities held by these conduits 98.5% was rated triple-A as of September of 2007. The widely publicized troubles of two regional German banks (Landesbank Sachsen and IKB) were linked to their exposure to their ABCP programs (Fitch, 2007).

The investors that are known to have bought below triple-A mortgage-backed securities seem *a priori* more specialized and better informed. If we look at the composition of specialized mortgage securities funds (such as FMSFX run by Fidelity) the distribution of non-government backed mortgage-backed securities is about 20% triple-A, 70% other investment-grade securities and 10% non-investment grade classes (although this fund in particular invested mostly in government-backed securities). Mezzanine CDOs, often put together by the same banks that arranged MBS deals, were by definition buyers of mezzanine classes of mortgage-backed securities (typically BBB).

[21] Since 2002 triple-A and double-A asset-backed and mortgage-backed securities are assigned a risk weight of 20% for calculating the risk-weighted assets of US depository institutions, the same as government-sponsored entities (Federal Register, November 29, 2001, Risk-Based Capital Guidelines; Capital Adequacy Guidelines; Capital Maintenance: Capital Treatment of Recourse, Direct Credit Substitutes and Residual Interests in Asset Securitizations; Final Rules).

1.6 Conclusion

This paper shows that investors did not rely exclusively on ratings when they priced residential mortgage-backed securities at of origination. Yield spreads at issuance predict future performance of mortgage-backed securities even after the information contained in the ratings is taken into account. These results are robust to a number of specifications and performance measures (downgrade, default and rating change).

The ability of yield spreads to predict future performance is driven exclusively by ratings below triple-A. This holds for both downgrades and rating changes, as well as for different subsamples (fixed and floating rate classes and for most cohorts). Also, in cases where ratings and yield spreads disagree about the quality of securities, yield spreads predict downgrades better for low rated securities but results are mixed for triple-A and double-A securities.

The inability of yield spreads of triple-A securities to predict outcomes suggests that investors may have performed less due diligence on these tranches than on lower rated securities. This indicates that there may be a non-trivial cost of acquiring information for securities that are relatively complex, resulting in different types of investor pools for bonds of different ratings.

There are many interesting questions about the informational content of prices that are left for future research. One question that remains is what information was reflected in prices that was not included in ratings. In particular, it is interesting to know whether we can fully characterize this additional information using publicly available "hard" information or whether prices have predictive ability even above publicly available coded data. One can equate the publicly available information at the time of issuance with datasets that have information on the loans included in each pool (such as the LoanPerformance dataset) and see to what extent adding this information eliminates yield spreads' ability to predict performance.

Another question is whether these results extend to other asset classes, namely CDOs, where the true quality of triple-A classes was harder to evaluate than MBS and also where securities were much harder hit by the crisis. Some of the buyers of triple-A mortgage-backed securities (such as the asset-backed commercial paper programs sponsored by banks) were also very active buyers of triple-A CDO securities, so it is possible that the patterns we find for MBS extend to CDOs. This would suggest that the relevant dimension of segmentation in structured finance (and possibly all debt markets) is the information sensitivity of the securities (reflected in the ratings) and not the underlying assets of the debt products themselves (MBS, CDOs or even corporate bonds).

1.7 Appendix 1 - Security Types in the Data

The dummies for types of bonds included in the regressions are listed below, as well as the frequency with which each occurs. All regressions that include security characteristics have dummies for each type of bond.

Name	Frequency Triple-A	Frequency Other Ratings
Accel Security	19932	749
Floater	13768	20940
Step Rate Bond	12208	22105
Sequential Pay	6045	166
Coll Strip Rate	5844	5787
Super Senior	5298	8
Available Funds	5239	1827
Int Rate Contract	4828	169
Pass Through	3221	256
Senior Support	2812	855
Nonaccel Security	2513	8163
Exchangeable	1576	29
Planned Amortization	1514	27
Accretion Directed	1323	9
Support	1087	10
Accrual	805	9
Non-Zero Delay	378	25
Retail	317	2
Targeted Amortiz.	286	2
Structured Coll	175	882
Component	120	33
Mandatory Redemp.	107	0
Extended Reset	100	6
Mezzanine	91	19306

1.8 Appendix 2 - Risk Weights for Securitization Exposures in Basel II RBA

	Risk weights for senior securitization exposures backed by granular pools (percent)	Risk weights for non-senior securitization exposures backed by granular pools (percent)	Risk weights for securitization exposures backed by non-granular pools (percent)
AAA	7	12	20
AA	8	15	25
A	10-20	18-35	35
BBB	35-100	50-100	50-100
BB	250-650	250-650	250-650

Source: Federal Register, December 7, 2007. "Risk-Based Capital Standards: Advanced Capital Adequacy Framework - Basel II; Final Rule"

Bibliography

[1] Ammer, J. and N. Clinton (2004). Good News Is No News? The Impact of Credit Rating Changes on the Pricing of Asset-Backed Securities. International Finance Discussion Papers Number 809.

[2] Ashcraft, A., P. Goldsmith-Pinkham and J. Vickery (2009). MBS ratings and the mortgage credit boom. Federal Reserve Bank of New York Working Paper.

[3] Bao, J. (2009). Structural Models of Default and the Cross-Section of Corporate Bond Yield Spreads. Working Paper.

[4] Becker, B. and T. Milbourn (2008). Reputation and competition: evidence from the credit rating industry. Harvard Business School Working Paper 09-051.

[5] Benmelech, E. and J. Dlugosz (2009). The Credit Rating Crisis. NBER Working Paper 15045

[6] Bolton, P., X. Freixas, and J. Shapiro (2009). The Credit Ratings Game. SSRN Working Paper 1342986.

[7] Bongaerts, D., M. Cremers and W. Goetzmann (2009). Multiple Ratings and yield spreads. Yale ICF Working Paper 08-27.

[8] Boot, A. W. A. and A. Thakor (1993). Security Design. *Journal of Finance*, 48:4, 1349-1378.

[9] Brennan, Michael J., J. Hein and S. H. Poon (2008). Tranching and Rating. SSRN Working Paper 1108556.

[10] Cavallo, E. A., A. Powell and R. Rigobon (2008) Do Credit Rating Agencies Add Value? Evidence from the Sovereign Rating Business Institutions. Inter-American Development Bank Research Department Working Paper 647.

[11] Coval, J. D., J. W. Jurek and E. Stafford (2009) Economic Catastrophe Bonds. *American Economic Review*, 99:3, 628–666.

[12] Cuchra, M. F. (2005) Explaining Launch Spreads on Structured Bonds. SSRN Working Paper 687800.

[13] Fitch Ratings (2007) Asset-Back Commercial Paper & Global Banks. Fitch Teleconference, August 23.

[14] Gabbi, G. and A. Sironi (2002) Which Factors Affect Corporate Bonds Pricing? Empirical Evidence from Eurobonds Primary Market Spreads. Bocconi U Newfin Research Center Working Paper; SSRN Working Paper 308299.

[15] Gerardi, K., A. Lehnert, S. M. Sherland and P. S. Willen (2008) Making Sense of the Subprime Crisis. Federal Reserve Bank of Boston Public Policy Discussion Paper No. 09-1.

[16] Golub, B. W. and C. C. Crum (2009) Risk Management Lessons Worth Remembering From the Credit Crisis of 2007-2009. Working Paper.

[17] Hand, J. R. M., R. W. Holthausen and R. W. Leftwich (1992) The Effect of Bond Rating Agency Announcements on Bond and Stock Prices. *The Journal of Finance*, 47:2, 733-752.

[18] Hull, J., M. Predescu and A. White (2004) The relationship between credit default swap spreads, bond yields, and credit rating announcements. *Journal of Banking & Finance*, 28:11, 2789-2811.

[19] Kisgen, D. J. and P. E. Strahan (2009) Do Regulations Based on Credit Ratings Affect a Firm's Cost of Capital? NBER Working Paper 14890.

[20] Luo, D., D. Y. Tang, and S. Q. Wang (2009). A Little Knowledge is a Dangerous Thing: Model Specification, Data History, and CDO (Mis)Pricing. Working Paper.

[21] Mason, J. and J. Rosner (2007). Where Did the Risk Go? How Misapplied Bond Ratings Cause Mortgage-backed Securities and Collateralized Debt Obligation Market Disruptions. SSRN Working Paper 1027475.

[22] Moody's Investors Service (2005) Default and Recovery Rates of Corporate Bond Issuers, 1920-2004. Special Comment.

[23] Moody's Investors Service (2007) Moody's Update on Bank-Sponsored ABCP Programs: A Review of Credit and Liquidity Issues. Special Report.

[24] Moody's Investors Service (2008) Structured Finance Rating Transitions: 1983-2007. Special Comment.

[25] Pagano, M. and P. Volpin (2009). Securitization, Transparency and Liquidity. CSEF Working Paper 210.

[26] Pinches, G. E. and J. C. Singleton (1978) The Adjustment of Stock Prices to Bond Rating Changes. *The Journal of Finance*, 33:1, 29-44.

[27] Riddiough, T. J. (1997) Optimal Design and Governance of Asset-Backed Securities. *Journal of Financial Intermediation*, 6, 121-152

[28] Securities and Exchange Commission (2008) Issues Identified in the Commission Staff's Examinations of Select Credit Rating Agencies. Summary Report.

[29] Skreta, V. and L. Veldkamp (2009). Ratings Shopping and Asset Complexity: A Theory of Ratings Inflation. SSRN Working Paper 1295503.

[30] Violi, R. (2004) Credit Ratings Transition in Structured Finance. CGFS Working Group on Ratings in Structured Finance Working Paper.

[31] Weinstein, M. I (1973) The Effect of a Rating Change Announcement on Bond Price. *Journal of Financial Economics*, 5, 329-350.

Table 1.1: Sumary Statistics

Panel A: By Cohort

	Amount issued (USD bn)	Count	Mean Spread (bp)	St. Dev. Spread (bp)	Av. No. classes
2003	496.5	8,574	197	114	9.9
2004	767.3	11,460	168	101	9.6
2005	1,058.5	17,135	134	72	12.2
2006	1,080.4	18,206	119	68	13.1
2007	802.1	12,037	137	74	14.2

Panel B: By Rating

	Amount issued (USD bn)	Count	Mean Spread (bp)	St. Dev. Spread (bp)	Av. No. classes
AAA	3,759.4	34,935	122	76	6.7
AA	210.8	9,638	113	46	1.2
A	109.2	8,845	146	58	1.2
BBB	82.6	9,663	228	101	1.2
BB	9.8	1,162	298	119	1.1
B	0.3	119	175	101	1.1

Panel C: Marginal Effect of Ratings on Yield Spreads

	(a)	(b)	(c)
AA	0.15***	0.17***	0.18***
	(0.02)	(0.02)	(0.02)
A	0.47***	0.48***	0.48***
	(0.02)	(0.02)	(0.02)
BBB	1.29***	1.30***	1.30***
	(0.02)	(0.02)	(0.02)
BB	1.93***	2.00***	1.99***
	(0.04)	(0.04)	(0.05)
B	0.75***	0.73***	0.73***
	(0.12)	(0.11)	(0.11)
Type FE	Y	Y	Y
Time FE	N	Y	Y
WAL	N	N	Y
Observations	67,412	67,412	67,412
R2	39.5%	48.2%	49.0%

Note: Dollar Amounts shown in Panels A and B in billions. Panel B excludes securities where there was disagreement between agencies about the rating class (e.g. AA vs A). Yield spreads shown in basis points in Panels A and B. Panel C shows coefficients of an OLS regression of yield spreads on rating class "dummies" and security characteristics. Standard errors (shown in parenthesis) are clustered at the issue level. All specifications include fixed effects for class type. Specifications (b) and (c) include quarter of origination fixed effects. Specification (c) also includes weighted average life (WAL) and WAL squared as controls.

Table 1.2: Frequency of Downgrades

Panel A: By Cohort

	June 2007		September 2007		December 2007		September 2008	
	Downgrade	Default	Downgrade	Default	Downgrade	Default	Downgrade	Default
2003	0.0%	0.0%	2.8%	0.2%	3.8%	0.4%	6.9%	1.2%
2004	0.1%	0.1%	2.3%	0.1%	5.0%	0.4%	13.6%	1.8%
2005	0.1%	0.0%	2.8%	0.2%	8.3%	0.7%	29.7%	4.4%
2006	0.1%	0.1%	7.2%	1.0%	22.2%	2.8%	61.5%	12.8%
2007	0.0%	0.0%	0.2%	0.0%	14.8%	0.1%	56.1%	6.8%

Panel B: By Rating

	June 2007		September 2007		December 2007		September 2008	
	Downgrade	Default	Downgrade	Default	Downgrade	Default	Downgrade	Default
AAA	0.0%	0.0%	0.3%	0.0%	0.5%	0.0%	16.2%	0.0%
AA	0.0%	0.0%	2.2%	0.0%	6.6%	0.1%	48.3%	1.6%
A	0.0%	0.0%	4.7%	0.2%	23.9%	0.9%	60.7%	7.3%
BBB	0.2%	0.1%	12.1%	1.4%	37.1%	3.7%	68.7%	24.8%
BB	2.2%	1.8%	31.5%	7.9%	62.0%	19.9%	84.4%	55.7%
B	1.7%	0.8%	4.8%	2.9%	18.2%	3.6%	49.6%	37.8%

Note: Table shows percentage of all classes in each category that were downgraded or defaulted. Downgrade is defined as a negative transition from one rating class to a lower class (e.g., AA to A). Default is defined as a transition into a CC rating or lower (Ca on Moody's scale). Ratings table excludes securities where there was disagreement between agencies about the rating class (e.g. AA vs A).

Table 1.3: Transition matrix for 2005 and 2006 cohorts

Panel A: 2005 cohort

	AAA	AA	A	BBB	BB	B	CCC or lower
AAA	96.2%	2.0%	1.2%	0.3%	0.2%		
AA	0.1%	73.0%	15.6%	7.0%	2.2%	1.4%	0.7%
A		0.2%	46.6%	16.8%	14.0%	10.6%	11.7%
BBB			0.1%	33.0%	14.5%	17.4%	35.0%
BB					15.7%	10.5%	73.8%
B						45.5%	54.5%

Panel B: 2006 cohort

	AAA	AA	A	BBB	BB	B	CCC or lower
AAA	74.5%	10.2%	8.1%	4.2%	2.0%	0.5%	0.4%
AA		24.2%	15.7%	15.2%	14.3%	19.6%	11.0%
A		0.1%	10.9%	8.7%	10.4%	20.4%	49.6%
BBB				7.8%	4.3%	9.8%	78.1%
BB					9.3%	0.3%	90.4%
B						35.7%	64.3%

Note: Includes all ratings agencies. Sample includes all securities issued in 2005 and in 2006. Paid off securities are recorded at their last rating before being paid off.

Table 1.4: Downgrade

Panel A: No Controls

	Specification 1		Specification 2		Specification 3	
	Q407	Q308	Q407	Q308	Q407	Q308
Yield Spread	3.43***	4.63***	3.82***	4.98***	5.22***	5.21***
	(0.19)	(0.39)	(0.18)	(0.38)	(0.28)	(0.39)
Rating FE	Y	Y	-	-	Y	Y
Rating Class FE	-	-	Y	Y	-	-
Lowest AAA only	-	-	-	-	Y	Y
Observations	66,554	67,395	66,556	67,405	43,762	44,262
Mean Prob of Downgrade	12.0%	37.4%	12.0%	37.4%	18.2%	49.3%
Effect of P75-P25 move in yield sp	3.1%	4.3%	3.5%	4.6%	5.5%	5.6%

Panel B: With controls

	Specification 1		Specification 2		Specification 3	
	Q407	Q308	Q407	Q308	Q407	Q308
Yield Spread	2.24***	5.87***	2.59***	6.20***	3.15***	5.66***
	(0.22)	(0.44)	(0.22)	(0.43)	(0.32)	(0.43)
Class characteristics	Y	Y	Y	Y	Y	Y
Rating FE	Y	Y	-	-	Y	Y
Rating Class FE	-	-	Y	Y	-	-
Lowest AAA only	-	-	-	-	Y	Y
Observations	61,657	67,288	61,659	67,298	42,220	44,210
Mean Prob of Downgrade	13.0%	37.4%	13.0%	37.4%	18.8%	49.4%
Effect of P75-P25 move in yield sp	2.1%	5.5%	2.4%	5.8%	3.4%	6.0%

Note: Table shows average marginal effects of logit regressions. Standard errors (shown in parenthesis) are clustered at the issue level. All regressions include quarter of origination fixed effects. Class characteristics include weighted average life (WAL), WAL squared and fixed effects for class type (e.g. accelerated security, step rate bond, planned amortization, among others). Specification 1 includes rating fixed effects and all classes available for each issue. Specification 2 includes rating class fixed effects and one dummy for (-) and another for (+) ratings. Specification 3 includes rating fixed effects and includes the lowest priority AAA tranche and all lower rated classes. The columns labeled Q407 include downgrades up to the end of 2007; the columns labeled Q308 include all downgrades up to the end of the third quarter of 2008. The last line of each panel shows the effect on the probability of downgrade of moving from the 25th to the 75th percentile in yield spreads.

Table 1.5: Downgrade, Linear Probability Model

	Specification 1		Specification 2		Specification 3	
	Q407	Q308	Q407	Q308	Q407	Q308
Yield Spread	0.037***	0.047***	0.044***	0.050***	0.047***	0.051***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
Class characteristics	Y	Y	Y	Y	Y	Y
Rating FE	Y	Y	-	-	Y	Y
Rating Class FE	-	-	Y	Y	-	-
Lowest AAA only	-	-	-	-	Y	Y
Observations	66,556	67,412	66,556	67,412	43,764	44,279

Note: Table shows coefficients of OLS regressions. Standard errors (shown in parenthesis) are clustered at the issue level. All regressions include quarter of origination fixed effects. Class characteristics include weighted average life (WAL), WAL squared and fixed effects for class type (e.g. accelerated security, step rate bond, planned amortization, among others). Specification 1 includes rating fixed effects and all classes available for each issue. Specification 2 includes rating class fixed effects and one dummy for (-) and another for (+) ratings. Specification 3 includes rating fixed effects and includes the lowest priority AAA tranche and all lower rated classes. The columns labeled Q407 include downgrades up to the end of 2007; the columns labeled Q308 include all downgrades up to the end of the third quarter of 2008.

Table 1.6: Default

	Specification 1		Specification 2		Specification 3	
	Q407	Q308	Q407	Q308	Q407	Q308
Yield Spread	0.71***	1.62***	0.72***	1.74***	0.71***	2.17***
	(0.11)	(0.18)	(0.10)	(0.17)	(0.11)	(0.24)
Class characteristics	Y	Y	Y	Y	Y	Y
Rating FE	Y	Y	-	-	Y	Y
Rating Class FE	-	-	Y	Y	-	-
Lowest AAA only	-	-	-	-	Y	Y
Observations	29,587	54,247	30,170	54,273	29,587	40,477
Mean Prob of Default	2.40%	7.80%	2.40%	7.80%	2.40%	10.40%
Effect of P75-P25 move in yield sp	0.80%	1.70%	0.80%	1.80%	0.80%	2.40%

Note: Table shows average marginal effects of logit regressions. Standard errors (shown in parenthesis) are clustered at the issue level. All regressions include quarter of origination fixed effects. Class characteristics include weighted average life (WAL), WAL squared and fixed effects for class type (e.g. accelerated security, step rate bond, planned amortization, among others). Specification 1 includes rating fixed effects and all classes available for each issue. Specification 2 includes rating class fixed effects and one dummy for (-) and another for (+) ratings. Specification 3 includes rating fixed effects and includes the lowest priority AAA tranche and all lower rated classes. The columns labeled Q407 include downgrades up to the end of 2007; the columns labeled Q308 include all downgrades up to the end of the third quarter of 2008.

Table 1.7: Rating Change

	Specification 1		Specification 2		Specification 3	
	Q407	Q308	Q407	Q308	Q407	Q308
Yield Spread	-0.25***	-0.46***	-0.28***	-0.49***	-0.32***	-0.53***
	(0.022)	(0.038)	(0.021)	(0.037)	(0.028)	(0.045)
Class characteristics	Y	Y	Y	Y	Y	Y
Rating FE	Y	Y	-	-	Y	Y
Rating Class FE	-	-	Y	Y	-	-
Lowest AAA only	-	-	-	-	Y	Y
Observations	67,412	67412	67412	67412	44279	44279
Mean change in rating	-0.60	-3.12	-0.60	-3.12	-0.90	-4.29
Effect of P75-P25 move in yield sp	-0.23	-0.43	-0.26	-0.45	-0.34	-0.57

Note: Table shows coefficients of OLS regressions where the dependent variable is the change in rating since origination. Standard errors (shown in parenthesis) are clustered at the issue level. All regressions include quarter of origination fixed effects. Class characteristics include weighted average life (WAL), WAL squared and fixed effects for class type (e.g. accelerated security, step rate bond, planned amortization, among others). Specification 1 includes rating fixed effects and all classes available for each issue. Specification 2 includes rating class fixed effects and one dummy for (-) and another for (+) ratings. Specification 3 includes rating fixed effects and includes the lowest priority AAA tranche and all lower rated classes. The columns labeled Q407 include downgrades up to the end of 2007; the columns labeled Q308 include all downgrades up to the end of the third quarter of 2008.

Table 1.8: Cohort Regressions (Year by Year)

Panel A: Downgrades

	2003 cohort		2004 cohort		2005 cohort		2006 cohort		2007 cohort	
	Q407	Q308	Q407	Q308	Q407	Q308	Q407	Q308	Q407	Q308
Yield Spread	1.51*	0.68	3.12***	4.06***	3.64***	8.29***	1.04	5.23***	4.77***	4.52***
	(0.80)	(0.66)	(0.53)	(0.65)	(0.53)	(1.20)	(1.03)	(1.61)	(1.08)	(1.67)
Class charact.	Y									
Rating FE	Y									
Observations	410	1,102	3,115	3,417	3,139	6,958	9,100	11,007	790	5,746

Panel B: Defaults

	2003 cohort		2004 cohort		2005 cohort		2006 cohort		2007 cohort	
	Q407	Q308	Q407	Q308	Q407	Q308	Q407	Q308	Q407	Q308
Yield Spread	2.55	2.50**	0.02	1.87***	1.75***	2.57***	1.58***	3.25***	0.71***	2.85***
	(1.88)	(1.15)	(0.28)	(0.43)	(0.33)	(0.46)	(0.31)	(0.75)	(0.24)	(0.94)
Class charact.	Y									
Rating FE	Y									
Observations	410	1,102	3,115	3,417	3,139	6,958	9,100	11,007	790	5,746

Note: Table shows average marginal effects of logit regressions. Standard errors (shown in parenthesis) are clustered at the issue level. All regressions include quarter of origination fixed effects. Class characteristics include weighted average life (WAL), WAL squared and fixed effects for class type (e.g. accelerated security, step rate bond, planned amortization, among others). Sample includes the lowest priority AAA class and all lower rated tranches (the equivalent to Specification 3 in the previous tables). The columns labeled Q407 include downgrades up to the end of 2007; the columns labeled Q308 include all downgrades up to the end of the third quarter of 2008.

