Asymmetric information and the supply-chain of mortgages: The case of Ginnie Mae loans*

Jonathan Becker[1], Kenneth Hendricks[1,2], Jean-François Houde[1,2] and Diwakar Raisingh[3]

[1]University of Wisconsin-Madison
[2]NBER
[3]Department of Justice

June 10, 2024

Abstract

This paper studies the cost of financial intermediation services provided by traditional banks and non-bank lenders in the market for mortgages. A unique feature of the supply chain of mortgages in the US is the fact that over 90% of "conforming" loans are securitized, and roughly 50% of these loans are originated and serviced by different lenders he cost of financial intermediation is therefore determined by the resale value of loans in the secondary and wholesale markets. We develop a simple modeling framework for loan valuation in these two markets. The prices banks are willing to pay for loans in the wholesale market depend upon their resale prices in the MBS market, and on the stream of fees that they can earn from servicing the loans. The value of the loan is common to all banks, and early prepayment is the primary source of risk that they face. The main friction is private information about this risk. The goals of the paper are to validate the model of loan valuation and to test for adverse selection in the wholesale and MBS markets. We do so using a proprietary auction dataset from one of the largest loan exchange platform intermediating wholesale transactions between loan originators and banks responsible of securitizing and servicing loans. Our empirical results support the hypothesis that banks are privately and asymmetrically informed about prepayment risk, consistent with the existence of a Winner's Curse in the mortgage market.

1 Introduction

The voluminous literature studying the mortgage market has focused mostly on the retail (origination) segment in which financial intermediaries acquire loans from borrowers. According to

*Pleminary and incomplete. Please do not cite without the permission of the authors.

the traditional banking model, vertically integrated financial institutions (traditional banks) fund mortgages using deposits and then keep the resulting loans on their balance sheets. Retail margins capture the difference between the interest rate on loans and the marginal cost of attracting deposits. In practice, in the U.S., only a relatively small share of conforming loans are funded this way.[1] Instead, the majority of conforming loans made by traditional banks are eventually pooled into mortgage-backed securities (MBS) and sold on secondary markets to investors such as pension funds, hedge funds or foreign banks. The banks retain the servicing rights and earn a monthly fee for collecting and distributing the monthly payments of the borrower. This fee is equal to the note rate of the loan minus the agency's fee for insuring the loan against default and the coupon paid to the investors. This structure corresponds to the *originate-to-distribute* model. Securitization greatly increases the supply of funds to the U.S. mortgage market. It also opens the door to non-depository mortgage specialists (brokers and correspondent lenders) that compete with traditional banks to attract borrowers. For mortgage specialists, the primary funding source is the wholesale market, in which other financial institutions (both traditional and shadow banks) compete to acquire loans via auctions or posted-prices.

This paper provides the first comprehensive analysis of the originate-to-distribute supply chain for conforming mortgages. We advance a simple modeling framework for loan valuation in the wholesale and secondary markets for mortgages. The prices that banks are willing to pay for loans in the wholesale market depend upon their resale prices in the secondary market, and on the stream of fees that they can earn from servicing the loans. These prices in turn depend on the duration of the loan: the flow of payments to the bank for servicing the loan and to the investors for funding the loan end when the loan is prepaid, either by the borrower or by the agency in the event of default. Thus, early prepayment is the main source of risk that banks and investors face. The main friction is private information about this risk: originators who sell in the wholesale market have more information about loan duration than the banks, and the banks have more information about loan duration than investors when they sell the loans in the secondary market. These asymmetries in information can give rise to adverse selection, increase the costs of intermediation, and prevent the efficient flow of funds.

The main goals of this paper are to validate our model of loan valuation, test for adverse selection in the two markets, and quantify their effects. The empirical analysis requires data covering the entire lifespan of a loan, from origination, to possible sale in the wholesale market, to sale in the secondary market, and to payment history and duration. Until now, such data have been unavailable to researchers, but we have been able to obtain access to three data sources that, once combined, allow us to construct the full history of a loan. A unique aspect of our study is the use of proprietary data on loan acquisition by auction from the FinTech company Optimal Blue (OB), which operates one of the largest loan exchange platforms. Sale by auction has facilitated entry

[1] Loans are said to be conforming if they satisfy the underwriting criteria of three federal agencies Freddie Mac, Fannie Mae, and Ginnie Mae.

into the primary market of lenders with limited capital who specialize in mortgage origination, and is a rapidly growing segment of the wholesale market.