Table 1.9: Rating by Rating

	Downgrade	Default	Rat. Change
AAA yield	0.09	0.18	-0.04
	(0.78)	(0.59)	(0.03)
Observations	11,752	1,103	11,802
Mean of LHS variable	19.5%	0.6%	-1.3
Effect of P75-P25 move in yield spread	0.1%	0.0%	0.0
AA yield	6.29***	6.46***	-0.69***
	(1.62)	(1.38)	(0.14)
Observations	9,633	5,158	9,638
Mean of LHS variable	48.4%	3.0%	-4.0
Effect of P75-P25 move in yield sp	2.8%	2.2%	-0.3
A yield	5.59***	4.77***	-0.41***
	(1.15)	(0.81)	(0.13)
Observations	8,843	7,908	8,845
Mean of LHS variable	60.7%	8.1%	-5.4
Effect of P75-P25 move in yield sp	3.9%	3.1%	-0.3
BBB yield	5.26***	5.27***	-0.72***
	(0.58)	(0.60)	(0.07)
Observations	9,662	9,658	9,663
Mean of LHS variable	68.7%	24.8%	-6.3
Effect of P75-P25 move in yield sp	6.3%	6.3%	-0.9
BB yield	5.12***	6.80***	-0.69***
	(1.09)	(1.62)	(0.14)
Observations	1,154	1,144	1,162
Mean of LHS variable	84.9%	56.5%	-8.3
Effect of P75-P25 move in yield sp	5.8%	7.7%	-0.8

Note: The first two columns show average marginal effects of logit regressions, the third column shows coefficients of OLS regressions. Standard errors (shown in parenthesis) are clustered at the issue level. All regressions include quarter of origination fixed effects and class characteristics including weighted average life (WAL), WAL squared and fixed effects for class type (e.g. accelerated security, step rate bond, planned amortization, among others). Sample includes the lowest priority AAA class and all lower rated tranches. Outcome variables include data up to the end of the third quarter of 2008. The fourth row for each rating refers to the mean of the dependent variable in each regression and the fifth row shows the effect of moving from the 25th to the 75th percentile in yield spread of each subsample.

Table 1.10: Triple-A vs Other Classes by Type

Panel A: All classes

	Q407		Q108		Q208		Q308	
	Downgrade	Rat. Change	Downgrade	Rat. Change	Downgrade	Rat. Change	Downgrade	Rat. Change
Spread of AAA	1.64*	0.01**	-6.20***	0.07***	-4.17***	0.03	0.09	-0.04
	(0.87)	(0.00)	(1.33)	(0.01)	(1.02)	(0.02)	(0.78)	(0.03)
Observations	4,466	11,802	9,662	11,802	11,285	11,802	11,752	11,802
Mean of LHS variable	2.0%	-0.06	5.0%	-0.29	10.0%	-0.61	20.0%	-1.31
Spread of non-AAA	4.85***	-0.46***	4.74***	-0.41***	7.46***	-0.67***	7.01***	-0.70***
	(0.42)	(0.04)	(0.47)	(0.04)	(0.51)	(0.05)	(0.48)	(0.06)
Observations	31,216	32,477	32,356	32,477	32,408	32,477	32,430	32,477
Mean of LHS variable	25.0%	-1.20	35.0%	-2.37	54.0%	-4.64	60.0%	-5.38

Panel B: Floating Rate Classes

	Q407		Q108		Q208		Q308	
	Downgrade	Rat. Change	Downgrade	Rat. Change	Downgrade	Rat. Change	Downgrade	Rat. Change
Spread of AAA	0.039**	0.000	-0.264***	0.006***	-0.214***	0.004	-0.049	0.002
	(0.019)	(0.000)	(0.067)	(0.001)	(0.076)	(0.003)	(0.047)	(0.003)
Observations	2,709	5,715	5,073	5,715	5,427	5,715	5,697	5,715
Mean of LHS variable	4.0%	-0.11	10.0%	-0.53	19.0%	-1.13	26.0%	-1.73
Spread of non-AAA	0.063***	-0.006***	0.057***	-0.004***	0.053***	-0.006***	0.054***	-0.007***
	(0.007)	(0.001)	(0.007)	(0.001)	(0.008)	(0.001)	(0.008)	(0.001)
Observations	21,519	21,746	21,695	21,746	21,720	21,746	21,732	21,746
Mean of LHS variable	31.0%	-1.59	45.0%	-3.12	66.0%	-5.68	70.0%	-6.37

Panel C: Fixed Rate Classes

	Q407		Q108		Q208		Q308	
	Downgrade	Rat. Change	Downgrade	Rat. Change	Downgrade	Rat. Change	Downgrade	Rat. Change
Spread of AAA	-0.53	0.00	-1.31**	0.01	-0.31	0.00	7.31***	-0.14***
	(0.50)	(0.00)	(0.56)	(0.01)	(0.29)	(0.01)	(1.32)	(0.03)
Observations	410	6,087	749	6,087	3,774	6,087	4,996	6,087
Mean of LHS variable	1.0%	0.00	1.5%	-0.07	1.5%	-0.13	16.1%	-0.91
Spread of non-AAA	3.95***	-0.04	5.40***	-0.09	9.73***	-0.66***	8.61***	-0.74***
	(0.50)	(0.06)	(0.64)	(0.08)	(0.91)	(0.10)	(0.82)	(0.10)
Observations	9,680	10,731	9,680	10,731	10,668	10,731	10,690	10,731
Mean of LHS variable	12.1%	-0.42	17.0%	-0.85	30.0%	-2.55	41.2%	-3.37

Note: Sample includes only the lowest priority AAA class and all lower ratings. Rating change information up to the quarter shown for each specification. The first column of each quarter (referring to downgrades) shows average marginal effects of logit regressions. The second column in each quarter shows coefficients of OLS regressions. Standard errors (shown in parenthesis) are clustered at the issue level. All regressions include quarter of origination fixed effects and class characteristics including weighted average life (WAL), WAL squared and class types as well as rating class fixed effects.

Table 1.11: Predictive Power of Yield Spreads by Size of the Issue

	Smallest Quartile	2nd Quartile	3rd Quartile	Largest Quartile
AAA yield	-1.50	-0.40	-2.65	5.37**
	(1.20)	(1.51)	(2.02)	(2.11)
Observations	2,974	2,936	2,749	2,726
Mean of LHS variable	15.3%	19.7%	22.6%	23.5%
Non-AAA yield	4.96***	9.43***	9.09***	10.02***
	(0.78)	(1.09)	(0.99)	(1.26)
Observations	8,779	7,961	8,087	7,571
Mean of LHS variable	46.7%	61.2%	68.4%	66.5%

Note: Table shows marginal effects of logit regressions where the dependent variable is downgrade. Standard errors (shown in parenthesis) are clustered at the issue level. All regressions include quarter of origination fixed effects and class characteristics including weighted average life (WAL), WAL squared and fixed effects for class type (e.g. accelerated security, step rate bond, planned amortization, among others). Sample includes the lowest priority AAA class and all lower rated tranches. Outcome variables include data up to the end of the third quarter of 2008. The fourth row for each rating refers to the mean of the dependent variable in each regression and the fifth row shows the effect of moving from the 25th to the 75th percentile in yield spread of each subsample.

Table 1.12: Triple-A Classes by Cohort

	Q407	Q108	Q208	Q308
2003		-6.16	-4.51**	-4.64**
		(4.25)	(2.10)	(2.10)
Obs.		597	1,146	1,146
Mean of LHS variable		3.0%	6.0%	6.0%
2004		-3.66***	-7.29***	-7.57***
		(1.00)	(1.72)	(1.79)
Obs.		2,363	2,380	2,380
Mean of LHS variable		2.0%	4.0%	4.0%
2005	0.229	-2.754	-2.792	4.436***
	(0.183)	(1.821)	(2.143)	(1.366)
Obs.	1,099	2,203	2,203	2,985
Mean of LHS variable	0.0%	2.0%	4.0%	10.0%
2006	6.010	-18.642***	-2.694	18.211***
	(4.401)	(7.142)	(4.205)	(4.968)
Obs.	1,129	1,687	2,463	2,533
Mean of LHS variable	6.0%	17.0%	21.0%	43.0%
2007	4.79	2.83	10.89*	12.11***
	(3.07)	(5.65)	(5.57)	(4.42)
Obs.	668	883	1,201	1,710
Mean of LHS variable	5.0%	13.0%	25.0%	44.0%

Note: Sample includes only the lowest priority AAA class and all lower ratings. Rating change information up to the quarter shown for each specification. The first column of each quarter (referring to downgrades) shows average marginal effects of logit regressions. The second column in each quarter shows coefficients of OLS regressions. Standard errors (shown in parenthesis) are clustered at the issue level. All regressions include quarter of origination fixed effects and class characteristics including weighted average life (WAL), WAL squared and class types as well as rating class fixed effects.

Table 1.13: Ratings vs Spreads, Downgrade probability

	AAA vs AA		AA vs A		A vs BBB		BBB vs BB	
Panel A: Median cutoff								
Higher Rat.	-0.25***	-0.07**	-0.13***	-0.05***	0.11***	0.04**	0.13***	0.03
	(0.01)	(0.03)	(0.01)	(0.01)	(0.01)	(0.02)	(0.04)	(0.05)
Obs.	9,429	9,429	5,133	5,133	4,350	4,350	576	576
R2	42%	45%	48%	51%	43%	45%	26%	38%
Panel B: 60th percentile cutoff								
Higher Rat.	-0.22***	-0.03	-0.09***	-0.02	0.14***	0.06**	0.16***	0.01
	(0.02)	(0.05)	(0.02)	(0.02)	(0.02)	(0.02)	(0.04)	(0.06)
Obs.	5,085	5,085	3,246	3,246	2,538	2,538	514	514
R2	42%	46%	49%	52%	40%	41%	24%	36%
Panel C: 70th percentile cutoff								
Higher Rat.	-0.16***	0.07	-0.03	0.00	0.16***	0.11***	0.15***	0.05
	(0.02)	(0.09)	(0.02)	(0.02)	(0.03)	(0.04)	(0.05)	(0.07)
Obs.	1,735	1,735	1,832	1,832	1,319	1,319	449	449
R2	45%	50%	52%	55%	34%	35%	23%	37%
Panel D: 80th percentile cutoff								
Higher Rat.	-0.05*	0.08	0.04	0.06**	0.17***	0.19***	0.19***	0.05
	(0.03)	(0.12)	(0.03)	(0.03)	(0.04)	(0.07)	(0.06)	(0.09)
Obs.	525	525	1,013	1,013	650	650	370	370
R2	28%	37%	49%	51%	35%	37%	24%	40%
Panel E: 90th percentile cutoff								
Higher Rat.	0.05	0.35***	0.04	0.04	0.26***	0.16*	0.36***	0.21*
	(0.04)	(0.13)	(0.04)	(0.04)	(0.06)	(0.10)	(0.07)	(0.11)
Obs.	317	317	586	586	350	350	289	289
R2	26%	39%	45%	48%	34%	39%	29%	45%
Controls	N	Y	N	Y	N	Y	N	Y

Note: Table shows coefficients on a dummy for the higher rating in an OLS regression. A negative coefficient means that the higher rated securities experienced fewer downgrades until the end of Q3:2008 and a positive coefficient means that the higher rated securities were downgraded more. The first column of each ratings pair shows unconditional OLS (essentially a t-test) and the second colum includes controls. Standard errors (shown in parenthesis) are clustered at the issue level. All regressions include quarter of origination fixed effects. Controls include weighted average life (WAL), WAL squared and fixed effects for class type as in the previous tables.

Table 1.14: Fixed and Floating Triple-A IV Approach

	Q407	Q108	Q208	Q308
Floater (OLS)	0.007	-0.053***	-0.067***	-0.027
	(0.007)	(0.015)	(0.023)	(0.032)
Fixed (OLS)	0.000	-0.001	-0.003*	0.006
	(0.000)	(0.001)	(0.002)	(0.005)
Floater (IV)	0.027	-0.115**	-0.158*	0.010
	(0.040)	(0.056)	(0.089)	(0.101)
Fixed (IV)	0.000	0.000	0.001	0.084***
	(0.001)	(0.001)	(0.004)	(0.015)

Note:The rows labeled IV refer to two-stage least squares regressions where the triple-A yield spreads are instrumented for using the spreads of double-A, A and triple-B classes. Sample includes only the lowest priority AAA class. Rating change information up to the quarter shown for each specification. Standard errors (shown in parenthesis) are clustered at the issue level. All regressions include quarter of origination fixed effects and class characteristics including weighted average life (WAL), WAL squared and fixed effects for class type (e.g. accelerated security, step rate bond, planned amortization, among others).

Table 1.15: Duration (Cox) Model

	Haz. Ratio	Std. Error	t-statistic	Number of classes
Whole sample	1.233	0.018	14.13	44,230
AAA	0.949	0.056	-0.89	11,802
Non-AAA	1.241	0.018	14.63	32,428

Note:Table reports hazard ratios of a Cox model where failure is a downgrade. The period at risk starts in the middle of 2007 and each period is one quarter (i.e. securities are at risk for at most 5 quarters). Sample includes only the lowest priority AAA class. Standard errors are clustered at the issue level. All regressions include quarter of origination fixed effects and class characteristics including weighted average life (WAL), WAL squared and fixed effects for class type (e.g. accelerated security, step rate bond, planned amortization, among others).

Chapter 2

Why Don't Lenders Renegotiate More Home Mortgages? The Role of Securitization

2.1 Introduction

Many commentators have attributed the severity of the foreclosure crisis in the United States in the 2007–2009 period to the unwillingness of lenders to renegotiate mortgages, and, as a consequence, have placed renegotiation at the heart of the policy debate. Several major policy actions to date have involved encouraging lenders, in one way or another, to renegotiate loan terms in order to reduce borrower debt loads. The Obama Adminstration's Making Home Affordable Plan, announced in February of 2009, provided financial incentives to servicers to renegotiate loans on the condition that the lenders reduce the interest rate for a significant period of time.[1]

The appeal of renegotiation to policy makers is simple to understand. Foreclosure can be a costly outcome post delinquency for both the borrower and the lender. If a lender makes a concession to a borrower by renegotiating a loan this can be beneficial to both parties, as long as the value to the lender of the new contract exceeds the value of the property in foreclosure. According to proponents, renegotiation of home mortgages is a type of public policy holy grail, in that it helps both borrowers and lenders at little or no cost to the government (some examples of this view include Congressional Oversight Panel (2009) and Geanakoplos and Koniak (2008)). However, if this logic is correct, lenders should find renegotiation attractive, even in the absence of government prodding. Yet, we observe very little renegotiation—less than two percent of all delinquent loans receive payment reducing modifications in the first year following delinquency, whereas approximately 30 percent of delinquent loans are in foreclosure one year after missing two mortgage payments.

The leading explanation attributes the reluctance of lenders to renegotiate to the process of securitization.

> Unfortunately, this win-win solution is not possible today. Your mortgage has been sold and repackaged in an asset-backed security pool and sold in tranches

[1] See "$275 Billion Plan Seeks To Address Crisis In Housing," *New York Times*, Feb. 18, 2009.

> with different priorities. There is disagreement on who has the right to renegotiate and renegotiation might require the agreement of at least 60% of the debt holders, who are spread throughout the globe. This is not going to happen. Furthermore, unlike your local bank, distant debt holders cannot tell whether you are a good borrower who has been unlucky or somebody just trying to take advantage of the lender. In doubt, they do not want to cut the debt for fear that even the homeowners who can easily afford their mortgage will ask for debt forgiveness. (Zingales 2008)

With securitization the claims on the cash flows from large pools of mortgages become dispersed. The existing theory on dispersed debt ownership suggests that this makes renegotiation following default harder (Bolton and Scharfstein 1996 and Asquith, Gertner, and Scharfstein 1994). For loans that are held on banks' balance sheets this problem should not arise, leading many to conclude that the existence of securitization trusts is one of the key obstacles to mortgage renegotiation.

More precise institutional evidence appears to confirm the role of securitization in impeding renegotiation. As mentioned in more detail below, PSAs do sometimes place global limits on the number of modifications a servicer can perform for a particular pool of mortgages. In addition, the rules by which servicers are reimbursed for expenses may provide a perverse incentive to foreclose rather than modify. Furthermore, because servicers do not internalize the losses on a securitized loan, they may not behave optimally. Another issue is the possibility that those investors whose claims are adversely affected by modification will take legal action. Finally, historically, SEC rules have stated that contacting a borrower who is fewer than 60-days delinquent constitutes an ongoing relationship with the borrower and jeopardizes the off-balance sheet status of the loan.

But some market observers express doubts about the renegotiation-limiting role of securitization. One important difference between dispersed corporate debt and mortgage-backed securities is that delinquent borrowers have to negotiate with multiple parties in the case of corporate debt, whereas securitization trusts have an agent (the servicer) who acts on their behalf. Hunt (2009) conducted an exhaustive review of a sample of PSAs and concluded, "it appears that large-scale modification programs may be undertaken without violating the plain terms of PSAs in most cases." Although some servicers have expressed concern about lawsuits, of the more than 800 lawsuits filed by investors in subprime mortgages through the end of 2008, not one involved the right of a servicer to modify a loan.[2] Even the Congressional Oversight Panel (2009), which did view securitization as a problem in general, conceded, "The specific dynamics of servicer incentives are not well understood." Finally, the SEC ruled in 2008 that if default was "reasonably foreseeable," then contact with a borrower prior to 60-day delinquency would not affect the accounting status of the loan.

In this paper, we explore the role of securitization in preventing the renegotiation of home mortgages using a dataset from Lender Processing Services (LPS), a large, detailed sample of residential mortgages.[3] Measuring renegotiation in the LPS data is a challenge because there is no field in the data that identifies whether or not a servicer has changed the terms of, or "modified," the loan. We overcome this difficulty by developing an algorithm to

[2]Navigant report, Congressional Oversight Panel (2009).

[3]Until 2008, the dataset was known as McDash.

identify modifications and we validate this algorithm on an unrelated dataset that includes a modification flag.

The LPS dataset includes loans that are serviced for private securitization trusts that are not sponsored by any of the government sponsored enterprises (GSEs), so-called "private-label" loans, which are subject to all of the contract frictions described above. It also includes loans owned by servicers, so-called "portfolio" loans, which are immune to such problems. We compare renegotiation rates of seriously delinquent loans (i.e. that missed two or more payments) in the two groups, controlling for observable characteristics of the loans. Our LPS data run from the first quarter of 2005 through the third quarter of 2008, and thus has the benefit of covering much of the housing and foreclosure crisis period, while excluding the period of dramatic government intervention in financial markets that occurred beginning in September of 2008 with the conservatorship of the housing GSEs and the failure of Lehman Brothers that was followed by the large-scale capital injections in the banking sector from the Troubled Asset Relief Program (TARP).

Our empirical analysis provides strong evidence against the role of securitization in preventing renegotiation. For our narrowest definition of renegotiation, payment-reducing modifications[4], we find that the differences in the likelihood of renegotiation in the 12 months subsequent to the first 60-day delinquency between the two types of loans is neither economically nor statistically significant. When we consider a broader definition that includes any modification, the data even more strongly reject the role of securitization in preventing renegotiation. It is worth noting that the pooling and servicing agreements (PSAs), which govern the conduct of servicers when loans are securitized, often place limits on the total number of modifications a servicer can perform, so we would expect the metric of "all modifications" to be most affected by securitization. We also find no differences between private-label and portfolio loans when we include in our definition of renegotiation the transactions whereby lenders allow borrowers to extinguish their liabilities by repaying less than the outstanding balance of the loan.[5]

Our results are highly robust. One potential problem with the data is that there is unobserved heterogeneity in the characteristics of portfolio and private-label loans. To address this, we exploit subsets of the LPS data, in which servicers provide an exceptional amount of information about borrowers. When we exclude observations where the servicer failed to report whether the borrower fully documented income at origination, or what the debt-to-income ratio was at origination, our results even more strongly reject the role of securitization in inhibiting renegotiation. When we focus only on loans for which the borrower fully documented income, we obtain results that are broadly consistent or, in some cases, stronger than the results for the full sample. Finally, we limit our sample to only subprime loans (as defined in LPS). These loans comprise only 7 percent of the LPS data, but they account for more than 40 percent of all serious delinquencies and almost 50 percent of the modifications that we identify in the data. The results that we obtain for the subprime sample are also consistent with our results for the full sample.

Another potential issue with our focus on 60-day delinquent loans is that portfolio lenders

[4]These include any changes to the contract terms that reduce the borrower's monthly payment.

[5]These transactions are known as short payoffs, short sales, or deeds-in-lieu of foreclosure, depending on the structure. We measure this component of renegotiation by counting the number of seriously delinquent loans that the servicer reports as "paid off."

can contact borrowers at any time, whereas some securitization agreements forbid lenders from contacting borrowers until they are at least 60 days delinquent (two missed payments). When we shift our focus to 30-day delinquent borrowers (one missed payment), our results continue to show no meaningful difference between renegotiation of private-label and portfolio loans.

One other possibility is that our algorithm for identifying modifications is somehow missing a class of loss-mitigation actions taken by servicers. Forbearance agreements and repayment plans, for example, would not necessarily show up in our data. However, neither of these actions constitutes renegotiation in any classic sense, because the lender still expects the borrower to repay in full, including interest on any delayed payment. In addition, unlike modifications, PSAs never place any limits on the use of forbearance agreements or repayment plans, so, *a priori*, we would have less reason to expect a difference in their use across private-label and portfolio loans. Finally, most successful forbearance agreements conclude with a modification to allow the borrower to repay the arrears incurred in forbearance. With all of that said, we test the proposition that servicers engage in other loss mitigation actions by looking at the "cure rate." This is the percentage of loans that transition to current status in the year subsequent to becoming 60-days delinquent. We find that in the full sample, private-label loans are unconditionally less likely to cure, but that when we correct for observable characteristics, we estimate a cure rate of around 55 percent for the typical portfolio loan and a cure rate about two percentage points *more* for an otherwise equivalent private-label loan. This difference is similar in other subsamples of the data. For example, in the subsample of subprime loans, the subsample with information about documentation and debt-to-income (DTI) status, and the sample of fully documented loans, we again find that private-label loans are significantly more likely to cure (both unconditionally and conditioning on observable characteristics).