We use the data on loan duration and auctions to establish several important facts..First, loan survival is positively correlated with loan and borrower characteristics. We estimate the probability of a loan surviving for different periods and find that the characteristics can explain a significant percentage of the variation in loan survival. Second, bidders value loan duration. We regress bids on the loan characteristics observed by bidders and find they bid more for loans with longer expected duration. Third, the auction is a common value auction in which the bidders are asymmetrically and privately informed about loan survival. Using the residuals of the bid regressions as measures of the bidders' signals, we find that a bidder's signal, and the maximum among it's rivals, are positively correlated with loan survival. The correlation with own signal is lower than the correlation with the maximum rival signal, and varies across the bidders, which suggests that the bidders are differentially informed. These facts provide strong support for our model of valuations.

Buyers in the wholesale market observe a subset of the characteristics observed by originators selling the loan. For example, the buyers do not observe the borrower's age, race, gender, or choice of fees, which are significant predictors of loan duration. Since these characteristics are not individually priced by the buyers, sellers have an incentive to use their private information to keep and securitize loans that they believe are more likely to survive, and to sell the others. Bidders in turn should respond to this adverse selection by lowering their bids. We use exogenous variation in the capacity of sellers to act on this incentive to determine whether, and by whom, the loans sold in the wholesale market are adversely selected.

A similar situation arises in the secondary market. Banks with a large volume of loans can choose to sell their loans in a To-Be-Announced (TBA), multi-issuer security or in a custom pool security. The TBA market is a highly liquid, forward market in which a seller and a buyer agrees to trade a volume of loans at a specified price and future date. Importantly, the investors do not observe the characteristics of the individual loans, because the loans are not selected at the time of the trade. The custom pool market is less liquid, but investors observe the characteristics of the individual loans in the custom pool security when the trade is made. As a result, the large banks have an incentive to sell loans that are less likely to survive in the TBA market, and loans that more likely to survive in the custom pool market, where they can be sold at higher price.

Banks must also decide whether to sell the loan in high coupon security or a low coupon security. They can get more money up front by selling a loan in a high coupon security (i.e., high price, low fee) or more money later by selling the loan in a low coupon security (i.e., low price, high fee). Clearly, this choice depends on the bank's assessment of the loan's duration. It should place loans that are more likely to survive in a low coupon security, and loans that are more likely to be prepaid early in a high coupon security. We test this hypothesis for loans sold in the TBA market only since we observe the security prices.in this market but not in the custom pool market.

To assess the importance of adverse selection in the secondary market, we construct a test that is commonly used in the insurance literature (Chiappori & Salanie 2000). Specifically, we test whether the 12-month survival rate is higher for loans placed in custom pool securities than in TBA securities, and for loans placed in low coupon TBA securities than for high coupon TBA securities. The results strongly support the hypothesis of adverse selection, both on observables and unobservables. The magnitude of the adverse selection effects are especially large for pool choice, even after controlling for loan characteristics.

The correlation between coupon choice and survival can be due to moral hazard. Lenders may be causing borrowers to refinance their loans early in order to obtain higher service income on the new loan. We explore this hypothesis using the sample of Ginnie Mae loans where the coupon is uniquely determined by the note rate, so banks have no choice. We find that the unconditional correlation between survival and service income is indeed positive but small, and it is essentially zero.once we condition on loan attributes. Therefore, we cannot reject moral hazard on observables but we can reject it on unobservables, which suggests that the correlation between coupon choice and survival is likely due to adverse selection.

This paper contributes to three important literatures at the intersection of Industrial Organization and Finance. First, it contributes to a large literature measuring the importance of adverse selection in financial and insurance markets.[2] Second, it builds on a growing number of papers analyzing the Industrial Organization of the mortgage industry (e.g. Stanton et al. (2014), Allen et al. (2019), Robles-Garcia (2019), Buchak et al. (2020)).