The policy debate has focused exclusively on the manner in which securitization impedes renegotiation and implicitly assumes that portfolio lenders face no institutional impediments, but this is not realistic. Portfolio lenders complain about accounting rules, including the need to identify modifications, even when the borrowers are current on their mortgage payments prior to the modification, as "troubled debt restructurings", which leads to reduction of the amount of Tier II capital and increased scrutiny from investors and cumbersome accounting requirements. The shortage of qualified staff, an oft-heard complaint from borrowers seeking renegotiation, affects servicers of portfolio loans and private label loans equally. Finally, the interests of the managers of a loan portfolio are not necessarily any more likely to be aligned with their investors than are the interests of the trustees of a mortgage pool; many have attributed the catastrophic failures of financial institutions like AIG in 2008 to misaligned incentives of managers and shareholders.

Our results are consistent with the hypothesis that securitization does impede renegotiation but that a different set of impediments leads to similar problems with portfolio loans and generates our finding that there is no difference. However, the small differences would represent a remarkable coincidence.[6] More importantly, the low overall levels of renegotiation

[6]Yet another possible explanation is that equal treatment provisions in PSAs force servicers to modify similar numbers of portfolio and private-label loans and that servicers are reluctant to modify portfolio loans in spite of the fact that they internalize the benefits because they must then modify private label loans for which they do not.

mean that even if contract frictions cut the overall number of concessionary modifications in half, 94 percent of seriously delinquent borrowers would still fail to receive a concessionary modification. So the puzzle remains why so few loans are renegotiated.

To our knowledge, this paper is the first to estimate directly the likelihood of renegotiation of private-label and portfolio-held mortgages. Piskorski, Seru, and Vig (2009) address the question of the effects of securitization on renegotiation, but rather than directly identifying renegotiation, they argue that observed differences in foreclosure rates *imply* differences in renegotiation activity. Our results contradict this interpretation. For renegotiation to explain the differences in foreclosure rates, there would have to be large errors in our algorithm for identifying renegotiation, and those errors would have to be significantly biased toward portfolio loans, a possibility that is particularly problematic given that the renegotiations we focus on are precisely the type that PSAs supposedly prevent. In addition, most of the loan histories in the LPS sample are right-censored, meaning that the borrowers have neither lost their homes nor paid off their mortgages when the data end, making it impossible to equate the absence of a foreclosure with successful renegotiation. By contrast, a "cure" is a necessary condition for renegotiation, and thus the differences we report in cure rates across portfolio and private-label loans that are neither large nor of consistent sign contradict the claim that securitization is a major obstacle to renegotiation.

If contract frictions induced by securitization are not a significant problem, then what is the explanation for why lenders do not renegotiate with delinquent borrowers more often? Researchers have puzzled over this since well before the advent of non-agency securitization and the the subprime crisis. Wang, Young, and Zhou (2001) write that:

> ...Given the literature, a bank should be more likely to work out a solution with the mortgage borrower than to foreclose on the property in the case of a mortgage default. However, Riddiough and Wyatt (1994b) observe that in dealing with mortgage defaults, many lenders take a hard-line approach and are reluctant to grant workout concessions." (pp. 960-1)

Wang, Young, and Zhou (2001) and Riddiough and Wyatt (1994b) both argue that information issues make foreclosure a better option from the lender's perspective and we view this as a plausible explanation.

The first way in which information issues affect the decision to modify is through the uncertainty that the lender faces about the outcomes of each loan, namely the possibility that borrowers will "self-cure" or that borrowers will re-default after receiving a modification. As we mentioned above, about 55 percent of seriously delinquent borrowers in our sample "cure" without receiving a modification; if taken at face value, this means that, in expectation, 55 percent of the money spent on a given modification is wasted.. Regarding the possibility of redefault, our results show that a large fraction of borrowers who receive modifications end up back in serious delinquency within six months. For them, the lender has simply postponed foreclosure; in a world with rapidly falling house prices, the lender will now recover even less in foreclosure. Documentary evidence shows that lenders explicitly took both self-cure risk and redefault risk into account when making their renegotiation decisions.

The second problem is that borrowers have private information about their willingness and ability to repay the loan and the value of the collateral. In making a renegotiation decision, lenders must weigh the benefits of preventing a particular foreclosure with the cost of providing that same concession to all observationally equivalent borrowers. Suppose the lender knows that some of the seriously delinquent borrowers need a 30% reduction in principal to make repayment attractive but others only need a 10% reduction. The lender may only offer 10% because the cost of giving 30% reductions to borrowers who only need 10% offsets the benefit of fewer foreclosures. Furthermore, the lender may worry that more generous renegotiation will induce borrowers currently making payments to join the ranks of the seriously delinquent.

In sum, our research suggests that in understanding distressed debt renegotiation, the case of residential mortgages and other debts need to be viewed differently. As Wang, Young, and Zhou (2001) write:

> ...previous models [of debt renegotiation] addressed the coordination and agency issues when a distressed firm faces multiple debt holders. Our model addresses the problems when one debt holder faces many potential defaulters. In this regard, our model is most relevant to an environment where a debt holder could face a large group of potential defaulters, as in a residential property market with falling prices... (p. 962)

Our results provide some insight into the perceived failure of the the U.S. governments Home Affordable Modification Program (HAMP) to stem the tide of foreclosure in 2009 and 2010.[7] HAMP was based on the premise that the barriers to renegotiation were largely institutional and not economic – i.e. renegotiation was in the interests of the owners of the loans but institutional barriers including securitization were preventing them from occuring. To address these institutional issues, HAMP diverted significant resources toward the servicers as opposed to the investors.[8] Subsequent updates to HAMP implicitly recognized this by providing increased financial incentives to investors to renegotiate loans.[9]

2.2 Related Literature and Existing Evidence

Our research draws on existing literature in several different fields. First, there has been substantial interest in the question of renegotiation of home mortgages among real estate economists, both prior to, and as a result of the current crisis. Riddiough and Wyatt (1994a), Riddiough and Wyatt (1994b), Ambrose and Capone (1996), and Wang, Young, and Zhou (2001) addressed informational issues that inhibit efficient renegotiation. We draw extensively on this research in Section 2.5. Springer and Waller (1993), in an early example, explores patterns in the use of forbearance as a loss mitigation tool. Capone (1996) and Cutts and Green (2005) both discuss the institutional issues, with the former study providing historical evidence and focusing on issues in the mid-1990s, and the latter study discussing innovations since then.

The issue of dispersed ownership and debt renegotiation has received a fair amount of attention in the corporate finance literature. Gan and Mayer (2006), for example, focus on commercial mortgages, and find that servicers delay liquidation of delinquent mortgages

[7] "U.S. Loan Effort Is Seen as Adding to Housing Woes," *New York Times*, Jan 1, 2010.
[8] http://www.ustreas.gov/initiatives/eesa/homeowner-affordability-plan/FactSheet.pdf
[9] "Mortgage Plan Remodeled Again ," *Wall Street Journal*, March 29, 2010.

when they are also the holders of the equity tranche of the deal. This suggests that participating in the losses due to liquidation may alleviate some of the agency problems posed by the separation of ownership and servicing pointed out before. However, it may also lead to conflicts of interest between holders of different tranches. In their setting, Gan and Mayer (2006) find that the servicers' behavior is consistent with asset substitution, as servicers seek to benefit from the option-like payoff of their position. Also, the contractual restrictions imposed by PSAs (discussed above) and standard economic arguments on the effects of dispersed ownership of debt (as in Bolton and Scharfstein 1996 and Asquith, Gertner, and Scharfstein 1994) further reduce the incentives of servicers to modify mortgages.

The start of the subprime crisis in 2007 led to a resurgence of interest in the topic of modifications among economists and aroused new interest from other fields, in particular, the field of law. Legal researchers, White (2008) and White (2009), for example, have addressed empirical questions about the frequency and characteristics of loan modifications, closely related to the analysis in this paper. In addition, they have also looked at issues related to the restrictions imposed by contracts (Hunt 2009 and Gelpern and Levitin 2009) and the interactions among foreclosure, renegotiation, and personal bankruptcy (Levitin 2009a and Levitin 2009b). In real estate, Quercia, Ding, and Ratcliffe (2009), Cutts and Merrill (2008), Stegman, Quercia, Ratcliffe, Ding, Davis, Li, Ernst, Aurand, and Van Zandt (2007), and Mason (2007), all discuss issues with contemporary loss mitigation approaches.

More broadly, previous research explored the factors that lead delinquent mortgages to transition to foreclosure or to cure, one of which is renegotiation. Pre-crisis papers include Ambrose and Capone (1998), Ambrose, Buttimer Jr., and Capone (1997), Ambrose and Capone (2000), Lauria, Baxter, and Bordelon (2004), Danis and Pennington-Cross (2005), Pennington-Cross (2009), and Pennington-Cross and Ho (2006). Mulherin and Muller (1987) discusses conflicts between mortgage insurers and owners that may lead servicers to induce or postpone foreclosure inefficiently. In light of the crisis, Piskorski, Seru, and Vig (2009) and Cordell, Dynan, Lehnert, Liang, and Mauskopf (2008a) have revisited the question.

2.3 Data

We use a dataset constructed by LPS. This is a loan-level dataset that covers approximately 60 percent of the U.S. mortgage market and contains detailed information on the characteristics of both purchase-money mortgages and mortgages used to refinance existing debt.[10] This dataset is especially useful in the context of this paper, as it includes both securitized mortgages and loans held in portfolio.[11] The LPS data specifically denote whether a mortgage is held in portfolio, or securitized by a non-agency, private institution.[12] If institutional constraints are restricting the modification process for private-label, securitized loans,

[10]We use a 10 percent random sample of the LPS data when estimating all of our empirical models. The dataset is simply too big to use in its entirety from a computational standpoint. However, we have checked the robustness of our results to using different sample sizes, and we do not find substantial differences.

[11]For a more detailed discussion of the LPS data, we direct the reader to Foote, Gerardi, Goette, and Willen (2009).

[12]The LPS data also denote when a loan is securitized by a GSE (Government Sponsored Enterprise) such as Freddie Mac or Fannie Mae. We eliminate this class of loans, since the GSEs hold all credit risk, and thus are not subject to any modification restrictions.

we would expect to see relatively few modifications among them, as compared to portfolio loans. Unfortunately, our LPS sample does not include direct information regarding loan modifications.[13] However, LPS does provide monthly updates to loan terms, so it is possible to identify loan modifications indirectly (and imperfectly). Table 2.3 shows two examples of modifications in the data. In the first example, the servicer cuts the interest rate, capitalizes arrears into the balance of the loan, and extends the term of the loan to 40 years. In the second example, the servicer just capitalizes arrears into the balance of the loan. In both cases the loan is reported as "current" after the modification, whereas before it was 90+ days delinquent.

We denote a loan as being modified if there is a change in its terms that was not stipulated by the initial terms of the contract. Such modifications include interest-rate reductions, principal-balance reductions, and term extensions. We can also identify principal-balance and mortgage-payment *increases* that reflect the addition of arrears into the balance of a loan.[14] We spell out our algorithm for identifying modifications in more detail in Appendix A.

There are two potential mistakes we can make in this exercise. First, we may falsely identify modifications ("false positives") because of measurement error in the data (for example, a mistake in the updated balance or interest rate) or some endogenous behavior on the part of the borrower (for example, a borrower making extra principal payments). Second, we could miss modifications ("false negatives") because our algorithm for finding modifications is incomplete. In order to test our algorithm, we use data from the Columbia files put together by Wells Fargo's CTSLink service. This dataset includes a similar set of variables to those in the LPS dataset (on performace of the loans and characteristics of the borrower at origination) but is limited to private-label loans. These files do include, however, explicit flags for modifications. This allows us to use the same algorithm described in Appendix A and compare the modifications we identify to the "true" modifications. Results are reported in Table 2.4. Overall our algorithm performs well, with 17 percent false negatives (that is, we do not identify around 17 percent of the "true" modifications) and around the same percentage of false positives (that is, approximately 17 percent of the modifications we identify are not flagged as modifications on the CTSLink data). By type of modification, our algorithm performs best for principal reductions, term increases, and fixed-rate mortgage reductions, and comparatively worse for ARM rate reductions and for principal increases.

We explore several different definitions of renegotiation in the data. Our first definition of "renegotiation" is concessionary modifications that serve to reduce a borrower's monthly payment. These may be reductions in the principal balance or interest rate, extensions of the term, or combinations of all three. This definition of renegotiation is a key focus of our

[13]The Office of Thrift Supervision (OTS), in collaboration with the Office of the Comptroller of Currency (OCC), used data from LPS to analyze the outcomes of recent mortgage modification programs (OCC and OTS Mortgage Metrics Report, Third Quarter 2008). In this report, they had access to supplementary data from servicers that include the identification of loans in the LPS data that had been modified. We have not been able to obtain access to this data.

[14]One of the major types of loan modifications that we are largely unable to identify are interest rate freezes for subprime ARMs, which reset after two or three years. However, the reason that we cannot identify those freezes is because many are not binding; the fully-indexed rate is lower than the initial rate. These modifications will have no major effect on the current terms of the mortgage, so we do not view this as a major drawback.

analysis because there is a consensus among many market observers that concessionary modifications are the most, or possibly the only, effective way of preventing foreclosures. As the Congressional Oversight Panel (COP) for the Troubled Asset Recovery Program (TARP) has written, "Any foreclosure mitigation plan must be based on a method of modifying or refinancing distressed mortgages into affordable ones. Clear and sustainable affordability targets achieved through interest rate reductions, principal write-downs, and/or term extensions should be a central component of foreclosure mitigation."[15]

Because the pooling and servicing agreements (PSAs), which govern the conduct of servicers when loans are securitized, often place limits on the number of modifications a servicer can perform, we broaden our definition of renegotiation to include any modification, regardless of whether it lowers the borrower's payment. Modifications are often thought to always involve concessions to the borrower, but many, and in some subsets most, modifications involve the capitalization of arrears into the balance of the loan, and thus lead to increased payments.

Finally, we attempt to include in our definition of renegotiation the transactions whereby lenders allow borrowers to extinguish their liabilities by repaying less than the outstanding balance of the loan.

We focus on loans originated after January of 2005, in line with most studies that use the LPS dataset.[16] We restrict attention to loans that have missed two or more payments and track the performance of those loans (and whether they were modified) until the end of the third quarter of 2008. Given that the purpose of this paper is to test for the effect of contract frictions and dispersed debt ownership on the renegotiation of mortgages, we choose to exclude the period following September of 2008. While the financial crisis of 2008-2009 and the bankruptcy of Lehman Brothers may be an interesting setting to study the behavior of financial institutions, the post-September 2008 period is also a time when political intervention in financial markets is likely to have significantly changed the incentives of banks and servicers. Our interest in this paper is to test for the effect of the incentives put in place by securitization and not the interaction of those incentives with political intervention and financial distress. The advantage from a research perspective of the period we consider is that it includes an early time period where house prices are increasing and a later time period (from December of 2006) where national house prices declined significantly.

We exclude loans that enter the database more than three months after being originated, as well as those with missing FICO, origination amount, interest rate information, and LTV, as well as those with missing information about whether they are subprime or prime and whether they are purchase or refinance loans.

[15]See the Congressional Oversight Panel (2009). This view is widely held and is the main focus of the Administration's Making Home Affordable foreclosure prevention plan was to encourage servicers to modify loans to reduce monthly payments to 31 percent of income.

[16]LPS added a few large national servicers in January of 2005, which created an attrition bias for the mortgage data prior to 2005. When a servicer enters the LPS dataset it only provides information on active loans. Thus, if one uses data prior to 2005 from these services, it will include only the mortgages that survived until 2005.

2.3.1 Summary Statistics from the Data

Table 2.5 reports the number of modifications performed each quarter from the first quarter of 2007 through the third quarter of 2008, disaggregated by the type of modification. Each of the numbers is a multiple of 10 because we used a 10 percent random sample and scaled up the numbers we found. The first column of Table 2.5 simply reports the total number of loan modifications made. Not surprisingly, modifications have become more common as the housing market has weakened. There appear to be more than 6 times as many modifications performed in the third quarter of 2008 as in the first quarter of 2007. In addition to the rapid growth in loan modifications, the composition of modifications has changed over time. This can be seen in the remaining columns of Table 2.5, which list the incidence of modifications of different types.[17]

An interesting finding is that most modifications entailed *increases* in the principal balance of a mortgage. Such increases are likely due to the addition of arrears to the outstanding mortgage balance for delinquent borrowers, and these often increase the monthly mortgage payment by a nontrivial amount. While the absolute numbers of balance-increasing modifications are still rising, they are falling as a percentage of total modifications. In the last few quarters, interest-rate reductions, which necessarily involve a decrease in the monthly mortgage payment, have become more frequent, rising to more than 20 percent of all modifications performed in 2008:Q3. Table 2.5 provides further information regarding the behavior of monthly mortgage payments for loans that have undergone a modification. There are several notable patterns in this table. First, as of 2008:Q3, modifications that involved payment decreases were more common than those that involved payment increases. Furthermore, the average and median magnitude of payment decreases has recently increased in our sample. From 2007:Q1 to 2008:Q2, the median payment decrease ranged from approximately 10 percent to 14 percent, but then increased to approximately 20 percent in 2008:Q3. Based on the logic from our simple framework above, it is likely that these will have more success than modifications involving increases in the payment and/or balance.

Another interesting observation from Table 2.5 is that the incidence of principal reductions is quite low in our data. This is likely due to two factors. First, the LPS dataset under-represents the subprime mortgage market.[18] A few servicers that focus almost exclusively on subprime mortgages have recently begun modification programs that involve principal reduction.[19] In addition, from a theoretical perspective, principal reduction plans suffer from the severe incomplete-information problem noted earlier. Balance reductions are appealing to both borrowers in danger of default and those who are not. In a recent paper, we argued that to avoid such moral hazard concerns, lenders have a strong incentive to only

[17] In many cases a mortgage will experience multiple types of modifications at the same time. For example, we see cases in the data in which the interest rate is decreased and at the same time the term of the loan is extended. Thus, the percentages in Table 2.5 are not calculated with respect to the number of loans modified, but rather with respect to the number of modifications performed.

[18] The majority of subprime mortgages are securitized by non-agency firms, and the LPS dataset includes approximately 35 percent of mortgages securitized by non-agency corporations.

[19] According to an October 2008 report by Credit Suisse, Ocwen Loan Servicing, LLC and Litton Loan Servicing LP were the only subprime servicers that had performed a nontrivial number of principal reduction modifications. Neither of these servicers contributes to the LPS dataset.

provide modifications to those borrowers who are most likely to default.[20] Table 2.5 contains summary statistics regarding the characteristics at origination of both the sample of modified mortgages and the sample of all loans in the LPS dataset. The patterns that emerge from the table are consistent with such an argument. We discuss this point in more detail below. The sample of modified mortgages is characterized by substantially lower credit scores, higher loan-to-value (ltv) ratios, and slightly higher debt-to-income ratios. The discrepancy in ltv ratios may be underestimated, as the percentage of mortgages with an ltv ratio of exactly 80 percent is significantly higher in the modification sample than in the full sample. This likely implies a larger fraction of highly leveraged loans, for which the second liens are not observable in the data. In addition, the modification sample includes a higher fraction of mortgages with non-traditional amortization schedules, such as interest-only loans, option ARMS, hybrid ARMs, and subprime loans.

In Table 2.6 we compare the size of payment decrease and payment increase modifications for loans held in private-label trusts and loans held in portfolio. The size (as a percentage of the original payment) of the median payment decrease due to modification is larger for private-label loans in the first three quarters of 2008. We see a similar pattern for the median payment increase due to modification, while the differences are small for the mean and median payment increase.

2.4 Differences in Modification Behavior

In this section, we directly address the question of whether the incidence of modification is impeded by the process of securitization. We show evidence that private-label loans and portfolio loans perform similarly, both unconditionally and when observable differences between securitized and portfolio-held loans are controlled for, using both a logit model with a 12-month horizon and a Cox proportional hazard model that takes into account the problem of right censoring in the data.

To make sure that our results are robust to the type of modification performed, we use several different definitions of modification in this section. Our first measure is the number of concessionary modifications, which we define as reductions in the interest rate, reductions in the principal balance, extensions of the term, or combinations of all three. Any or a combination of these serves to reduce a borrower's monthly mortgage payment. We use this as our primary definition of modification in our analysis, as there is a consensus among most market observers that concessionary modifications are the most, or perhaps the only, effective way of preventing foreclosures. Because pooling and servicing agreements, which govern the conduct of servicers when loans are securitized, often limit modifications that change *any* of the contract terms (not just those that result in payment decreases), we broaden our definition of renegotiation to include any modification, regardless of whether it lowers the borrower's payment. As we discussed above, many, and in some subsets, most modifications, involve the capitalization of arrears into the balance of the loan and thus lead to increased payments. Finally, we attempt to include in our measure of renegotiation the number of times that lenders allow borrowers to extinguish their liabilities by repaying less than the outstanding balance of the loan. These transactions are known as short payoffs, short sales,

[20]See Foote, Gerardi, and Willen (2008) for a more detailed discussion.

or deeds-in-lieu of foreclosure, depending on the structure. We do this by counting the number of seriously delinquent loans that the servicer reports as paid off, and including these observations in our definition of modification.

Before turning to the regressions, however, it is instructive to look at the unconditional frequencies of modifications in the data. Panel A of Table 2.7 shows the unconditional frequencies for each type of investor. The first takeaway from the table is the extremely low percentages of modifications for *both* types of mortgages. Less than 2 percent of 60-day delinquent loans received concessionary modifications in the 12 months following the first serious delinquency, and only 8.6 percent of the delinquent loans received *any* type of modification in the same period. These are low levels of modifications compared to the frequency of foreclosure (foreclosure proceedings were initiated on approximately half of the loans in the sample and completed for almost 30 percent of the sample.), and they suggest that even if there are contract frictions that are preventing modifications in securitized trusts, the economic effects are small.

The second takeaway from the table is that the unconditional differences between portfolio loans and private-label loans are very small in absolute terms. There is a difference of approximately 0.8 percentage points and 0.1 percentage points for concessionary modifications and all modifications, respectively. These are very small differences, and they suggest that contract frictions do not play an important role in inhibiting the renegotiation process for loans in securitized trusts. However, these are unconditional statistics, and it is possible that once observable differences in the characteristics of each type of loan and borrower are accounted for, the results may change.[21] Thus, we now estimate differences in modification behavior while controlling for observable loan and borrower characteristics. These characteristics include the contract interest rate at origination; the credit score of the borrower at origination; the loan-to-value ratio of the mortgage (not including second or third liens) at origination[22]; the logarithm of the nominal dollar amount of loan; an indicator of whether the purpose of the loan was a refinance of a previous mortgage or a home purchase; an indicator of whether the loan was considered to be subprime[23]; a measure of the amount of equity in the property at the time of delinquency, specified as a percentage of the original loan balance and updated by state-level house price indexes calculated by the Federal Housing Finance Agency (FHFA)[24] (and an indicator for a borrower who is in a position of negative equity at the time of delinquency, where the value of the mortgage exceeds the value of the home); and the unemployment rate of the county in which the borrower resides, calculated

[21]For example, if private-label loans are significantly riskier, and thus better candidates for modification on average, then the unconditional difference will significantly understate things.

[22]Because of the lack of information on second liens in the LPS data and the prevalence of second mortgages as a way to avoid paying mortgage insurance, we include an indicator variable if the ltv ratio is exactly equal to 80 percent. These are the borrowers who likely took out second mortgages, as the requirement for mortgage insurance occurs at ltv ratios above 80 percent. Our experience with other, more complete datasets also confirms that many of these borrowers are likely to have second mortgages that bring the cumulative ltv ratio up to 100 percent.

[23]This definition of subprime comes from the mortgage servicers that contribute to the LPS dataset.

[24]House prices are measured at the state level using the FHFA index. We also tried using Case-Shiller metropolitan area house price indexes and found no substantive differences. We chose to use the FHFA prices for our primary specifications because of their greater sample coverage.

by the Burea of Labor Services (BLS).[25] We also include, but do not show because of space considerations, a set of cohort dummies that control for the quarter when the mortgage was originated, information regarding the amortization schedule of the mortgage (interest-only or negative amortization, including mortgages commonly referred to as option ARMs), an indicator for whether the size of the mortgage is greater than the GSE conforming loan limits, an indicator for whether the house is a primary residence, an indicator for adjustable rate mortgages that contain a reset provision (so-called "hybrid ARMs"), and, finally, an indicator for a borrower who does not use the corresponding property as a principal residence (this includes both properties used strictly for investment purposes, and vacation homes).

2.4.1 Canonical Specification Results

Panel B of Table 2.7 displays the estimated marginal effects from a set of logit models for the three different types of modification definitions. The dependent variable is 1 if a 60-day deliquent loan is modified at any point in the 12 months following the first delinquency.[26] The first column considers payment-reducing (concessionary) modifications, the second column includes both payment-reducing and payment-increasing modifications, and the third column contains all modifications considered before, as well as prepayments. In all regressions, the group of portfolio-held loans is omitted from the estimation and is thus assumed to be the reference group. We cluster the standard errors at the zip code level to account for the fact that loans in the same geographical area are likely to suffer correlated (unobserved) shocks.

According to the estimates in the first column, private-label loans were approximately 0.6 percentage points more likely to receive concessionary modifications than loans held in portfolio. This estimate is economically small but statistically significant at the 10 percent level. When we consider all modifications the point estimate becomes one percentage point (statistically significant at the 10 percent level), while for the third specification, private-label loans were 2 percentage points more likely to receive concessionary modifications (statistically significant). As discussed above, all of these specifications include a number of additional loan characteristics that are important in the underwriting process and, thus, likely to play an important role in the modification decision. The first observation to make regarding the results reported in Panel B is that the difference between the incidence of modification for portfolio-held loans and private-label loans becomes even smaller when these variables are controlled for in the estimation. The results also imply that loans with higher credit scores were modified less, loans more likely to have a second mortgage were modified less, larger loans were modified more (although jumbo loans were modified less), and loans with more equity at the time of delinquency were modified less. We find a sizeable difference in terms of the frequency of modification for both refinances and subprime loans for the two specifications in the second and third columns. Conditional on being 60-days delinquent, subprime loans were modified about 2 percentage points less than prime loans (second column). We estimate a model separately for subprime loans in Table 2.8.

[25]Equity and periods of unemployment are very important determinants of a borrower's decision to default, and thus should also be important factors in the modification decision.

[26]The 12-month horizon implies that only mortgages that become delinquent before September 2007 are considered in the logit estimation.

Censoring is an important issue for any sample of mortgages, as there are currently many delinquent loans that are, or will soon be, good candidates for modification, as the housing market continues to decline. For this reason, we estimate a Cox proportional hazard model of the transition from serious delinquency to modification. The Cox model is very common in the survival analysis literature, and it has the advantage of being both flexible in terms of functional form considerations, as the baseline hazard function can be treated as an incidental parameter, and easy to estimate in terms of computational considerations. The results, expressed as hazard ratios, are reported in Panel C. A hazard ratio less than 1 indicates that private-label loans were less likely to receive a modification compared to portfolio loans, while a ratio greater than 1 signifies the opposite. The estimates are consistent with what we report for the logits in the previous panel. Private-label loans were less likely to receive concessionary modifications, but this coefficient estimate is statistically insignificant. For the other two modification definitions the sign flips, but again the result is not statistically significant. All three specifications include the same covariates that were included in the logit models.

2.4.2 Subsample Results

Table 2.8 contains further logit estimation results for various subsamples of interest to see if there are different probabilities than in the full sample. Since the subprime indicator seems to be an important predictor of modification conditional on serious delinquency in Table 2.7, we report the estimated marginal effects for only the sample of subprime loans in the second column of Table 2.8. The subprime sample also has the advantage that the agencies (Fannie Mae and Freddie Mac) were unlikely to be the marginal investor for this type of loan, so it is less likely that the portfolio and private-label samples differ significantly on unobservable characteristics. In the third column, we report results from the sample of LPS mortgages for which the borrower had a FICO score of less than 620, since automated underwriting systems generally instruct lenders to engage in increased scrutiny for such loans because of increased default risk. In the fourth and fifth columns, we focus on samples of loans that we believe contain the most information regarding the borrowers, in order to try to minimize the amount of unobservable heterogeneity that could potentially be biasing the results. In the fourth column, we focus on the sample of loans for which both the DTI ratio and the documentation status contain non-missing values, while the fifth column contains results for only the loans that were fully documented (in terms of income and assets) at origination. Panel A contains both unconditional means and estimated marginal effects for concessionary modifications, while Panel B contains results for the broader definition that also includes non-concessionary modifications.

The results are largely consistent with those contained in Table 2.7. We redisplay the results from the full sample in the first column of Table 2.8 for ease of comparison. The difference in modification frequency between private-label and portfolio-held, subprime mortgages for 60-day delinquent loans is small, and not statistically different from zero for both definitions of modification. Using a FICO cutoff of 620 as an alternative definition of subprime does not seem to make much difference. The unconditional means are similar (for both types of loans) compared to the LPS subprime sample, as are the marginal effects of private-label loans estimated from the logit models (although the result for all mods is statistically significant at the 10 percent level). Finally, we find small but statistically significant results at the 10 percent level for the last two subsamples (displayed in the fourth and fifth columns of Table 2.8), indicating that private label loans were *more* likely to be modified, not less.

2.4.3 Alternative Delinquency Definition

As an additional robustness check, we broaden our definition of delinquency and focus on modifications performed on loans subsequent to their first 30-day delinquency, which corresponds to one missed mortgage payment. While waiting until a borrower becomes seriously delinquent (defined as 60-days) to renegotiate is common practice in the servicing industry, there are no direct contractual stipulations (to our knowledge) that restrict a servicer from modifying the loan of a borrower who is 30-days delinquent. Thus, in Table 2.9 we repeat our analysis of Tables 2.7 and 2.8, but condition on 30-days delinquency rather than 60-days. The table contains three panels of estimation results, one for each of our modification definitions, and all of the subsamples considered in Table 2.8. The unconditional means, logit marginal effects, and Cox hazard ratios are all reported for each combination of subsample and modification definition.

The results are very similar to those from the analysis of 60-day delinquent loans. According to the full sample and subprime sample logit models, private-label loans received slightly more concessionary modifications, and the difference for the subprime subsample (0.5 percentage points) is statistically significant at the 10 percent level. Similarly, in the non-missing documentation and DTI sample and the full documentation sample Cox models, private-label loans received more concessionary modifications, although those differences are also small.[27] The results for our second modification definition are similar, and we find no significant differences between the incidence of portfolio and private-level modifications. The samples of portfolio loans with non-missing information for DTI and documentation status were modified at similar rates as the corresponding sample of private-label loans, with small differences both when we consider concessionary modifications and all modifications. Finally, in Panel C, we see strong evidence for both the logit and Cox specifications, that delinquent private-label loans prepayed more often than portfolio loans. The differences are statistically significant for every one of the subsamples.

2.4.4 Redefault Probabilities and Cure Rates

In the previous subsections, we showed that there is little difference in the frequency of mortgage loan modifications between servicers of loans held in a private trust versus loans held in portfolio. There are two potential reasons that may explain the failure of those exercises to pick up important differences in servicer behavior that may truly exist. First, it may be that contract frictions in securitization trusts do not result in substantial differences in the frequency of modifications (the extensive margin) but do result in significant differences in the intensive margin, with respect to the types of modifications performed, the extent to which contract terms are modified, and, more broadly, the care or effort expended in each

[27]The logit marginal effects correspond to percentage point differences, while the Cox hazard ratios correspond to percent differences. If one expresses the logit marginal effects as a percent change of the unconditional means, those percent changes are very similar in magnitude to the Cox results.

modification by private-label servicers compared to that expended by portfolio servicers. Second, there may be a type of renegotiation that our algorithm does not identify, but that is used to a large extent in loss mitigation efforts and used differently by servicers of private-label loans than by servicers of portfolio loans. For example, forms of forbearance, which are often called repayment plans in the industry, would not be picked up by our algorithm.[28] In this subsection, we use the LPS data to attempt to address these possibilities.

We perform two separate empirical exercises to address each of these concerns in turn. First, we compare redefault rates of private-label modified loans with those of portfolio modified loans. We define redefault as a loan that is 60 days delinquent or more, in foreclosure process or already foreclosed and now owned by the lender (REO for "real-estate-owned") six months after the time of the modification. If there are important differences in the manner by which servicers of private-label loans modify mortgages relative to the foreclosure procedures of servicers of portfolio loans, then we would expect to see significant differences in the subsequent performance of modified loans.

Second, to address the possibility that our algorithm misses an important aspect of renegotiation, we compare the cure rates of seriously delinquent, private-label loans to those of seriously delinquent portfolio loans. The idea behind this exercise is that any appreciable difference in servicer renegotiation behavior will manifest itself in differences in cure rates. As we mentioned above, Piskorski, Seru, and Vig (2009) focus on differences in foreclosure rates in an attempt to identify differences in the extent of renegotiation between servicers of portfolio loans and servicers of private-label loans. We view the cure rate as the more relevant outcome to study, as differences in foreclosure rates could be due to many other explanations, such as the unwillingness of banks to recognize losses or political pressure to reduce foreclosure numbers, perhaps by imposing foreclosure moratoria. In the current version of their paper, Piskorski, Seru, and Vig also estimate very similar cure regressions to those we report below, and after finding similar differences in cure rates compared to differences in foreclosure rates, interpret it as evidence that securitization may impede renegotiation. We find small differences (much smaller than the differences in foreclosure rates documented in Piskorski, Seru, and Vig (2009)), due in part to the different treatment of right-censored observations and to different control variables, which we discuss in detail in Appendix B and in Adelino, Gerardi, and Willen (2010). It is important to stress however, that differences in servicer renegotiation behavior are also only one potential explanation for differences that may exist in cure rates. To put this idea in the terms of logical reasoning, differences in cure rates are a necessary condition for significant differences in renegotiation behavior, but they are not a sufficient condition.

Table 2.10 contains the results of the redefault analysis. The first observation to note from the table is that the unconditional probability that a modified mortgage redefaults in this six-month period is very large, at about 35–45 percent for payment-reducing modifications (Panel A), and about 45–60 percent for all modifications (Panel B). We argue below that the high level of redefault rates could explain why we observe so few modifications — very often they do not lead to successful outcomes even as little as six months after the modification.

[28] However, as we argued above, PSAs do not contain restrictions on repayment plans, because they do not involve changing the terms of the mortgage. Thus, we would argue that differences in forbearance behavior that might exist could not be the result of contract frictions in securitization trusts.

The second observation to note is that there is no statistically significant difference between the redefault rates of private-label loans and those of portfolio loans, once the observable characteristics of the mortgages are taken into account (except for the subprime subsample, and when we consider all modifications, where private label is actually associated with fewer defaults). These results, combined with the statistics displayed in Table 2.6 suggest that there are no substantial differences in either the type of modification employed or in the care/effort expended by the two types of servicers.

Table 2.11 shows the results of both logit and cox proportional hazard models for the probability that a seriously delinquent loan subsequently cures. Accounting for right-censored observations is very important here, which is why we also include results from the hazard specification in Table 2.11 (we discuss this issue in detail in Appendix B). Our definition of a cure in the logit models is that the loan is either current, 30-days delinquent, or prepaid at any point during the first 12 months following the first 60-day delinquency. The first important point to make is that the unconditional cure probabilities in our sample are large (around 55 percent). Given that the unconditional modification probability is about 8 percent, this means that many loans cured without any intervention on the part of servicers. The second important observation to note in this table is that the cure probabilities for portfolio loans and private-label loans are quite similar. The unconditional cure probability is smaller by about 1 percentage point for private-label loans in the whole sample, but that difference reverses to 3 percentage points larger (statistically significant) when we control for observable characteristics of the loans and borrowers, which is approximately 5 percent of the unconditional mean for private-label loans. The estimated hazard ratio shows the opposite sign, as it implies that private-label loans are about 3 percent less likely to cure relative to portfolio loans, but is not statistically significant at conventional levels.

We also include results for the subsamples of interest in columns 2–5. For many of the subsamples private-label loans were *more* likely to cure. This is an important robustness check, as we argued above that unobserved heterogeneity is likely to be less of a problem in the subsamples (especially for the non-missing documentation status and DTI ratios sample and the full documentation sample). Thus, the change in the sign of the differences in cure rates between private-label servicers and portfolio servicers suggests that unobserved heterogeneity between the two loan types plays an important role.

In addition to the subsamples reported above, we also include results from three other subsamples that focus on borrowers with high credit scores at origination (FICO $>$ 680) in columns 6–8. Piskorski, Seru, and Vig (2010) focus much of their attention on the sample of loans for which the borrower's FICO score was above 680, based on their belief that borrowers with good credit histories are better candidates for renegotiation from the lender's perspective. We do not share the same sentiment, and in fact believe that this sample of borrowers likely contains a significant amount of unobserved heterogeneity. The reason is that the majority of mortgages in this subsample are below the GSE conforming loan limits, and thus are potential candidates for acquisition and securitization by Fannie Mae or Freddie Mac. We believe the fact that these mortgages were not sold to the GSEs implies that they are likely characterized by some risky element (unbeknownst to us) that makes them unattractive. Column 6 shows that the difference in cure rates between private-label and portfolio loans is only negative in the high FICO subsample (approximately 17 percent) compared to the full sample of loans. But, the results from the hazard estimates in columns

7 and 8 suggest that the difference between the high FICO sample and the full sample is largely being driven by the sample of conforming loans that we believe to be characterized by significant unobserved heterogeneity. The difference in cure rates for the non-conforming sample is consistent with the difference in cure rates in the full sample of mortgages (column 1).

2.4.5 Causality

We have been careful to interpret the estimates discussed above in terms of correlations rather than as causal effects. That is, we are not claiming that causality runs explicitly from securitization or portfolio lending to modification or cure probabilities, because we do not have any exogenous variation in the decision to securitize loans to obtain clean identification. We are simply saying that the fact that we do not see large observable differences in modification and cure probabilities does not support the story that the institution of securitization is impeding mortgage renegotiation. With that said, we believe that this study offers more than a purely descriptive analysis, because unobserved heterogeneity likely works against our finding of small differences in modification rates and cure rates. To the extent that privately securitized loans are worse on unobservables compared to portfolio-held loans, we would expect to see lower modification and cure rates for private-label loans, and thus a larger difference in renegotiation between the two types of loans. Of course if the opposite relationship held, then we'd expect lower modification rates of portfolio loans, which is why we are reluctant to interpret our estimates in a causal manner. But we would argue that the majority of the literature regarding this crisis suggests that securitized loans suffer from this issue more than portfolio-held loans.[29]

Piskorski, Seru, and Vig (2010) claim to have found a source of exogenous variation in the decision to securitize a mortgage, which they then use to study differences in foreclosure rates between private-label and portfolio loans. The source of the variation comes from rules in securitization contracts regarding mortgages that default early in their payment history, or so-called "early-payment defaults." In Adelino, Gerardi, and Willen (2010) we argue in detail that this is not a valid instrument for securitization as lenders choose which loans to repurchase and, in fact, repurchase very few. The identification thus hinges on the assumption that lenders randomly choose which loans to repurchase which is no more plausible than the assumption that lenders randomly select loans to select loans for securitization in the first place.

2.5 Understanding the Empirical Results

If securitization does not block renegotiation, then why is it so rare? In this section, we discuss the incentives behind the renegotiation decision from the lender's point of view, which, in a stylized way, mirrors the net present value (NPV) calculation that servicers are supposed to perform when deciding whether to offer a borrower a modification. We show that servicer uncertainty about the true ability or willingness of borrowers to repay their loans (and thus their ability to cure without needing a modification), as well as about

[29]See Krainer and Laderman (2009).

whether the borrower will redefault even after successful renegotiation can dramatically affect the NPV calculation, ruining what a naive observer might think of as a "win-win" deal for the borrower and lender. In addition, we argue that moral hazard may play an important role in the lender's modification decision. Specifically, as a lender offers a more generous modification to its eligible borrowers, it produces a financial incentive for ineligible borrowers to take hidden actions in order to gain eligibility. For example, many lenders require that borrowers are delinquent on their mortgage before qualifying for a modification. If the benefits of missing mortgage payments to qualify outweigh the costs, which include restricted access to future credit, then even borrowers who are not seriously considering default will have an incentive to become delinquent.

We also provide institutional evidence in this section that supports our arguments and findings above. This includes evidence of low modification frequencies in previous housing busts, well before the advent of securitization trusts; the equal treatment provision statements contained in the PSAs, which direct the servicer to behave as if it was in fact the investor of the mortgage-backed security and thus the owner of the mortgages; and finally, the absence of lawsuits to date directed at servicers by investors in mortgage-backed securities, which one would expect to find if modifications were unambiguously better than foreclosures.

2.5.1 Renegotiation and Asymmetric Information

In this section, we argue that avoiding the deadweight loss of foreclosure through renegotiation is far more difficult than it may at first appear. Many commentators have argued that the large losses in foreclosure relative to the costs of modification mean that the only possible explanations for foreclosures are either contract frictions from securitization or simple economic irrationality. Foreclosure can dominate renegotiation even when the costs of foreclosure far exceed the proposed concession by the lender, in particular when there is uncertainty about the effectiveness of renegotiation and when borrowers have private information about their ability and willingness to repay or about the value of their property.

There are two sets of problems. The first relate to uncertainty about the outcome of the loan with and without renegotiation and emerge even in situations where the borrower has no private information. Lenders take into account two important facts: (i) not all borrowers who renegotiate avoid foreclosure and (ii) some borrowers who fail to get modifications still manage to avoid foreclosure. Both of these "errors" make renegotiation less attractive, the former, known as redefault risk, because the renegotiation postpones foreclosure, delaying recovery and leading to further deterioration of the property and the latter, known as self-cure risk, because the lender will make a concession to a borrower who would have repaid the loan in full without assistance. It is important to understand that both redefault risk and self-cure risk emerge in situations where the lender and borrower have symmetric information. In other words, the borrower has no better idea whether he or she will cure or redefault than the lender.

Adding asymmetric information makes renegotiation much more difficult as shown in Riddiough and Wyatt (1994b) and Wang, Young, and Zhou (2001). Adelino, Gerardi, and Willen (2010) (henceforth AGW) consider a model in which borrowers have private information about their willingness and ability to repay the loan. Formally, AGW summarize a borrower's private information as a reservation amount that they are willing to repay; if the cost of repaying the loan exceeds the amount the borrower owes, default ensues and otherwise he or she repays. A lender that can observe the borrower's reservation value will set the repayment amount equal to it and thus avoid foreclosure and the associated deadweight losses. But a lender who cannot observe the borrower's valuation faces the familiar problem of a non-discriminating monopolist. In this case the marginal cost of renegotiation equals the liquidation value of the property but the profit-maximizing lender, like the monopolist, sets the price above marginal cost and thus allows foreclosures to proceed even in the absence of self-cure or redefault risk. The intuition is that by increasing the concession in renegotiation, the lender prevents additional foreclosures but incurs losses because it must extend that same concession to borrowers with higher reservation valuations, who would have accepted a much smaller concession.

It is crucial to understand that in this setup, the lender will walk away from positive NPV modifications. If we assume that there is no self cure risk and no redefault risk, then as long as the renegotiated repayment amount exceeds the liquidation value of the property, renegotiation has positive NPV and is, in that sense, profitable. But the goal of the lender is to maximize profits and not just to achieve non-negative profits, and the optimal concession to the borrower typically exceeds the minimum profitable concession.

In practice, asymmetric information has a pernicious effect that will be felt by borrowers who receive modification offers that are too small to allow them to keep their homes. Such borrowers will promise to default unless a larger modification is offered but the lender will rationally proceed with foreclosure. The costs savings on smaller modification offers accepted by observationally equivalent borrowers will offset the losses from the foreclosure.

But the problem is even worse than the previous paragraphs indicates because we assumed, up until now, that the distribution of reservation values is invariant to the size of the modifications offered. AGW show that the costs of modification grow even larger relative to foreclosure if we allow the distribution to change. In this case, the decision to offer more generous terms in renegotiation induces some borrowers who previously did not seek renegotiation to do so by, for example, deliberately becoming delinquent on their loans. This problem is well-known in the industry and is the moral hazard problem of modifications. As one lender described it in an article in *ABA Banking Journal*, "We have not to date forgiven any principal. We are wary of the consequences of being known as a bank that forgives principal."[30]

2.5.2 Institutional Evidence

While the results from Section 2.4 may be surprising to market commentators who believe that contract frictions inherent in securitization trusts are preventing large-scale modification efforts in mortgage markets, we argue in this section that both historical evidence and evidence from securitization contracts actually support our findings.