Finally, it contributes to the empirical auction literature; in particular we use insights from important empirical papers studying the design of auctions in financial markets, and auctions with common values (e.g. Hendricks et al. (2003), Haile (2002), Hortaçsu & McAdams (2010), Hortaçsu & Kastl (2012), Cassola et al. (2013)).

2 Supply Chain of Mortgages in the US

Figure 1 illustrates the flow of funds from investors and depositors to borrowers. In the retail market, banks and mortgage specialists provide loans to borrowers seeking to buy a home or refinance an existing mortgage. Borrowers are presented with a rate sheet of note rates and upfront payments, known as "discount points", that specify the amount that the borrower would need to pay at closing to lower the note rate (andreas et al. 2013). Based on these rate sheets, the borrower selects a lender and a note rate. The lender must then decide whether to keep the loan on its books, sell the loan to another lender in the wholesale market, or securitize the loan and sell it in the secondary market. Loans that are originated, securitized, and sold in the secondary market by vertically integrated banks are called *retail* loans. Loans that are sold in the wholesale

[2] Closely related to this paper include the work of Agarwal et al. (2012) and Downing et al. (2009). See Einav et al. (2010) for a review of the insurance literature.

Figure 1: Supply chain of mortgages in U.S.

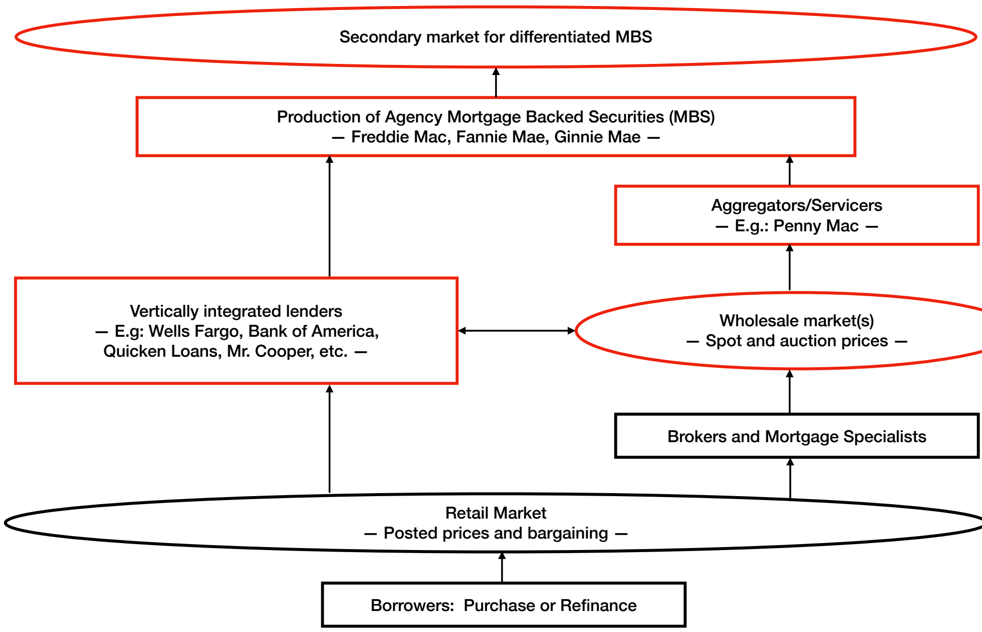

market and then securitized and resold in the secondary market are called *non-retail* loans. The wholesale channel is split between brokers and correspondent lenders. A *broker* matches a borrower to a bank who underwrites and funds the loan at closing. A *correspondent lender* is a mortgage specialist or bank who underwrites and funds the mortgage at closing and then sells the loan to an ““aggregator” or sponsor bank for securitization, typically a few days after the closing date. In our sample of Ginnie Mae loans, which covers the period from 2013 to 2022, roughly 40% of all securitized loans were originated by correspondent lenders, and 10% by brokers. The remaining 40% of loans were retail.