First, we look at history. If securitization, or more precisely private-label securitization, inhibits renegotiation, then we would expect that renegotiation would have been common in the 1990s, when there was little private-label securitization, or in the 1970s, when se-

[30] "Bankers' view of the new Hope for Homeowners program," *ABA Banking Journal*, October 2008.

curitization itself was rare. But, the historical evidence we have does not bear that out. In 1975, Touche Ross surveyed loss mitigation activities at savings and loans and found, "Lenders... were unwilling to either modify loans through extended terms or refinancing to a lower rate."[31] In the 1990s, a report commissioned by Congress to study foreclosure alternatives, said, "Along with loan modifications, long-term forbearance/repayment plans are the most under utilized foreclosure avoidance tools currently available to the industry."[32]

Second, many observers have focused on institutional factors that inhibit loan modification when the loan is securitized, but other factors may play a similar role for portfolio lenders as well. In particular, accounting rules force lenders (i) to take writedowns at the time of the modification (reducing Tier II capital), (ii) to identify modified loans as troubled debt restructurings (under FAS 15), and (iii) to impose burdensome reporting requirements on modified loans including loan-specific allowances for potential losses (under FAS 114). Additionally, payments made by borrowers for loans that are subject to "troubled debt restructurings" are recognized only as principal repayments and generate no interest income until the bank can demonstrate that a borrower is "performing." All of the above accounting requirements potentially make modifications costly for a bank. Downey Financial, for example, attempted to refinance current borrowers out of risky option ARMs into safer, fixed-rate instruments and argued that the change should not affect their balance sheet because the borrowers had never missed payments. However, their accountants viewed the refinancings as "troubled debt restructurings," and forced the firm to restate the share of nonperforming assets for November 2007 to 5.77 percent from 3.65 percent.[33]

If modifications were truly in the best financial interest of investors in mortgage-backed-securities (MBS) as many commentators have alleged, we would expect to see concern on their part regarding the low levels of modifications performed to date. But, according to Cordell, Dynan, Lehnert, Liang, and Mauskopf (2008b), who interviewed a number of MBS investors, they (the investors) are not concerned that servicers are foreclosing on many more mortgages than they are modifying. Thus, there does not seem to be much concern by market participants that either incentives or contract frictions are inhibiting servicers from performing loan modifications. The evidence in the literature seems to suggest a small role for contract frictions in the context of renegotiation. In a 2007 study of a small sample of PSAs, Credit Suisse found that fewer than 10 percent of the contracts ruled out modifications completely, while approximately 40 percent allowed modifications, but with quantity restrictions,[34] and the rest, about half, contained no restrictions on renegotiation behavior. Hunt (2009) also analyzed a sample of subprime PSAs and concluded that outright modification bans were extremely rare. A 2008 report by the COP analyzed a number of securitized mortgage pools with quantity restrictions and concluded that none of the restrictions were binding. In terms of incentive issues, Hunt (2009) found that most of the contracts in his sample explicitly instructed the mortgage servicer to behave as if it were the owner of the pool of the loans:

> The most common rules [in making modifications] are that the servicer must

[31] Capone (1996), p. 20–21.

[32] Capone (1996).

[33] http://www.housingwire.com/2008/01/14/downey-financial-accounting-rules-suck/

[34] The quantity restrictions often took the form of a limit (usually 5 percent) on the percentage of mortgages in the pool that could be modified without requesting permission from the trustee.

follow generally applicable servicing standards, service the loans in the interest of the certificate holders and/or the trust, and service the loans as it would service loans held for its own portfolio. Notably, these conditions taken together can be read as attempting to cause the loans to be serviced as if they had not been securitized. (p. 8, insertion added)

2.6 Conclusion

There is widespread concern that an inefficiently low number of mortgages have been modified during the current crisis, and that this has led to excessive foreclosure levels, leaving both families and investors worse off. We use a large dataset that accounts for approximately 60 percent of mortgages in the United States originated between 2005 and 2007, to shed light on the claim that delinquent loans have different probabilities of renegotiation depending on whether they are securitized by private institutions or held in a servicer's portfolio. By comparing the relative frequency of renegotiation between private-label and portfolio mortgages, we are able determine whether institutional frictions in the secondary mortgage market are inhibiting the modification process from taking place.

Our first finding is that renegotiation in mortgage markets during this period was rare. In our full sample of data, less than 2 percent of the seriously delinquent borrowers received a concessionary modification in the year following their first serious delinquency, while approximately 8 percent received some type of modification. These numbers are low, considering that foreclosure proceedings were initiated on approximately half of the loans in the sample and completed for almost 30 percent of the sample.

Our second finding is that a comparison of renegotiation rates for private-label loans and portfolio loans, while controlling for observable characteristics of loans and borrowers, yields economically small, and for the most part, statistically insignificant differences. This finding holds for a battery of robustness tests we consider, including various definitions of modification, numerous subsamples of the data, including subsamples for which we believe unobserved heterogeneity to be less of an issue, and consideration of potential differences along the intensive margin of renegotiation.

Since we conclude that contract frictions in securitization trusts are not a significant problem, we attempt to reconcile the conventional wisdom held by market commentators, that modifications are a win-win proposition from the standpoint of both borrowers and lenders, with the extraordinarily low levels of renegotiation that we find in the data. We argue that the data are not inconsistent with a situation in which, on average, lenders expect to recover more from foreclosure than from a modified loan. At face value, this assertion may seem implausible, since there are many estimates that suggest the average loss given foreclosure is much greater than the loss in value of a modified loan. However, we point out that renegotiation exposes lenders to two types of risks that are often overlooked by market observers and that can dramatically increase its cost. The first is "self-cure risk," which refers to the situation in which a lender renegotiates with a delinquent borrower who does not need assistance. This group of borrowers is non-trivial according to our data, as we find that approximately 50 percent of seriously delinquent borrowers "cure" in our data without receiving a modification. The second cost comes from borrowers who default again after receiving a loan modification. We refer to this group as "redefaulters," and our results show that a large fraction (between 40 and 60 percent) of borrowers who receive modifications, end up back in serious delinquency within six months. For this group, the lender has simply postponed foreclosure, and, if the housing market continues to decline, the lender will recover even less in foreclosure in the future. In addition to self-cure and redefault risk, we discuss how asymmetric information between the borrower and lender can significantly reduce the financial incentives for a lender to renegotiate.

We believe that our analysis has some important implications for policy. First, "safe harbor provisions," which are designed to shelter servicers from investor lawsuits, are unlikely to have a material impact on the number of modifications, and thus will not significantly decrease foreclosures. Second, and more generally, if the presence of self-cure risk, redefault risk, and moral hazard stemming from asymmetric information do make renegotiation less appealing to investors, the number of easily "preventable" foreclosures may be far smaller than many commentators believe.

2.7 Appendix: Identifying Modifications in the LPS Dataset

In this section we discuss in detail the assumptions that we used to identify modified loans in the LPS dataset. The LPS dataset is updated on a monthly basis, and the updated data include both new mortgages originated and a snapshot of the current terms and delinquency status of outstanding mortgages. Essentially, for a given mortgage, we compare the updated terms to the terms at origination, as well as the change in terms from the proceeding month, and if there is a material change over and above the changes stipulated in the mortgage contract, then we assume that the contract terms of the mortgage have been modified.

2.7.1 Interest Rate Reductions

We use a different set of rules to identify reduced interest rates for fixed-rate mortgages (FRM) and adjustable-rate mortgages (ARM). In principle, identifying a rate change for an FRM should be easy, since by definition the rate is fixed for the term of the mortgage. However, after a detailed inspection of the LPS data, it became apparent that some of the smaller rate fluctuations were likely due to measurement error rather than to an explicit modification. Thus, we adopt a slightly more complex criterion: The difference between the rate at origination and the current rate must be greater than 50 basis points; *and* the difference between the rate in the previous month and the current rate must be greater than 50 basis points; *and* either the mortgage must be 30-days delinquent with the loan currently in loss mitigation proceedings (as reported by the servicer) or the difference between the rate in the previous month and the current rate must be greater than 300 basis points (which allows for the possibility that a loan that is current could feasibly qualify for a modification).

Identifying interest rate reductions for ARMs is slightly more complicated, since by definition the interest rate is variable and can move both up and down. The LPS data contain the information necessary to figure out how much the interest rate should move from month to month. This rate is often referred to as the fully indexed rate, as it is normally specified as a fixed spread above a common nominal interest rate. The LPS dataset contains information regarding the initial rate, the appropriate index rate, and the spread between the index and the mortgage rate. In addition, the majority of ARMs are characterized by a period at the beginning of the contract in which the interest rate is held constant (these mortgages are often referred to as hybrid ARMs). At the end of this period, the interest rate adjusts (or resets) to a certain spread above an index rate and then subsequently adjusts at a specific frequency. The LPS dataset also contains information regarding the length of the initial fixed period, enabling us to identify this period in the data and determine the point at which the interest rate should begin to adjust (we refer to this period as the reset date). Our criterion for identifying an interest rate reduction for an ARM is as follows: The difference between the rate at origination and the current rate must be greater than 50 basis points; *and* the difference between the rate in the previous month and the current rate must be greater than 50 basis points; *and* if the reset date has passed, then the difference between the fully-indexed rate and the current rate must be at least 100 basis points ; *and* either the mortgage must be 30-days delinquent with the loan currently in loss mitigation proceedings (as reported by the servicer) or the difference between the rate in the previous month and the current rate must be greater than 300 basis points (which allows for the possibility that a loan that is current could feasibly qualify for a modification). In addition, we allow for more modest month-to-month decreases in the interest rate (200 to 300 basis points) as long as there is also a positive change in the delinquency status of the loan (that is, the loan is reported to be less delinquent). Our inspection of the data suggests that the majority of modifications involve a resetting of the delinquency status back to current, or a minor delinquency, so conditioning on this change likely eliminates many false positives.

2.7.2 Term Extensions

In theory, it should be straightforward to identify term extensions in the LPS data, but it can be tricky to do so because of possible measurement error in the variable that measures the remaining maturity of each loan. We defined a term extension in the LPS dataset to be a case in which the loan was at least 30-days delinquent at some point and the number of years remaining increases by at least 20 months *or* the change in number of years remaining is greater than the difference between the original term of the loan and the remaining term (for example, if the original maturity is 360 months, and the loan has 350 months remaining, then the increase in length must be at least 10 months) and, finally, either the monthly payment decreases *or* the principal balance increases *or* the loan is in loss mitigation.

2.7.3 Principal Balance Reductions

A reduction in the remaining balance of a mortgage is perhaps the most difficult type of modification to identify because of the prevalence of "curtailment" or partial prepayment among mortgage borrowers. For example, it is common for borrowers to submit extra mortgage payments in order to pay down the loan at a faster rate. For this reason, we were forced to adopt strict criteria to limit the number of false positives. Our criterion for identifying a principal balance reduction is as follows: The month-to-month decrease in the remaining principal balance must be at least -10 percent and cannot be more than -30 percent (the upper bound does not matter as much as the lower bound—we experimented with -40 percent and -50 percent, but did not find a substantial difference); the principal balance recorded in the previous month must be greater than $25,000 (since we throw second liens out, and look only at mortgages originated after 2004, this cutoff does not bind often); the month-to-month payment change must be negative (there are only a few cases in which the principal balance is reduced without a corresponding decrease in the payment, but in these cases the term is extended, and thus is picked up in our code for identifying term extensions); and, finally, the mortgage must be either 30-days delinquent or currently in loss mitigation proceedings (as reported by the servicer).

2.7.4 Principal Balance Increases

For interest-only and fully-amortizing mortgages, identifying an increase in the principal balance due to the addition of arrears is relatively straightforward. It becomes trickier for mortgages that allow for negative amortization, as the principal balance is allowed to increase over the course of the contract, by definition. For interest-only and fully-amortizing mortgages our criterion is: The month-to-month principal balance must increase by at least 0.5 percent (to rule out measurement error in the data); the loan must have been at least 30-days delinquent at the time of the balance increase; and, finally, the month-to-month payment change must be positive unless there is also a corresponding increase in the term of the loan. For mortgages that allow for negative amortization, the criterion is similar, except that the balance increase must be at least 1 percent and there must be a positive change in the delinquency status of the loan.

2.8 Appendix: Transferred Loans

There are loans in the LPS data for which the servicing rights are transferred to a servicer that does not contribute to the LPS dataset. When the servicing rights are transferred, the loans drop out of the data, and we do not observe their subsequent performance. While they only account for a small subsample of the loans in the data, unfortunately the transferred loans do not appear to be randomly selected. This is evident from a comparison of their observable characteristics at origination and their relative performance up until the time of transfer. In addition, a significantly higher percentage of portfolio-held loans are transferred compared to private-label loans. These observations imply that the sample of transferred loans cannot be simply dropped from the data, as doing so could result in a non-trivial bias of the estimated difference in cure rates (and foreclosure rates) between portfolio and private-label mortgages. In this appendix, we document some of the observable differences between transferred loans and non-transferred loans, show the bias introduced into the estimates of cure rate differences when transferred loans are dropped from the data, and discuss the appropriate way to deal with this sample of loans. Most of this discission is taken from a correspondence with Piskorsi, Seru, and Vig. In earlier versions of their paper, they in fact dropped the transferred loans, while in earlier versions of our paper (as well as in this draft), we kept them in the sample, but assumed that they did not cure. As we discuss below, dropping the transfers altogether creates a large bias, while assuming that the transfers do not cure does not seem to create bias, and in our opinion is a valid assumption for reasons discussed below.

In our sample, there are 3,922 loans that are transferred within 12 months of the date of the first 60-day delinquency. Approximately 4 percent of seriously delinquent, private-label loans transfer within 12 months, while 6.5 percent of seriously delinquent, portfolio loans transfer. Table 2.1 displays the the mean and median of the FICO score and DTI ratio at origination for the sample of transferred loans. The mean and median FICO score is lower for the entire sample of transferred loans compared to non-transferred loans, but the difference between portfolio transfers and non-transfers is much larger than the difference between private-label transfers and non-transfers. The same pattern emerges for DTI ratios.

In addition to the transferred loans having worse observable characteristics, they also perform significantly worse before they are transferred than loans that are not transferred.

Table 2.2 shows that transferred loans are significantly worse than non-transferred loans in the period immediately before they are transferred. The table compares the status of the transferred loans in the month before transfer to the status of all loans in the first twelve

Table 2.1: FICO and DTI Distribution of Transferred versus Non-transferred Loans

		Transfers			Non-Transfers		
		#	Mean	Median	#	Mean	Median
FICO	Private-label	1,276	628	627	68,442	632	632
	Portfolio	708	622	620	9,793	652	654
	Total	1,984	626	624	78,235	634	634
DTI	Private-label	1,636	42	41	40,821	40	41
	Portfolio	457	39	41	7,913	35	36
	Total	2,093	42	41	48,734	39	40

months after origination. A much lower percentage of transferred loans are current, and a much higher percentage are in foreclosure or seriously delinquent. This is especially true of the portfolio transfers.

Table 2.2: Status of Loan in Preceding Period: Transferred versus Non-transferred Loans

	Transfer Observations		All Observations	
Status in Previous Month	Private-label (%)	Portfolio (%)	Private-label (%)	Portfolio (%)
30-days delinquent	7.1	6.8	8.6	9.3
60-days delinquent	23.0	16.0	12.9	12.9
90-days delinquent	20.8	27.4	20.9	21.6
Current	9.5	7.5	14.0	23.1
In Foreclosure Proceedings	32.9	39.3	24.6	20.6
REO Status	6.6	2.8	16.6	10.1

As a result of the differences in observed characteristics and performance, the transferred loans cannot simply be thrown out of the sample. Doing so creates a serious bias in the comparison between the cure rates of portfolio loans versus private-label loans. The bias works in the direction of the portfolio loans, since the portfolio loans that are transferred are of worse quality than the private-label loans that are transferred. Thus, by throwing out all transferred loans, the cure rate of seriously delinquent portfolio mortgages is artificially inflated relative to the cure rate of private-label loans. This can actually be clearly seen from Panel A2 in the August 2009 version of Piskorski, Seru, and Vig. The table shows the results of logit cure regressions at different horizons. At each horizon the authors estimated two regressions, with each regression corresponding to a different sample of loans. The first sample excludes the transferred loans completely (this is the column titled "original sample"), while the second sample excludes loans that were transferred before the relevant horizon, but includes loans that were transferred after the horizon (this is the column titled "original + transferred"). While this does not solve the problem completely for horizons greater than 1 month (since transferred loans are still being excluded), the 1 month logits should be free of bias, as most of the transferred loans are still present. The estimated marginal effect falls from 0.045 to 0.026, when the transferred loans are kept in the sample which corresponds to a 46% decrease. This implies that holding the horizon constant (at 1 month), excluding transferred loans increases the estimated difference in cure rates between portfolio and private-label from just over 7% of the unconditional mean of the cure rate for portfolio loans to over 12%. In addition, at the twelve month horizon (the last two columns of the table) the bias can no longer be seen because almost all loans that are transferred in the LPS data are transferred within 12 months of the first serious delinquency.[35]

As a result of the bias discussed above from throwing out the sample of transferred loans, we conclude that the appropriate way to deal with the issue is to estimate a hazard model rather than a discrete-choice model like a logit or probit. A hazard model appropriately accounts for right-censored observations, and thus, has the advantage of being able to use observable information before the loans are transferred. Of course, it must be noted that the hazard model would not be able to control for scenarios in which the transferred loans perform worse after they are transferred (which we cannot observe). An implicit assumption in a hazard model is that right-censored observations perform the same as non-censored observations after they drop out of the data. Given that transferred loans perform worse while we observe them in the data, it is likely the case that they perform worse after we stop observing them. As a result, a hazard model specification for cure rates might still contain a bias in favor of portfolio loans. Piskorski, Seru, and Vig instead argue in an earlier draft of their paper (August 2009) that the sample of transferred portfolio loans may have been transferred to other servicers of portfolio loans who renegotiate many of their loans, in which case it would be more difficult to show that portfolio loans cure more than private-label loans. They say:

> Second, there are at least two large servicers (Ocwen Loan Servicing and Litton Loan Servicing) who are widely known to renegotiate substantial amount of loans that are not in the LPS database. Consequently, not including a sample of largely bank-held loans that might be transferred for renegotiation to these servicers might actually make it harder for us to demonstrate that portfolio loans are renegotiated more intensively.

However, the servicers that they mention, Ocwen and Litton (who are not providers to the LPS data), service private-label loans almost exclusively. Thus, if the private-label loans are being transferred to Ocwen and Litton, then the estimates of the hazard model will actually be biased in favor of finding larger cure rates for portfolio loans, rather than smaller cure rates!

As mentioned above, in earlier versions of this paper, as well as in the current draft, we estimate logit models and treat the transfers as non-cures (as opposed to throwing them out of the data). We made this assumption on the grounds that the servicer that we observe in the data did not modify these loans while they were in the data. While some may argue that this assumption could have the effect of biasing the estimates of the cure differences in the opposite direction – that is, creating a downward bias in the probability of cure for

[35] In fact the bias is present in both columns.

portfolio loans – the estimates from the hazard specification in Table 2.11 are very similar to the estimates from the logit models, implying little bias.

Bibliography

Adelino, M., K. Gerardi, and P. Willen (2010). Mortgage Renegotiation and Asymmetric Information. manuscript.

Ambrose, B., R. Buttimer Jr., and C. Capone (1997). Pricing mortgage default and foreclosure delay. *Journal of Money, Credit & Banking 29*(3).

Ambrose, B. and C. Capone (1996). Cost-benefit analysis of single-family foreclosure alternatives. *The Journal of Real Estate Finance and Economics 13*(2), 105–120.

Ambrose, B. and C. Capone (1998). Modeling the conditional probability of foreclosure in the context of single-family mortgage default resolutions. *Real Estate Economics 26*(3), 391–429.

Ambrose, B. and C. Capone (2000). The hazard rates of first and second defaults. *The Journal of Real Estate Finance and Economics 20*(3), 275–293.

Asquith, P., R. Gertner, and D. Scharfstein (1994). Anatomy of financial distress: An examination of junk-bond issuers. *Quarterly Journal of Economics 109*(3), 625–658.

Bolton, P. and D. S. Scharfstein (1996). Optimal debt structure and the number of creditors. *Journal of Political Economy 104*(1), 1–25.

Capone, C. A. (1996). Providing alternatives to mortgage foreclosure: A report to Congress. US Department of Housing and Urban Development.

Congressional Oversight Panel (2009). Foreclosure crisis: Working toward a solution. March Oversight Report.

Cordell, L., K. Dynan, A. Lehnert, N. Liang, and E. Mauskopf (2008a). The incentives of mortgage servicers: myths and realities. *Federal Reserve Board Finance and Economics Discussion Series* 2008-46.

Cordell, L., K. Dynan, A. Lehnert, N. Liang, and E. Mauskopf (2008b). The incentives of mortgage servicers: Myths and realities. *Finance and Economics Discussion Series, Federal Reserve Board 46*.

Cutts, A. and R. Green (2005). Innovative servicing technology: smart enough to keep people in their houses? *Building Assets, Building Credit: Creating Wealth In Low-Income Communities*, 348–377.

Cutts, A. and W. Merrill (2008). Interventions in mortgage default: Policies and practices to prevent home loss and lower costs. Technical report, Working Paper.

Danis, M. A. and A. N. Pennington-Cross (2005). The delinquency of subprime mortgages. *SSRN eLibrary*.

Foote, C., K. Gerardi, L. Goette, and P. Willen (2009, June). Reducing foreclosures: No easy answers. Working Paper 15063, National Bureau of Economic Research.

Foote, C., K. Gerardi, and P. Willen (2008). Negative equity and foreclosure: Theory and evidence. *Journal of Urban Economics 6*(2), 234–245.

Gan, Y. H. and C. Mayer (2006). Agency conflicts, asset substitution, and securitization. *NBER Working Paper No. 12359.*

Geanakoplos, J. and S. Koniak (2008). The best way to help homeowners — and the economy. Available at http://cowles.econ.yale.edu/ gean/crisis/index.htm.

Gelpern, A. and A. Levitin (2009). Rewriting Frankenstein contracts: Workout prohibitions in mortgage-backed securities. http://ssrn.com/abstract =1323546.

Hunt, J. P. (2009). What do subprime securitization contracts actually say about loan modification? *Berkeley Center for Law, Business and the Economy working paper.*

Krainer, J. and E. Laderman (2009). Mortgage loan securitization and relative loan performance. *Federal Reserve Bank of San Francisco Working Paper 22.*

Lauria, M., V. Baxter, and B. Bordelon (2004). An investigation of the time between mortgage default and foreclosure. *Housing Studies 19*(4), 581–600.

Levitin, A. (2009a). Helping homeowners: Modification of mortgages in bankruptcy. *Harvard Law & Policy Review Online 3.*

Levitin, A. (2009b). Resolving the foreclosure crisis: Modification of mortgages in bankruptcy. *Wisconsin Law Review*.

Mason, J. (2007). Mortgage loan modification: Promises and pitfalls.

Mulherin, J. and W. Muller (1987). Volatile interest rates and the divergence of incentives in mortgage contracts. *Journal of Law, Economics, and Organization 3*(1), 99–115.

Pennington-Cross, A. (2009). The duration of foreclosures in the subprime mortgage market: a competing risks model with mixing. *The Journal of Real Estate Finance and Economics*, 1–21.

Pennington-Cross, A. and G. Ho (2006). Loan servicer heterogeneity and the termination of subprime mortgages. *Federal Reserve Bank of St. Louis Working Paper 2006-024A.*

Piskorski, T., A. Seru, and V. Vig (2009). Securitization and distressed loan renegotiation: Evidence from the subprime mortgage crisis. *Chicago Booth School of Business Research Paper No. 09-02.*

Quercia, R., L. Ding, and J. Ratcliffe (2009, March). Loan modifications and redefault risk. Working paper, Center For Community Capital: The University of North Carolina at Chapel Hill.

Riddiough, T. and S. Wyatt (1994a). Strategic default, workout, and commercial mortgage valuation. *The Journal of Real Estate Finance and Economics 9*(1), 5–22.

Riddiough, T. and S. Wyatt (1994b). Wimp or tough guy: Sequential default risk and signaling with mortgages. *The Journal of Real Estate Finance and Economics 9*(3), 299–321.

Springer, T. and N. Waller (1993). Lender forbearance: evidence from mortgage delinquency patterns. *Real Estate Economics 21*(1), 27–46.

Stegman, M., R. Quercia, J. Ratcliffe, L. Ding, W. Davis, W. Li, K. Ernst, A. Aurand, and S. Van Zandt (2007). Preventive servicing is good for business and affordable homeownership policy. *Housing Policy Debate 18*(2).

Wang, K., L. Young, and Y. Zhou (2001). Nondiscriminating Foreclosure and Voluntary Liquidating Costs. *The Review of Financial Studies 15*(3), 959–985.

White, A. M. (2008). Rewriting contracts, wholesale: Data on voluntary mortgage modifications from 2007 and 2008 Remittance Reports. *Fordham Urban Law Journal forthcoming.*

White, A. M. (2009). Deleveraging the American homeowner: The failure of 2008 voluntary mortgage contract modifications. *Connecticut Law Review forthcoming.*

Zingales, L. (2008). Plan b. *Economists' Voice 5*(6), 4.

Table 2.3: Examples of modifications in the data.

Example 1: Servicer cuts interest rate, capitalizes arrears in the balance of the loan and extends term to 40 years.

Date	MBA Delinq. Stat.	Interest Rate	Monthly Payment	Outstanding Balance	Remaining Term in Months
2008m4	9	6.5	907	141,323	340
2008m5	9	6.5	907	141,323	339
2008m6	9	6.5	907	141,323	338
2008m7	C	4.5	660	146,686	479

Example 2: Servicer capitalizes arrears into the balance of the loan but otherwise leaves the loan unchanged.

Date	MBA Delinq. Stat.	Interest Rate	Monthly Payment	Outstanding Balance	Remaining Term in Months
2008m5	6	9.25	1,726	208,192	346
2008m6	9	9.25	1,726	208,192	346
2008m7	9	9.25	1,726	208,192	346
2008m8	C	9.25	1,815	218,316	341
2008m9	C	9.25	1,815	218,184	340

Table 2.4: Robustness of the modifications algorithm

False positives by type of modifications

	# of Modifications Using WF CTS Data	False Positives
FRM Rate Reduction	5,381	8.0%
ARM Rate Reduction	8,951	22.0%
Principal Reductions	470	1.9%
Principal Increases	13,010	12.8%
Term Increases	394	2.3%

Overall success of algorithm

	No Mod Using Our Algorithm	Mod Using Our Algorithm	Total
No Mod in WF Data	2,329,187	3,559	2,332,746
Mod in WF Data	3,627	17,514	21,141
Total	2,332,814	21,073	2,353,887

Notes: We test our algorithm on a dataset of securitized mortgages in which the trustee has identified modifications (data is from Wells Fargo Trustee Services). The lower panel shows that about 17.2% of our modifications are false positives, meaning that we identify modifications but the trustee does not and about 16.9% are false negatives, meaning that the trustee identifies a modification but we do not.

Table 2.5: Modification Statistics

(1) By Type of Modification: 2007:Q1–2008:Q4

	# Loans Modified	Interest Rate Reductions		Principal Balance Reductions		Principal Balance Increases		Term Extensions	
		#	(% total)	#	(% total)	#	(% total)	#	(% total)
2007:Q1	10,940	600	5.3	700	6.2	8,660	76.4	1,380	12.2
2007:Q2	14,600	820	5.4	550	3.7	11,630	77.3	2,050	13.6
2007:Q3	17,720	770	4.1	810	4.3	15,170	81.2	1,940	10.4
2007:Q4	27,150	2,990	9.7	700	2.3	22,520	72.8	4,740	15.3
2008:Q1	36,230	6,010	13.8	900	2.1	32,100	73.8	4,500	10.3
2008:Q2	44,750	9,050	16.4	1,300	2.4	39,750	72.1	5,030	9.1
2008:Q3	62,190	16,280	20.3	940	1.2	56,940	70.9	6,110	7.6

(2) By Payment Change

	Payment Decreases					Payment Increases				
	#	mean Δ		median Δ		#	mean Δ		median Δ	
		$	%	$	%		$	%	$	%
2007:Q1	2,080	-492	-13.2	-157	-10.0	5,020	106	6.7	62	4.4
2007:Q2	2,060	-464	-12.7	-141	-9.6	7,710	120	7.0	63	4.4
2007:Q3	2,470	-290	-12.9	-125	-9.7	10,380	110	6.7	60	4.3
2007:Q4	5,600	-367	-15.3	-159	-11.7	14,540	100	5.9	59	3.9
2008:Q1	11,500	-358	-14.0	-210	-13.2	18,720	108	6.5	62	4.3
2008:Q2	18,660	-425	-16.1	-239	-14.1	20,770	124	7.4	69	4.1
2008:Q3	31,770	-562	-21.5	-365	-20.2	26,400	124	6.3	63	3.6

(3) Loan Characteristics of Modified Mortgages

	All Loans					Modifications				
	#	mean	p25	p50	p75	#	mean	p25	p50	p75
FICO (at origination)	1,892,777	706	660	713	762	17,533	622	580	621	662
LTV (at origination)	2,250,162	75	67	79	85	21,675	82	78	80	90
DTI (at origination)	1,346,093	37	28	38	45	13,945	41	35	41	47
Mortgage balance (at origination)	2,267,497	231K	121K	185K	288K	21K	234K	121K	186K	294K
% characterized as										
LTV = 80	14.4					21.7				
Subprime	6.8					47.4				
Fixed	71.2					39.7				
Hybrid ARM	7.7					26.2				
IO-ARM	11.3					13.1				
IO-Fixed	2.1					2.7				
Option-ARM	5.1					12.0				
Option-Fixed	0.3					1.4				
Owner	89.3					96.0				
Investor	7.1					2.6				
Vacation Home	3.7					1.1				
Purchase	51.9					49.0				
Low/no documentation	29.2					20.4				

Notes: These statistics were computed using a 10% random sample of the LPS data. Quantities obtained from the data are multiplied by a factor of 10. The percentages in panels (1) and (2) are taken with respect to the total number of modifications, and *not* loans modified. Thus, there is double-counting in the sense that some loans received multiple types of modifications in a given quarter.

Table 2.6: Modification Comparison by Payment Change

Private-label Modifications

	Payment Decreases					Payment Increases				
	#	mean		median		#	mean		median	
		$	%	$	%		$	%	$	%
2007:Q1	106	-614	-14.42	-162	-10.85	239	121	6.02	76	3.37
2007:Q2	110	-505	-12.02	-222	-9.30	364	168	7.96	76	3.49
2007:Q3	128	-261	-11.82	-131	-8.42	558	145	7.52	75	3.65
2007:Q4	288	-313	-13.38	-163	-12.36	741	125	6.24	74	3.52
2008:Q1	634	-393	-16.12	-261	-15.65	938	133	6.76	79	4.08
2008:Q2	1,014	-540	-18.94	-334	-17.89	1,241	152	8.14	83	4.08
2008:Q3	1,778	-641	-22.01	-423	-19.95	1,805	137	6.22	70	3.31

Portfolio Modifications

	Payment Decreases					Payment Increases				
	#	mean		median		#	mean		median	
		$	%	$	%		$	%	$	%
2007:Q1	28	-759	-20.90	-428	-17.19	128	106	7.78	52	5.46
2007:Q2	19	-1172	-25.17	-656	-28.07	222	81	6.11	55	5.28
2007:Q3	31	-395	-17.13	-168	-15.29	255	71	6.13	43	5.37
2007:Q4	90	-474	-11.11	-90	-2.48	292	70	5.50	37	4.29
2008:Q1	187	-369	-10.00	-183	-8.08	331	80	6.59	33	3.97
2008:Q2	309	-304	-10.90	-117	-6.64	405	63	5.59	34	3.56
2008:Q3	376	-585	-25.19	-295	-17.85	359	105	7.04	39	4.26

Table 2.7: Modifications (Main Sample)

Panel A: Unconditional Percentages

	Concessionary Mods	All Mods	All Mods + Prepayments
Portfolio	0.010	0.086	0.154
Private-label	0.018	0.087	0.159

Panel B: Logit Regressions (12 month horizon)

	Concessionary Mods	All Mods	All Mods + Prepayments
Private-label	0.006	0.010	0.021
	1.88	1.97	3.24
Initial Rate	0.002	-0.001	-0.004
	3.00	-1.14	-2.62
LTV Ratio	0.000	-0.001	-0.002
	2.44	-2.02	-2.86
LTV = 80	-0.001	-0.012	-0.035
	-0.79	-3.14	-7.78
FICO	0.000	-0.001	-0.003
	-0.88	-1.02	-4.00
$FICO^2$	0.000	0.000	0.000
	0.78	0.96	3.80
FICO $<$ 620	0.009	0.058	0.060
	1.27	3.76	3.56
$620 \leq$ FICO $<$ 680	0.005	0.022	0.024
	1.11	2.34	2.15
Log Original Amount	0.004	0.014	0.029
	2.33	4.00	6.41
Equity at Delinquency	0.000	-0.032	0.001
	-0.22	-2.58	0.12
Negative Equity	-0.004	-0.020	-0.034
	-0.58	-1.14	-1.32
Unemployment	0.000	-0.004	-0.007
	0.35	-3.64	-4.69
Refi	0.002	0.007	0.034
	0.95	1.89	6.22
Subprime	0.001	-0.018	-0.017
	0.23	-3.86	-2.70
Other Controls	Y	Y	Y
# Mortgages	28,574	28,574	28,574

Panel C: Duration Model

	Concessionary Mods	All Mods	All Mods + Prepayments
Private-label	0.91	1.00	1.02
	1.43	0.10	0.68
# Mortgages	77,676	77,676	77,676

Notes: Other controls include indicator variables for Jumbo, Option, Hybrid and Interest-Only mortgages, as well as for condos and multifamily homes. Panel B shows the marginal effects of logit regressions with a 12-month horizon, t-statistics shown below the coefficients. Standard errors are clustered at the zip code level. Panel C shows hazard ratio estimates from a Cox proportional hazards model.

Table 2.8: Modifications (Robustness tests with alternative samples)

Panel A: Concessionary Modifications

	All Loans	Subprime	$FICO < 620$	Non-missing Documentation and DTI	Full Documentation
Portfolio Mean	0.010	0.014	0.012	0.008	0.008
Private-label Mean	0.018	0.020	0.022	0.017	0.019
Private Label Mg. Eff. (Logit)	0.006	0.004	0.004	0.007	0.009
	1.88	0.87	0.76	1.81	1.74
Private Label Haz. Ratio (Cox)	0.91	0.86	0.89	1.19	1.07
	1.43	1.56	1.05	1.99	0.63
# Mortgages	28,574	18,099	13,101	15,206	10,900

Panel B: All Modifications

	All Loans	Subprime	$FICO < 620$	Non-missing Documentation and DTI	Full Documentation
Portfolio Mean	0.086	0.069	0.080	0.086	0.076
Private-label Mean	0.087	0.092	0.104	0.093	0.104
Private Label Mg. Eff. (Logit)	0.010	0.010	0.017	0.014	0.027
	1.97	1.12	1.74	1.89	2.53
Private Label Haz. Ratio (Cox)	1.00	0.96	1.04	1.17	1.13
	0.10	0.69	0.55	3.28	1.94
# Mortgages	28,574	18,242	13,223	15,317	10,981

Panel C: All Mods + Prepayment

	All Loans	Subprime	$FICO < 620$	Non-missing Documentation and DTI	Fully Documented
Portfolio Mean	0.154	0.145	0.150	0.152	0.141
Private-label Mean	0.159	0.164	0.189	0.164	0.179
Private Label Mg. Eff. (Logit)	0.021	0.026	0.032	0.025	0.039
	3.24	2.47	2.61	2.68	3.13
Private Label Haz. Ratio (Cox)	1.02	1.01	1.06	1.13	1.12
	0.68	0.24	1.18	2.96	2.15
# Mortgages	28,574	18,242	13,223	15,317	10,981

Notes: Portfolio and private-label means are unconditional probabilities of modification in each sample. Marginal effects are computed from logit models with a 12-month horizon that include all the controls in Table 2.7. Standard errors are clustered at the zip code level. t-statistics are reported below the marginal effects.

Table 2.9: Modifications Conditional on 30 Days Delinquency (Logits)

Panel A: Concessionary Mods

	All Loans	Subprime	*FICO* < 620	Non-missing Documentation and DTI	Fully Documented
Portfolio Mean	0.007	0.005	0.005	0.005	0.005
Private-label Mean	0.010	0.011	0.012	0.010	0.011
Private Label Mg. Eff. (Logit)	0.001	0.005	0.005	0.003	0.004
	0.55	1.77	1.68	1.65	1.71
Private Label Haz. Ratio (Cox)	0.988	0.941	0.949	1.128	1.034
	0.39	1.16	0.94	2.86	0.62
# Mortgages	64,857	31,402	24,004	33,341	23,752

Panel B: All Mods

	All Loans	Subprime	*FICO* < 620	Non-missing Documentation and DTI	Fully Documented
Portfolio Mean	0.036	0.037	0.042	0.036	0.031
Private-label Mean	0.036	0.043	0.047	0.037	0.041
Private Label Mg. Eff. (Logit)	-0.002	-0.001	-0.001	-0.004	0.000
	-0.98	-0.20	-0.28	-1.60	-0.05
Private Label Haz. Ratio (Cox)	0.94	0.91	0.91	1.22	1.10
	1.04	1.08	0.94	2.75	1.01
# Mortgages	65,256	31,513	24,197	33,590	23,969

Panel C: All Mods + Prepayment

	All Loans	Subprime	*FICO* < 620	Non-missing Documentation and DTI	Fully Documented
Portfolio Mean	0.158	0.188	0.158	0.156	0.133
Private-label Mean	0.184	0.196	0.213	0.183	0.191
Private Label Mg. Eff. (Logit)	0.026	0.014	0.027	0.021	0.033
	6.28	1.80	3.08	3.65	4.29
Private Label Haz. Ratio (Cox)	1.12	1.00	1.09	1.16	1.20
	6.32	0.12	2.51	6.22	5.73
# Mortgages	65,256	31,513	24,197	33,590	23,969

Notes: Portfolio and private-label means are unconditional probabilities of modification in each sample. Marginal effects are computed from logit models with a 12-month horizon that include all the controls in Table 2.7. Hazard ratios are computed from Cox proportional hazard models with the same controls as in Table 2.7. z-statistics are shown below the coefficients, and t-statistics are reported below the marginal effects. Standard errors are clustered at the zip code level. Sample sizes refer to the logit regressions. The sample sizes for the Cox models are slightly larger.

Table 2.10: Redefault Conditional on Modification

Panel A: Concessionary Mods

	All Loans	Subprime	$FICO < 620$	Non-missing Documentation and DTI	Fully Documented
Portfolio Mean	0.348	0.479	0.392	0.299	0.338
Private-label Mean	0.447	0.473	0.465	0.457	0.458
Private Label Mg. Eff. (Logit)	0.037	-0.037	-0.019	0.030	-0.001
	1.43	-0.84	-0.37	0.82	-0.01
# Mortgages	2,970	1,848	1,338	1,633	1,239

Panel B: All Mods

	All Loans	Subprime	$FICO < 620$	Non-missing Documentation and DTI	Fully Documented
Portfolio Mean	0.455	0.661	0.520	0.461	0.467
Private-label Mean	0.553	0.591	0.590	0.573	0.577
Private Label Mg. Eff. (Logit)	0.011	-0.066	-0.041	0.011	0.001
	0.71	-2.34	-1.52	0.50	0.02
# Mortgages	8,561	5,209	4,626	4,855	3,809

Notes: redefault is defined as loans that are 60 days delinquent, 90 days delinquent, in the process of foreclosure or in REO 6 months after the modification. Marginal Effects refer to the marginal effects of a logit model with a horizon of 6 months. t-statistics shown below the marginal effects. Standard errors are clustered at the zip code level.

Table 2.11: Cure Conditional on 60 Days Delinquency

	All Loans	Subprime	$FICO < 620$	Non-missing Documentation and DTI	Fully Documented	$FICO > 680$		
						Full Sample	Conforming	Non-conforming
Portfolio Mean	0.558	0.503	0.586	0.553	0.575	0.534	0.549	0.493
Private-label Mean	0.544	0.561	0.634	0.567	0.614	0.436	0.449	0.409
Private Label Mg. Eff. (Logit)	0.028	0.068	0.079	0.047	0.062	-0.034	-0.030	-0.046
	3.40	5.53	6.00	4.16	4.39	2.20	1.56	1.71
Private Label Haz. Ratio (Cox)	0.97	1.08	1.05	1.00	1.04	0.93	0.91	0.99
	1.53	2.71	1.59	0.20	1.61	2.29	2.42	0.27
# Mortgages	28,574	18,242	13,223	15,317	10,981	5,702	3,842	1,845

Notes: The dependent variable ("Cure") is defined as a loan that is either current, 30 days delinquent, or prepaid 12 months after the first 60-day delinquency. Portfolio and Private-label means are unconditional probabilities of modification in each sample. Marginal effects are computed from logit models with a 12-month horizon that include all the controls in Table 2.7. Standard errors are clustered at the zip code level. t-statistics are reported below the marginal effects.

Chapter 3

Reservation Prices in Online English Auctions: A Field Experiment

3.1 Introduction

Auctions are widely used to sell a variety of goods, from art and antiques in auction houses to electronics and web advertising in online auctions. One important question economists have focused on is how best to design auctions in order to maximize revenues for the auctioneer, including how to set the optimal minimum bid. Riley and Samuelson (1981) answer this question in the setting of independently distributed private values and show that the optimal reservation price exceeds the auctioneer's true valuation and is independent of the number of bidders. Levin and Smith (1996) show that in a setting of correlated values the optimal reservation price converges to the seller's true value. In reality, however, sellers often either set no reservation prices or set reservation prices below what theory would predict (McAfee, Quan and Vincent, 2002).

There has been significant interest in evaluating empirically the effect of reservation prices on the outcomes of auctions. While there are many studies done in a laboratory setting (for an overview of the existing experimental studies see Kagel, 1995 and Ockenfels, Reiley and Sadrieh, 2006), there is significantly less experimental work done using field data. The advent of online auctions has made the study of this topic easier, by providing both lower cost environments to obtain data and to set up field experiments. Three papers in particular run field experiments to test the effect of introducing and varying reservation prices in online auctions. Reiley (2006) runs first-price, sealed bid auctions of collectible trading cards and finds that introducing (and increasing) reservation prices reduce the probability of sale and increase the revenue conditional on sale. Ostrovsky and Schwarz (2009) study a large field experiment to assess the effect of altering reservation prices in Internet advertising auctions in line with theory of optimal auctions and find that this improves revenues for the auctioneering firm. Brown and Morgan (2009) conduct auctions of collectible coins on Yahoo! and eBay and find that setting positive reservation prices (as opposed to no reservation prices) increases seller revenues.

In this paper we analyze the outcomes of field experiments using second-hand motorcycles sold by a large commercial bank in India in online ascending price auctions. The experiments involved two main interventions: changing the reservation price used by the bank (mostly reducing the initial minimum by 10 percent, although we also performed a few increases in initial reservation price) and bundling vehicles together. We ran four auctions in three separate locations (Chennai, Hyderabad and Mumbai) with a total of 637 motorcycles put up for sale and a total of 269 sold items.

We find that reductions in the reservation price used by the bank lead to a higher probability of sale - a one percent decrease in the initial price leads to a 1.5-2 percentage points increase in the probability of sale (for a baseline probability of sale of about 42 percent). At the same time, a one percent drop in the reservation price reduces the final sale price for the vehicles by about 0.4-0.8 percentage points (for a baseline appreciation of 6 percent). On balance, the bank would be better off reducing the minimum bid in terms of total revenue if it were not able to put the vehicle up for sale in later auctions. By reducing the reservation price on the motorcycles by around 10 percent, the increase in probability of sale on a given auction is around 15-20 percentage points and the reduction in final sale amount is approximately 4-8 percentage points. In reality, however, the bank cannot commit to not attempting to re-sell the vehicles (and in fact always puts unsold vehicles up for sale in later auctions), so the final verdict on total revenues depends on the storage costs of each vehicle, the time it takes to sell vehicles in later auctions and the final price for each vehicle. The trade-off faced by the auctioneer when discounting the motorcycles is between 0.4-0.8 percent of the value of the vehicle for each percentage point reduction in the minimum bid and a 1.5-2 percentage point increase in the probability of having to bear the cost of waiting.

One issue that has been discussed in the literature that we can directly address with our experimental design is that of "auction fever" induced by a reduction in the minimum bid. This is the first paper to test the effect of changing minimum bids on auction fever in a field experiment and in the context of ascending-bid online auctions (as opposed to sealed bid auctions, discussed in Reiley, 2006). This phenomenon has been documented in other settings where bidders have become more aggressive (and placed higher bids) due to additional competition during an auction. Some examples include the results of Hubl and Popkowski-Leszczyc (2004), Ku, Malhotra and Murnighan (2005), and Ariely and Simonson (2003) who find auctions produce higher sale prices and higher revenues when bidders are in the presence of a larger number of rival bids and active bidders.

We find no evidence that reducing minimum bids induces auction fever on the part of participants. The vehicles that were discounted sold, on average, below the reservation price that had been set initially by the bank and we find just a small increase in the number of bidders for discounted items. Even when we exclude discounted vehicles that sold below the initial minimum bid there is no significant difference between the sale price of discounted and non-discounted vehicles. This contradicts the hypothesis that reducing the minimum bid would induce more competition and, consequently, higher final sale prices (as suggested by Heyman, Orhun and Ariely, 2004). We hypothesize that the reason for this discrepancy relative to earlier studies is (i) the fact that there is no significant increase in the number of bidders for discounted versus non-discounted items, a condition that is necessary for auction fever to occur and (ii) the fact that the bidders in these auctions are specialized and experienced buyers of used motorcycles in the locations where the auctions are held. It is important to note, however, that we have nothing to say about other potential forms of inducing competition such as directly altering the number of bidders or the sunk costs of each bidder as used by Ku, Malhotra, and Murnighan (2005).

The fact that the final sale price for the discounted vehicles is lower relative to the control group is consistent with the bidders inferring information about the quality of the vehicles from the initial minimum bid, which they know is also the reservation price for the bank. The inspection done by each bidder of each motorcycle before the auction is far from exhaustive, so participants may infer that the bank has information about the vehicles that is reflected in the minimum bid. It is less clear, however, why the information conveyed by the minimum bid alone would lead to a higher probability of sale of lower minimum bid vehicles (one possible explanation would be the existence of specialized dealers in lower quality vehicles, for example).

The joint result of reduced sale prices along with a higher probability of sale is most consistent with mispricing on the part of the auctioning bank, along with the presence of few potential bidders for each vehicle (i.e. some bidder specialization). If the minimum bid set by the bank is just above the highest bidder valuation, then the motorcycle will go unsold. Reducing the minimum bid may bring in these potential bidders. At the same time, if there are only a few potential bidders for each type of motorcycle, then as long as the reduction in the minimum bid still keeps the reservation price between the highest and the second highest bidder valuation, the final sale price will be lower with the discounted minimum bid (Riley and Samuelson, 1981, and Levin and Smith, 1996).

Regarding the bundles created in the first three auctions, selling two vehicles together reduces the probability of sale by anywhere between 5 and 22 percentage points and reduces the final sale price by about another 6 to 8 percentage points. The bundles we created had a specific structure, namely each bundle included what the auctioning bank termed one "fast-moving" and one "slow-moving" motorcycle model. These labels reflect the collections team's perception about the probability of sale of the "slow-moving" vehicles, namely that it was low given the minimum bid they set. The opposite was true for the "fast-moving" vehicles, namely the collections team thought these vehicles had a high probability of sale given their initial minimum bids. By bundling together both types of models, one might be able to reduce the average mispricing and thereby increase the net revenue. In fact, bundling increased the probability of sale for "slow-moving" vehicles, but the net effect of bundling on total revenue is negative. We cannot determine whether the underlying reason for the failure of bundling is due to preferences for types of vehicles or due to credit constraints on the part of bidders.

This paper contributes to the existing literature in a few different ways. First, unlike most of the existing experimental evidence (both in laboratory settings and in the field), the items for sale in the auctions we analyze (motorcycles) are large ticket items, with a mean value of 25 thousand rupees (approximately 550 dollars or one-fifth of the per capita GDP in India). Second, and also unlike most of the laboratory studies, the bidders in our setting are highly specialized and have extensive experience participating in this type of auctions. Third, we use the existing minimum bid set by the auctioneer as the benchmark rather than a zero reservation price. The minimum bids set by the bank officers are guided by both past auction experience and prevailing market prices, so it is reasonable to assume they already entail a significant amount of optimization. We look at changes (mostly reductions) to this initial reservation price, which directly allows us to test for "auction fever", something that has not yet been tested in this fashion in the literature. Fourth, the auctions we analyze are closer to common-value auctions than most laboratory experiments that have been analyzed in the literature, where in most cases the researchers try to create a context closer to independent private values auctions.

The remainder of the paper is organized as follows: Section 2 reviews the existing literature and theoretical predictions, Section 3 gives institutional background and delineates the experimental setup, Section 4 describes the results and Section 5 concludes.

3.2 Theoretical literature and empirical predictions

The main exercise in this paper involves the variation of minimum bids (or reservation prices) upwards and downwards in online affiliated-value ascending price auctions. The items for sale are second-hand motorcycles held in three cities in India and the seller is a large commercial bank. There are two main outcomes of the auctions that we will analyze: First, we consider bidder participation in the auctions and consequently the probability of sale of each item and, second, we analyze the revenue for the seller conditional on a sale.

The existing literature on the effect of minimum bids has considered two main sets of assumptions about the valuations of the participating bidders - independent private values and affiliated values (also called common values when the true value is the same for all bidders). The auctions run in this experiment are closest in spirit to affiliated-value auctions. Unlike the independent private values (IPV) setting, a higher valuation of one bidder is likely to be correlated with higher valuations of the other participating bidders, but the objects being auctioned do not have identical value to all bidders.

In the benchmark model with independent private values considered by Milgrom (2004), when the number of bidders is fixed, the seller setting a higher reservation price entails a trade-off: increasing the revenue conditional on a sale and, simultaneously, reducing the probability of a sale. The optimum reservation price in this setting is higher than the true value for the seller and this result is independent of the number of bidders in the auction (Riley and Samuelson, 1981, and Myerson, 1981). In the common value setting, revenue should also increase when a reservation price or minimum bid is set (Milgrom and Weber, 1982), but the optimal reservation price converges to the true reservation value to the seller as the number of participating bidders increases (Levin and Smith, 1996). The reason behind the difference between the result in the IPV setting and the common value setting is that the highest bidder valuation is not correlated with the second highest valuation under the IPV assumptions, whereas it is in the common value setting.

Note that, under both IPV and common value assumptions, positive reservation prices (or minimum bids) only affect the revenue conditional on sale when there are few bidders in an auction. In fact, only when the reservation price is set below the highest bidder valuation and above the second-highest bidder valuation does it influence the ultimate sale price of an item. As more bidders enter an auction, the gap between the highest and the second highest valuation becomes smaller in expectation and the reservation price becomes less important both for the revenue conditional on sale and the probability of sale (although it affects the probability of sale more in the case of common value auctions due to the correlation between bidder valuations, Levin and Smith, 1996).

The predictions from the theory regarding bidder participation are thus straightforward:

setting higher seller reservation prices (in the form of minimum bids) should reduce bidder participation, and, consequently, increase the probability that an item goes unsold. Conversely, reducing the minimum bid should increase participation and reduce the probability of not selling a vehicle. We will analyze the probability of sale in all four auctions as a function of the change in base price. We also look at the probability of sale separately for vehicles where the bank indicated that the ex ante likelihood of sale was high and low at the prices initially set (we refer to these two groups as "fast-moving" and "slow-moving" vehicles, respectively).

Within the vehicles that are sold, we compare the number of bidders that entered the auction for discounted versus undiscounted items whenever that information is available. Changing the reservation price should also change the number of bidders that participate in an auction when entry is endogenous and costly. In our experiment entry is not very costly, because most bidders are already registered on the website and regularly visit the vehicle warehouses, but the decision to attend an online auction is still likely to be a function of the reservation prices.

In terms of the price appreciation of each vehicle conditional on sale, there are three main alternative hypotheses, two of which derive directly from above: First, with few potential bidders, the predictions in Riley and Samuelson (1981) and Levin and Smith (1996) point to lower final sale prices with discounted minimum bids. The fact that new bidders enter the auction would reinforce this effect, as bidders in affiliated-value auctions should take the information learned from other bids into account and (in order to avoid the winner's curse) bid less than the expected value of the object conditional on each bidder's individual valuation or signal (Milgrom and Weber, 1982). This is especially true in our setting given the experienced and specialized nature of the participants in the vehicle auctions that we analyze, as well as the "real world" setting of the auctions (Kagel and Levin, 1986; Dyer and Kagel, 1996).[1] The second hypothesis holds when there are many potential bidders. In that case, the above theories would suggest that the minimum bid should not matter for the final sale price conditional on a sale occurring, leading to essentially no change in final sale prices for discounted reservation prices.

There are mechanisms other than those in Riley and Samuelson (1981) and Levin and Smith (1996) that lead to a positive correlation between minimum bids and the final sale price. First, if the buyers infer information from the seller's reservation price, the price appreciation should be positively correlated with the change in the minimum bid. Second, if buyers use the minimum bids as reference points (Hubl and Popkowski-Leszczyc, 2003 and Rosenkranz and Schmitz, 2007) lower minimum bids should lead to lower final sale prices. Note, however, that these two mechanisms don't have clear predictions about the probability of sale (as the bidders' valuations are a function of the reservation price itself) which makes them harder to evaluate empirically.

The third alternative hypothesis relies on the fact that bidding in these auctions is sequential and visible to all participants.[2] The fact that bidders in these motorcycle auctions

[1]This is in contrast to the winner's curse shown in the work of Kagel, Levin, Battalio, and Meyer (1989) and Lind and Plott (1991) where the experiments are run in laboratories with inexperienced subjects.

[2]The auctions in this experiment mimic English auctions in an online setting, so they are unlike the second-price auctions on eBay where bidders are asked to submit a bid equal to their reservation price and the intermediate bids are done automatically.

observe the bids of other participants and also the number of active bidders for each vehicle produces a feedback effect that allows learning about other participants' reservation prices. This could induce what the literature has called "auction fever", which means that subjects perceive the behavior of other bidders to be competitive and thus increase their subjective valuation of the auctioned object (Heyman, Orhun, and Ariely, 2004). If this phenomenon is strong enough, reducing minimum bids would lead to higher final sale prices. For this prediction to hold, we need the number of bidders to increase significantly for discounted items. At the same time, the fact that the bidders in the auctions we consider are fairly knowledgeable of the market value of the auctioned items makes the "auction fever" hypothesis less likely.

Given that we have randomly changed the minimum bids initially set by the bank, we will judge the impact of the changes of the minimum bid on final sale prices by looking at appreciation relative to the minimum bid initially set by the bank. This should be equivalent to looking at sale prices, except that the randomization procedure does not produce exactly equal average minimum bids suggested by the bank. By looking at appreciation relative to initial base prices we should eliminate some of the heterogeneity and obtain a better estimate of the impact on revenue conditional on sale.

3.3 Institutional Background

The partner for the field experiments we discuss in this paper is a large commercial bank in India. We worked with the collections department responsible for repossessing vehicles from borrowers who became seriously delinquent on car and motorcycle loans. This department was also responsible for managing the inventory of repossessed vehicles and for liquidating that collateral. The bank sold 17.6 thousand vehicles in 2005 and over 25 thousand in 2006, of which approximately 80 percent were motorcycles and the rest were cars and commercial vehicles (the weight of motorcycles in total revenue was approximately 30 percent). These vehicles were sold either in online or in physical auctions conducted in warehouses and storage sites across India. We don't have data on the exact split between online and physical auctions, but by 2006 more than 50 percent of the sales were done through online auction websites.

The bidders in our experiment are all either owners or employees of second-hand motorcycle dealerships in the cities the auctions were held in. The bidding is closed to individuals not affiliated with a professional dealership primarily because many of the vehicles put for sale need some type of refurbishing and making the bidding public might subject the bank to some type of litigation risk. Most of these vehicles are thus bought with the intention of re-sale after a short period of time. While it is the case that some dealers might have customers already lined-up for some motorcycles (and may have agreed on a sale price that could induce a "quasi-private" valuation for the item), the value of each of these motorcycles is strongly influenced by the existing second-hand market.

The re-possessed vehicles held by banks are an important source of inventory for the dealerships that participate in these auctions. Most participants' main activity consists of buying, refurbishing and re-selling these vehicles, so the size of the purchases and the potential profit they can make on each motorcycle is central to their business. As we pointed out before, the fact that the motorcycles put for sale are large ticket items bought by specialized and professional dealers is one of the elements that sets our study apart the most from the previous literature. Most of the existing experimental literature analyzes auctions of small-value items where bidders were generally inexperienced auction participants (without a clear perception of the market value of the objects).

The bidding in this bank's online auctions starts from a minimum bid set by the auctioneer and bids are visible to everyone. The vehicles are sold to the highest bidder so the minimum bid is in effect the bank's reservation price (no secret reservation price is set). This is different from what has been found in other online auctions, such as in Bajari and Hortasu (2003) for eBay coin auctions, where bidders in general set very low minimum bids and used secret reservation prices for high value items. The bank also never bids in the auctions, i.e. there is no shill biding.

The auctions that we consider in this paper have some additional characteristics that we abstract from. One such characteristic is that from discussions with some auction participants it became apparent that bidders sometimes collude in the auctions. In fact, owners of neighboring motorcycle shops often sit together when bidding on motorcycles that they are interested in. While this can have an effect on the dynamics of the auctions, we do not have information of the level of collusion that is going on in each auction, so we assume it just ads noise to the estimate of the true effect of changing reservation prices. Another characteristic that we abstract from is that there are instances in which bidders do not pay for (and consequently don't pick up) the items won in the auction. When a vehicle is sold in an auction, the winning bidder has 48 hours to pay the winning amount and collect the item from the bank's warehouse. Failing to do so usually leads the bank to re-auction the previously sold items. While the bank tried to punish bidders for these situations, as of the time of these experiments these mechanisms were ineffective at preventing these "fake" sales.

The information available on the website to bidders includes the registration number for the vehicle, whether or not the bank has been able to obtain the registration documents from the borrower, the model and make of each motorcycle, the year of manufacture and one picture of the vehicle. Prior to the auction, dealers typically visit the warehouse where the vehicles are stored. These visits are intended for dealers and their clients, who first obtain an authorization from the collections department to visit the warehouses. Most of this information (except obviously for the information collected during the warehouse visits) is available to us in the dataset.

3.3.1 Experimental Setup

The experiments involved randomly dividing the motorcycles in each auction into two or three groups - one groups in which the vehicles had the minimum bid reduced by 10 percent, another in which the reservation prices were left unchanged and (for the Mumbai auction only) a third group where we increased the minimum bid by 10 percent. We obtained the initial minimum bids from the collections team and these initial values were set as if no changes would take place (in particular, the collections team was not aware of which minimum bids would be manipulated). The collections team also provided us with a list of models that historically had a lower probability of sale, which they called "slow-moving" vehicles. We divided the "slow-moving" and "fast-moving" vehicles evenly into the two or three bins defined above, so the same proportion of both types of vehicles received a discount.

We also constructed bundles made up to two motorcycles each - one randomly selected "slow-moving" model and one randomly selected "fast-moving" model. The intention of constructing bundles was to test (a) whether there was, indeed, mispricing of "slow-moving" vehicles that would be minimized by bundling together "slow" and "fast" moving vehicles or, alternatively, (b) whether there were cash constraints or preferences that would lead dealers to prefer individually sold vehicles to the bundles. There were 17 bundles constructed in the 10/10/2006 auction, 28 bundles made in the 11/27/2006 auction and 48 bundles in the 1/27/2007 auction. Importantly, bundled items were also included in the randomization procedure for reducing the minimum bids.

3.4 Empirical results

3.4.1 Summary Statistics

The data used in this paper comes exclusively from the commercial bank in India mentioned in Section 3. The online auctions we analyze are all location- and type of vehicle-specific, which means that each auction only included motorcycles from one of the locations where the bank operates. We ran experiments in two online auctions held in Hyderabad, one held in Chennai and one held in Mumbai.

The mean minimum bids initially set by the bank are shown in Table 3.1. We split the vehicles into the motorcycles sold in a bundle and those sold individually. The counts for the discounted and unchanged groups show that treatment and control sample are close to evenly divided (the Chennai collections team had some reservations about the experiment and so a smaller number of vehicles were discounted in that auction). The mean minimum bids are significantly different between different auctions, with Hyderabad's minimum bids for individually sold items on average at 28 thousand rupees, Chennai at nineteen thousand and Mumbai at around 27 thousand rupees (all regressions in the analysis section include auction fixed effects to take into account these differences). The differences in the initial minimum bids set by the bank for the two groups ("unchanged" and "discounted") are within half a standard deviation of each other. The standard deviation of prices for individually sold vehicles is approximately eight thousand rupees.

In Table 3.2 we provide the test for the randomization procedure. If the randomization achieved its purpose, there should be no significant differences between the initial minimum bids set by the bank for the treatment and control samples (note again that the discount to the minimum bids was applied to these initial prices set by the bank). We show the results separately for each auction and then for all auctions pooled together. There are no statistically significant differences between the "discounted" and "unchanged" subsamples, so we can consider the two samples comparable.

Table 3.3 contains all other relevant summary statistics for the auctions in which we held reservation price experiments. We again separately show the summary statistics for bundles and motorcycles sold individually. The mean revised reservation price of the discounted items is, by construction, 10 percent below the mean reservation price for the items in the "unchanged" group. The one exception is the Mumbai auction, where we rounded the updated minimum bids to the closest 500 rupees with the intention of "hiding" the discounted items. In the other auctions it would be possible to identify the discounted items by looking at the reservation price, so bidders might infer some information from atypical minimum bids such as 23.4 thousand rupees, which would correspond to 90 percent of an initial reservation price of 26 thousand.

The appreciation for each vehicle is defined as the difference between the sale amount and the minimum bid (as a percentage of minimum bid) and is calculated based on the initial minimum bid set by the bank rather than on the revised minimum bids in order to make the figures comparable between the "discounted" and "unchanged" groups. The first fact about appreciation that is worth noting is that for all auctions the increase in price for "unchanged" items is larger than for "discounted" ones. This is consistent with the results found in the regressions, namely that discounted items suffered a reduction in the price conditional on sale relative to vehicles that used the initial minimum bid. The second interesting fact is that the mean difference between the reservation prices and the actual sale amounts varies widely between auctions. This reflects both cross sectional variation in demand (i.e. demand for motorcycles varies between different cities) as well as time series variation. We take this into account in the regression analysis by using auction fixed effects. Finally, the appreciation amount is slightly higher for the items for which we increased the minimum bid (70 motorcycles in the Mumbai auction).

Consistent with the results mentioned in the introduction and discussed below, the percentage of vehicles sold is significantly different for "discounted" and "unchanged" vehicles. For all auctions except Chennai, the discounted items had a higher probability of sale, in some cases of up to 100 percent higher (as in the Mumbai auction). Interestingly, for the subset of vehicles for which we increased the initial minimum bid there is almost no change in the fraction of vehicles that were sold.

The number of bidders participating in the bidding does not vary visibly for "discounted" and "unchanged" vehicles.[3] The data we have available only includes the top three bids for each vehicle, buy conditional on that data restriction we find no visible differences in the number of bids for the two subsamples. As we pointed out before, that could be one reason why we find no evidence of the "auction fever" phenomenon. We should note that this reflects some degree of bidder concentration, given the total number of bidders and the number of bids we see in each auction (Table 3.4). In fact, we have a total of 23 different bidders in the first Hyderabad auction and 15 in the second, only 6 different bidders in the whole Chennai auction and, finally, 13 winning bidders in the Mumbai auction. The concentration is quite different between different auctions. For example, the four largest winning bidders in the first Hyderabad auction buy between 7 and 8 motorcycles each, whereas the two top bidders in the second Hyderabad auction buy a combined total of 69 motorcycles.

3.4.2 Regression results

The first outcome of the auctions that we consider in the regression analysis is the probability of sale for each vehicle. The results are presented in Table 3.5 and show that discounted items have a higher probability of sale on the order of 1.5 to 1.8 percentage points for every

[3] This data is only available for the first three auctions – we were unable to obtain it for the Mumbai auction.

percentage point discount applied to the initial minimum bid. In the first four columns of Table 3.5 we included the results with an indicator variable for whether a vehicle was discounted or not. In columns (5) through (8) we repeat the same regressions with the amount each vehicle was discounted by included linearly. Given that most vehicles that received a discount were discounted either by 10 percent or by an amount very close to that, the results are very similar between the two specifications. The main controls we include are auction fixed effects to account for the differences between locations and specific dates, as discussed in the previous section. As expected given the summary statistics in Table 3.3, the two auctions in Hyderabad show the highest probability of sale, with Chennai and Mumbai having much lower conversion rates from auction to actual sale. We also estimate all regressions with an indicator variable for bundles. Bundles are less likely to be sold, although the point estimate is not statistically significant when we compare bundles to all other vehicles. The difference becomes large and statistically significant when bundles are compared to "fast-moving" vehicles (specifications 5, 6, 7 and 8). We also use fixed effects for "slow-moving" vehicles and an interaction of the "slow-moving" dummy with the discount variable. We confirm that the models termed by the bank to be "slow-moving" do have a significantly lower probability of sale of about 20-30 percent when compared to "fast-moving" vehicles.

Next we turn to price appreciation, or the difference between the sale price and the minimum bid as a percentage of the initial minimum bid. The regression results are presented in Table 3.6 and Table 3.7. The difference between the two tables is that the first uses an indicator variable for discounted vehicles as the independent variable of interest, whereas the second table (Table 3.7) uses the amount by which the vehicles were discounted as the independent variable. Overall we find significantly lower price increases relative to the minimum bid initially set by the bank for discounted when compared to non-discounted items. Discounted vehicles sold at an average 4-8 percentage points lower sale price when they were discounted by 10 percent. Note that the appreciation of each vehicle is measured relative to the minimum bid initially set by the bank, in order to make it comparable between treatment and control samples. When we exclude the discounted vehicles that sold below the initial minimum bid, we find no statistically significant difference in the sale price, which is more evidence against the "auction fever" hypothesis. We find no statistically significant difference between the price appreciation of "slow-moving" and "fast-moving" vehicles, although discounts seem to have a statistically significantly larger effect on price appreciation for "slow-moving" than for "fast-moving" vehicles. Regarding the bundles, we find those are subject to a lower price appreciation of about 5-8 percentage points.

We also consider the probability that a motorcycle is sold exactly at the minimum bid set by the auctioneer. This is one measure of the influence of the minimum bid on the ultimate auction outcome and may be an indication that the minimum bid was set between the highest and the second highest valuation. The results are shown in Table 3.8 and are consistent with what we obtained for the probability of sale and price appreciation. In particular, we find that discounted items are less likely to be sold exactly at the minimum bid set (11-16 percentage points less likely, not always statistically significant). This is consistent with higher reservation prices being more likely to be set between the highest and second highest reservation prices and thus inducing a higher price conditional on sale. Bundles are associated with an 18-39 percentage points higher likelihood of an item being sold exactly at the minimum bid set by the auctioneer, consistent with there being few available bidders for bundled items (due to specialization or maybe financing constraints). "Slow-moving" models are also much more likely to be sold at exactly the minimum bid set by the bank.

In order to obtain some more detail about bidder participation (other than just the probability of sale), we consider the number of bidders that participate in the auctions. We show the results in Table 3.9. We find only a small increase of around 0.2 bidders per sold vehicle on average for discounted versus non-discounted items. This difference is statistically significant only in specification (3) where we take into account the effect of "slow-moving" vehicles and then only at the 10 percent level. Bundles are associated with a drop of 0.3-0.8 bidders on average for sold items. We should note that we only have the top three bids for each vehicle, not the full bidding history for these auctions, which may bias the results against finding an increase in bidder participation.

We also consider the distribution of bids for the three subsamples in the experiment (discounted, unchanged and increased minimum bids). We follow Reiley (2006) and tabulate the number of bids for separate bins that indicate the distance from the original minimum bids set by the bank. Consistent with Tables 3.5-3.9, we find that discounted items are much more likely to sell below the original minimum bid and we find a similar relationship to Reiley (2006), namely that higher reservation prices induce higher bids, in our case in the setting of online English auctions rather than online sealed-bid auctions (results in Table 3.10). This holds up to 125 percent of the original minimum bid and then the relationship becomes less stable, although this could be due to very few observations above that threshold. The pattern of many bids within 10 percent of the reservation price set by the auctioneer (clearly visible for discounted and unchanged items) is consistent with bidders either inferring information from the reservation price or using the reservation price as a reference point as suggested by Rosenkranz and Schmitz (2007).

3.5 Conclusion

In this paper we analyze the impact of reservation prices on the probability of sale and the revenue conditional on sale in the setting of online motorcycle auctions. We run field experiments using vehicles from the repossessed inventory of a large commercial bank in India. We find that lower reservation prices (relative to the original minimum bids set by the auctioning bank) increase the probability of sale but reduce the price conditional on sale of these motorcycles. This is consistent with standard auction theory but contradicts the hypothesis that one can induce "auction fever" by using lower minimum bids in online ascending-price online auctions. We hypothesize that the fact that the bidders are specialized motorcycle dealers and experienced participants in online auctions similar to the ones we ran makes it harder to arouse competitive behavior than when facing less experienced subjects.

Most of our experiments involved lowering reservation prices by 10 percent or less. One question left for further research is whether larger discounts would produce similar results or whether, alternatively, one could induce more participation and, consequently, higher prices in line with the "auction fever" hypothesis.

Bibliography

[1] Ariely, D. and Simonson, I. (2003) "Buying, Bidding, Playing, or Competing?: Value Assessment and Decision Dynamics in Online Auctions." Journal of Consumer Psychology, 13(1-2), 113-123.

[2] Bajari, P. and Hortasu, A. (2003) "Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions." The RAND Journal of Economics 2. pp. 329-355.

[3] Brown, J. and Morgan, J. (2009) "How Much Is a Dollar Worth? Tipping versus Equilibrium Coexistence on Competing Online Auction Sites.", Journal of Political Economy, vol. 117, no. 4, pp. 668-700.

[4] Dyer, D. and Kagel, J. (1996) "Bidding in common value auctions: How the commercial construction industry corrects for the winner's curse." Management Science, Oct 1996, 42, 1463-1475.

[5] Hubl, G. and Popkowski-Leszczyc, P. (2003) "Minimum Prices and Product Valuations in Auctions." Marketing Science Institute Report 2003 (03-117), 115-41.

[6] Hubl, G. and Popkowski-Leszczyc, P. (2004) "Bidding Frenzy: Intensity of Competitive Interaction Among Bidders and Product Valuation in Auctions." Advances in Consumer Research, Volume 31, pp. 90-93.

[7] Heyman, J. E., Orhun, Y. and Ariely, D. (2004) "Auction Fever: The Effect of Opponents and Quasi-Endowment on Product Valuations. Journal of Interactive Marketing 18(4).

[8] Kagel, J. and D. Levin (1986) "The Winner's Curse and Public Information in Common Value Auctions." The American Economic Review, Vol. 76, No. 5.

[9] Kagel, J., D. Levin, R. Battalio, and D. Meyer (1989) "First-Price Common Value Auctions: Bidder Behavior and the Winner's Curse" Vol. 27, No. 2.

[10] Kagel, J.H. (1995) "Auctions: A Survey of Experimental Research." Handbook of Experimental Economics, pp. 501-585.

[11] Ku, G., Malhotra, D. and Murnighan, J. K. (2005) "Towards a Competitive Arousal Model of Decision-Making: A Study of Auction Fever in Live and Internet Auctions." Organizational Behavior and Human Decision Processes 96. 89-103.

[12] Levin, D. and Smith, J. L. (1996) "Optimal Reservation Prices in Auctions." The Economic Journal, Vol. 106, No. 438, pp. 1271-1283

[13] Lind, B., and C. Plott (1991) "The Winner's Curse: Experiments with Buyers and with Sellers". The American Economic Review, Vol. 81, No. 1.

[14] McAfee, R. P., Quan, D. C. and Vincent, D. R. (2002) "How to Set Minimum Acceptable Bids, with an Application to Real Estate Auctions." The Journal of Industrial Economics, Vol. 50, No. 4, pp. 391-416

[15] Milgrom, P. (2004) "Putting Auction Theory to Work." Cambridge University Press.

[16] Milgrom, P. and Weber, R. (1982) "A General Theory of Auctions and Bidding." Econometrica, 50, pp. 1089-1122.

[17] Myerson, R. B. (1981), "Optimal Auction Design," Mathematics of Operations Research,6(1), pp. 58-73.

[18] Ockenfels, A., Reiley, D., and Sadrieh, A. (2006) "Online Auctions." Prepared for the Handbook on Economics and Information Systems.

[19] Ostrovsky, M. and Schwarz, M. (2009) "Reserve Prices in Internet Advertising Auctions: A Field Experiment." Working Paper.

[20] Reiley, D. (2006) "Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet." The RAND Journal of Economics, Vol. 37, No. 1, pp 195-211.

[21] Riley, J. G. and Samuelson, W. (1981) "Optimal Auctions." American Economic Review, Vol. 71, pp. 381-392.

[22] Rosenkranz, S. and Schmitz, P. W. (2007) "Reserve Prices in Auctions as Reference Points."The Economic Journal, No. 117, 637-653

Table 3.1: Summary Statistics of Initial Reservation Price

	Sold Individually				Bundle		
	Increase	Unchanged	Discount	Total	Unchanged	Discount	Total
Hyderabad		25,179	26,716	25,899	55,643	57,393	56,518
10/10/2006		(8,579)	(6,805)	(7,789)	(8,848)	(11,087)	(9,883)
		42	37	79	14	14	28
Hyderabad		24,519	26,687	25,603	53,992	56,194	55,093
11/27/2006		(8,725)	(8,305)	(8,551)	(11,300)	(8,824)	(10,091)
		59	59	118	24	24	48
Chennai		19,105	16,000	18,185	37,556	42,250	39,765
1/27/2007		(6,172)	(7,651)	(6,724)	(10,818)	(10,082)	(10,432)
		38	16	54	9	8	17
Mumbai	27,871	28,434	26,880	27,818			
8/14/2007	(6,170)	(6,092)	(6,775)	(6,308)			
	70	76	54	200			
Total	27,871	25,075	25,726	25,749	51,336	54,134	52,720
	(6,170)	(8,047)	(8,036)	(7,825)	(12,361)	(11,036)	(11,745)
	70	215	166	451	47	46	93

Note: Columns indicate the change to the initial reservation price. First four columns refer to vehicles sold individually and the last three to bundled vehicles. First line for each date has the mean reservation price before it was manipulated. Standard deviation in parenthesis and number of observations below the standard deviation.

Table 3.2: Verification of Randomization Procedure

	Hyderabad 10/10/2006	Hyderabad 11/27/2006	Chennai 1/27/2007	Mumbai 8/14/2007	All
Increase				-563	281
				(1,045)	(1,210)
Discount	1,593	2,178	-973	-1,555	477
	(1,622)	(1,394)	(1,978)	(1,123)	(735)
Bundle FE	YES	YES	YES	YES	YES
Auction FE					YES
Number of obs.	107	166	71	200	544

Note: The dependent variable is the initial reservation price. The independent variables are indicator variables for whether the morotcycle was discounted or received a price increase. All regressions include one indicator variable for bundles. The last regression includes all four auctions with auction fixed effects as independent variables.

Table 3.3: Summary Statistics of Outcomes

		Sold Individually				Bundles		
		Increase	Unchanged	Discount	Total	Unchanged	Discount	Total
Hyderabad 10/10/2006	Revised Res. Price		25,179	24,045	24,647	55,643	51,654	53,648
	Sale Price		27,798	27,070	27,389	54,250	52,606	53,325
	Appreciation		12.0%	2.0%	6.4%	4.6%	-2.1%	0.8%
	Percent Sold		50.0%	73.0%	60.8%	50.0%	64.3%	57.1%
	No. Bidders		2.3	2.4	2.4	2.1	2.8	2.5
Hyderabad 11/27/2006	Revised Res. Price		24,519	24,019	24,269	53,992	50,574	52,283
	Sale Price		24,444	26,369	25,570	55,528	52,368	53,604
	Appreciation		10.2%	0.9%	4.8%	0.3%	-9.2%	-5.5%
	Percent Sold		45.8%	64.4%	55.1%	37.5%	58.3%	47.9%
	No. Bidders		2.1	2.2	2.2	1.2	1.6	1.5
Chennai 1/27/2007	Revised Res. Price		19,105	14,400	17,711	37,556	38,025	37,776
	Sale Price		21,231	22,850	21,681	28,750	33,300	31,025
	Appreciation		2.9%	0.5%	2.2%	7.0%	-10.0%	-1.5%
	Percent Sold		34.2%	31.3%	33.3%	22.2%	25.0%	23.5%
	No. Bidders		1.3	1.4	1.3	1.5	1.0	1.3
Mumbai 8/14/2007	Revised Res. Price	29,329	28,434	25,630	27,990			27,990
	Sale Price	31,781	29,000	28,011	29,438			29,438
	Appreciation	13.8%	11.9%	8.7%	11.1%			11.1%
	Percent Sold	22.9%	18.4%	40.7%	26.0%			26.0%
	No. Bidders	NA	NA	NA	NA			NA
Total	Revised Res. Price	29,329	25,075	23,621	25,200	51,336	48,720	50,042
	Sale Price	31,781	25,677	26,776	26,763	52,056	50,928	51,400
	Appreciation	13.8%	9.7%	3.1%	6.7%	2.7%	-6.7%	-2.8%
	Percent Sold	22.9%	34.9%	55.4%	40.6%	38.3%	54.3%	46.2%
	No. Bidders		2.0	2.3	2.1	1.6	2.0	1.8

Note: Columns indicate the change to the initial reservation price. The first four columns refer to vehicles sold individually and the last three to bundled vehicles. The first line for each auction indicates the mean reservation price after the manipulation in the experiment. The second line indicates the mean sale price for each auction. The third line shows the appreciation over the original reservation price (not the final reservation price). The fourth line shows the percentage of vehicles put for sale that were sold and the last line for each auction shows the average number of bidders on vehicles that were ultimately purchased. The auction in Mumbai has no information about the actual number of bids, just the final sale price and the ID of the winning bidder.

Table 3.4: Bidder Characterization

		Hyderabad 10/10/2006	Hyderabad 11/27/2006	Chennai 1/27/2007	Mumbai 8/14/2007
No. bidders by auction		23	15	6	13
No. buyers by auction		14	10	5	13
No. bids by participant	Mean	5.7	10.6	4.3	4
	Median	3	4	2	5
	Max	20	55	10	6
	P75	9	10	9	5
	P25	1	1	2	2
No. purchases by participant	Mean	3.9	7.6	4	4
	Median	3	2	2	5
	Max	8	39	9	6
	P75	7	6	7	5
	P25	1	1	1	2

Table 3.5: Probability of Sale

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Increase	2.06	2.39	3.64	2.86				
	(7.79)	(7.88)	(7.48)	(7.71)				
Discount	17.37***	18.05***	19.79***	18.20***				
	(4.53)	(5.10)	(5.10)	(6.63)				
Percent Discount					1.57***	1.61***	1.78***	1.53**
					(0.41)	(0.47)	(0.46)	(0.62)
Auction #2	-6.20	-6.21	-7.42	-7.48	-6.19	-6.19	-7.40	-7.50
	(5.21)	(5.21)	(5.02)	(5.03)	(5.23)	(5.23)	(5.04)	(5.04)
Auction #3	-22.28***	-22.22***	-28.35***	-28.50***	-22.39***	-22.36***	-28.46***	-28.69***
	(5.33)	(5.34)	(4.54)	(4.56)	(5.33)	(5.33)	(4.54)	(4.55)
Auction #4	-30.75***	-30.76***	-45.93***	-45.92***	-25.60***	-25.41***	-41.83***	-42.58***
	(4.93)	(4.93)	(4.69)	(4.71)	(5.12)	(5.21)	(4.90)	(5.06)
Bundle	-6.62	-5.16	-20.97***	-21.33***	-6.59	-5.69	-21.31***	-21.90***
	(5.16)	(7.16)	(5.89)	(5.94)	(5.17)	(7.16)	(5.87)	(5.91)
Bundle*Discount		-3.08	-5.21	-3.76		-1.92	-4.11	-1.72
		(10.12)	(9.42)	(10.30)		(10.29)	(9.62)	(10.64)
Slow			-29.16***	-30.15***			-29.11***	-30.65***
			(4.25)	(5.03)			(4.26)	(5.00)
Slow*Discount				3.75				5.95
				(9.88)				(10.12)
Number of obs.	544	544	544	544	544	544	544	544

Note: Table shows the marginal effects of logit regressions where the dependent variable is an indicator for whether the vehicle was sold in an auction. The first four specifications include two indicator variables for whether the initial minimum bid was increased or reduced. The last four specifications include the percentage discount for each vehicle as independent variable. All regressions include auction fixed effects. "Slow" indicates vehicles that were deemed by the auctioning bank to be harder to sell at the price initially set. Bundle indicates that the vehicle was included in a two-vehicle bundle. Standard errors are clustered at the bundle level and are shown in parenthesis

Table 3.6: Price Appreciation Conditional on Sale (Discount as Indicator Variable)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Increase	-0.86	-0.61	-0.54	1.25	-0.58	0.60	0.64	0.86
	(4.19)	(4.24)	(4.25)	(4.31)	(4.45)	(4.44)	(4.43)	(4.51)
Discount	-7.80***	-7.39***	-7.28***	-4.35*	-0.01	2.44	2.51	2.97
	(1.89)	(2.12)	(2.14)	(2.57)	(2.28)	(2.56)	(2.56)	(3.00)
Auction #2	-2.60	-2.58	-2.52	-2.74	-2.16	-2.08	-2.43	-2.45
	(2.20)	(2.21)	(2.21)	(2.20)	(2.62)	(2.60)	(2.60)	(2.61)
Auction #3	-5.86*	-5.78*	-6.22*	-5.77*	-6.00	-5.62	-4.48	-4.38
	(3.35)	(3.36)	(3.47)	(3.45)	(3.82)	(3.79)	(3.88)	(3.90)
Auction #4	3.42	3.41	2.83	2.43	2.39	2.18	3.54	3.50
	(2.86)	(2.86)	(3.07)	(3.06)	(3.36)	(3.33)	(3.48)	(3.49)
Bundle	-7.56***	-6.41*	-6.90*	-5.45	-11.41***	-6.70*	-5.25	-5.08
	(2.36)	(3.57)	(3.70)	(3.74)	(2.78)	(3.61)	(3.76)	(3.82)
Bundle*Discount		-2.00	-2.13	-5.03		-10.84**	-10.85**	-11.30**
		(4.66)	(4.67)	(4.85)		(5.37)	(5.36)	(5.59)
Slow			-1.30	4.18			4.02	4.67
			(2.52)	(3.69)			(3.07)	(3.80)
Slow*Discount				-9.22**				-1.69
Number of obs.	226	226	226	226	171	171	171	171
Pseudo R2	0.17	0.17	0.17	0.19	0.14	0.16	0.17	0.17

Note: Table shows the coefficients of OLS regressions where the dependent variable is the percentage difference between the sale price and the initial minimum bid set by the bank for each vehicle sold in an auction. The "discount" and "increase" variables are two indicator variables for whether the initial minimum bid was increased or reduced. The first four specifications include all vehicles put for auction, the last four specifications exclude discounted vehicles that sold below their original minimum bid. "Slow" indicates vehicles that were deemed by the auctioning bank to be harder to sell at the minimum bid initially set. Bundle indicates that the vehicle was included in a two-vehicle bundle. Standard errors are clustered at the bundle level and are shown in parenthesis.

Table 3.7: Price Appreciation Conditional on Sale (Discount as Continuous Variable)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Discount	-0.82***	-0.79***	-0.78***	-0.52**	-0.05	0.18	0.19	0.21
	(0.19)	(0.21)	(0.21)	(0.25)	(0.23)	(0.26)	(0.26)	(0.30)
Auction #2	-2.59	-2.58	-2.51	-2.72	-2.16	-2.09	-2.44	-2.45
	(2.19)	(2.20)	(2.20)	(2.20)	(2.61)	(2.59)	(2.60)	(2.61)
Auction #3	-5.96*	-5.90*	-6.32*	-5.92*	-6.13	-5.77	-4.63	-4.59
	(3.33)	(3.34)	(3.45)	(3.44)	(3.81)	(3.78)	(3.87)	(3.90)
Auction #4	0.04	0.20	-0.29	0.82	1.95	2.95	4.35	4.44
	(2.73)	(2.78)	(2.95)	(3.00)	(3.11)	(3.13)	(3.30)	(3.37)
Bundle	-7.54***	-6.69*	-7.16*	-5.86	-11.37***	-6.94*	-5.49	-5.41
	(2.35)	(3.55)	(3.68)	(3.72)	(2.77)	(3.60)	(3.76)	(3.81)
Bundle*Discount		-1.49	-1.61	-4.22		-10.21*	-10.22*	-10.43*
		(4.63)	(4.64)	(4.83)		(5.36)	(5.35)	(5.58)
Slow			-1.24	3.76			4.00	4.31
			(2.51)	(3.67)			(3.06)	(3.79)
Slow*Discount				-8.40*				-0.81
				(4.53)				(5.77)
Number of obs.	226	226	226	226	171	171	171	171
Pseudo R2	0.18	0.18	0.18	0.19	0.14	0.16	0.17	0.17

Note: Table shows the coefficients of OLS regressions where the dependent variable is the percentage difference between the sale price and the initial minimum bid set by the bank for each vehicle sold in an auction. The "discount" variable is the amount by which the original minimum bid was reduced or increased as a percentage of the original minimum bid. The first four specifications include all vehicles put for auction, the last four specifications exclude discounted vehicles that sold below their original minimum bid. "Slow" indicates vehicles that were deemed by the auctioning bank to be harder to sell at the minimum bid initially set. Bundle indicates that the vehicle was included in a two-vehicle bundle. Standard errors are clustered at the bundle level and are shown in parenthesis.

Table 3.8: Probability of No Appreciation

	(1)	(2)	(3)	(4)
Discount	-15.69**	-10.98	-15.11**	-13.99
	(6.37)	(7.63)	(7.50)	(11.39)
Auction #2	17.84**	17.96**	16.17**	16.14**
	(8.09)	(8.10)	(7.40)	(7.40)
Auction #3	42.73***	43.70***	52.19***	52.31***
	(10.02)	(9.96)	(8.11)	(8.17)
Bundle	17.88**	27.58**	39.03***	39.41***
	(8.36)	(12.62)	(10.87)	(11.45)
Bundle*Discount		-15.19	-10.02	-11.01
		(12.40)	(13.21)	(15.25)
Slow			34.91***	35.72***
			(8.77)	(12.16)
Slow*Discount				-1.90
				(15.56)
Number of obs.	174	174	174	174
Pseudo R2	0.11	0.12	0.20	0.20

Note: Table shows the marginal effects of logit regressions where the dependent variable is an indicator for whether the vehicle was sold at the minimum bid set by the auctioneer (i.e. appreciation = 0). All specifications include two indicator variables for whether the initial minimum bid was reduced. All regressions include auction fixed effects. Only the two Hyderabad auctions and the auction in Chennai are included, given that all sold vehicles in Mumbai were sold above the minimum bid. "Slow" indicates vehicles that were deemed by the auctioning bank to be harder to sell at the minimum bid initially set by the bank. Bundle indicates that the vehicle was included in a two-vehicle bundle. Standard errors are clustered at the bundle level and are shown in parenthesis.

Table 3.9: Number of Bidders

	(1)	(2)	(3)	(4)
Discount	0.20	0.15	0.23*	0.26
	(0.13)	(0.15)	(0.14)	(0.18)
Auction #2	-0.41***	-0.41***	-0.37***	-0.37***
	(0.13)	(0.14)	(0.12)	(0.12)
Auction #3	-1.06***	-1.07***	-1.34***	-1.33***
	(0.21)	(0.21)	(0.20)	(0.20)
Bundle	-0.34**	-0.48**	-0.77***	-0.75***
	(0.14)	(0.22)	(0.21)	(0.22)
Bundle*Discount		0.23	0.14	0.11
		(0.29)	(0.27)	(0.29)
Slow			-0.80***	-0.76***
			(0.14)	(0.22)
Slow*Discount				-0.07
				(0.28)
Number of obs.	174	174	174	174
Pseudo R2	0.18	0.19	0.32	0.32

Note: Table shows the coefficients of OLS regressions where the dependent variable is the number of bidders for each vehicle sold in an auction. All specifications include an indicator variable for whether the initial minimum bid was reduced. All regressions include auction fixed effects. Specifications (2) and (5) only include one vehicle of each bundle that was constructed. Only the two Hyderabad auctions and the auction in Chennai are included, given that there is no information on the number of bidders for the auction in Mumbai. "Slow" indicates vehicles that were deemed by the auctioning bank to be harder to sell at the minimum bid initially set by the bank. Bundle indicates that the vehicle was included in a two-vehicle bundle. Standard errors are clustered at the bundle level and are shown in parenthesis.

Table 3.10: Distribution of Bids

	Percentage of original minimum bid									
	$< 95\%$	$< 100\%$	$< 105\%$	$< 110\%$	$< 115\%$	$< 120\%$	$< 125\%$	$< 130\%$	$< 135\%$	$\geq 135\%$
Increased			0.19	0.31		0.25	0.13	0.06		0.06
			(0.1)	(0.12)		(0.11)	(0.09)	(0.06)		(0.06)
			3	5		4	2	1		1
Unchanged			0.85	0.34	0.25	0.08	0.11	0.03	0.01	0.11
			(0.08)	(0.08)	(0.06)	(0.03)	(0.04)	(0.02)	(0.01)	(0.05)
			79	32	23	7	10	3	1	10
Discount	0.78	0.27	0.38	0.24	0.07	0.04	0.04	0.04	0.03	0.06
	(0.08)	(0.06)	(0.08)	(0.06)	(0.03)	(0.02)	(0.03)	(0.02)	(0.02)	(0.03)
	91	32	45	28	8	5	5	5	4	7

Note: Table shows the average number of bids per sold motorcycle for each subsample. Standard deviation of the mean in parenthesis and the number of bids below the standard deviation. Each column gives (non-overlapping) ranges of the original minimum bid. Bundles are included as one observation.