The traditional way that correspondent lenders sold their loans was in a posted-price market. Buyers post wholesale rate sheets daily that specify the prices they were willing to pay for loans. These prices are known as*lock* prices because they vary with the lock-in period, which can be0, 30, 60, or 90 days. They are also seller-specific, and contingent on a very coarse binning of loan characteristics. Given these characteristics and a lock-in period, the correspondent lender selects a buyer, and that buyer then agrees to buy the loan shortly *after* closing. With this sales mechanism, both the seller and the buyer incur holding costs. They also bear the risk of “fall-out” which are loan applications that do not result in a closed loan, either because the borrower does not qualify or because they turn down the offer.

More recently, online platforms such as *Optimal Blue* have provided correspondent lenders with the option to sell their loans individually in online auctions. Buyers in the auction have the same loan information they would have in the posted price market. However, the auction allocation is likely to be more efficient (i.e., the buyer with the highest valuation is more likely to get the loan), since the bids that buyers submit are a continuous function of loan characteristics. The auction may also be less costly, because only the seller incurs holding costs. The data do not allow us to distinguish between posted price sales and auction sales, but auction sales represent a growing segment of the wholesale market. Specifically,auction sales on the OB platform accounted for over 75% of loan sales in2019.

Most mortgages are securitized. The securitization process involves pooling many different loans and issuing a mortgage-backed security (MBS) backed by these loans. The security is then sold (in tranches) to investors such as foreign banks, hedge funds, and pension funds. In the case of conforming loans, the GSEs, Fannie Mae and Freddie Mac, purchase the loans from the banks, insure them against default, and issue the MBSs. In the case of government-insured loans, Ginnie Mae insures the loans against default, but does not issue securities directly. Banks are responsible for delivering and managing loan packages, subject to Ginnie's underwriting rules. In both cases, the banks typically retain the servicing rights to the loans. They collect the interest payment from the borrower each month, pay the MBS coupon to investors and the guarantee fee to the agency, and keep the difference, known as the *service income*. Both investors and banks bear the risk of loan prepayment.

Most agency MBSs are sold in the To-Be-Announced (TBA) market. TBA trades are forward contracts: a seller and buyer agree to trade a volume of loans(par value) of agency MBS at a specified price and future date (settlement date). The time between the date of the trade and date of the settlement is typically several weeks. When the trade is made, the mortgages in the MBS are not known, in part because they often do not yet exist. Instead, the two parties agree upon the issuer (agency guarantor), coupon, and maturity of the loans. In selecting the pool of loans, the seller has an incentive to deliver loans that are more likely to default or be refinanced. However,according to Vickery & Wright (2013), the buyer understands this incentive,anticipates that loans will be adversely selected, and prices accordingly.[3] In principle, the market could unravel, but the lack of loan information makes the market more homogeneous and liquid, and the liquidity premium more than offsets the adverse selection discount (see Vickery & Wright (2013)). An important benefit of the TBA market for lenders in the primary and wholesale markets is that it locks in a resale price for the loans that they originate or buy. Banks with large volumes of loans can choose to sell their loans in another,less liquid, securities market. It is called the "specified-pool" or "custom-pool" market, because the characteristics of the loans in the agency MBS are known when the trade is made. During our sample period, roughly 20% of the securitized loans are pooled in custom TBA-eligible securities.

Some banks such as Quicken Loans or Bank of America rely almost exclusively on the retail

[3]This price is known in the finance literature as the "cheapest-to-deliver" price.

channel to acquire loans. In contrast, several large shadow-banks such as Pennymac originate a very small quantity of loans directly, and rely on the wholesale market to acquire loans. However, most banks and shadow banks manage a diversified portfolio of retail and non-retail loans.

As originators in the retail market, these lenders have to decide whether to keep a loan or sell it in the wholesale market and, as buyers in the wholesale market, they have to set lock prices in the posted price market and submit bids in the auction market. As sellers in the secondary market,the banks need to decide whether to sell in the TBA or the specified pool market and, in these markets, choose the coupon for their loan. In what follows, we study these decisions in the context of the auction market.

3 Data

In this section, we describe the three main sources of data used in the empirical analysis and how we track loans across these data sets. We focus on 30-year fixed rate mortgages insured by Ginnie Mae, a public corporation responsible of insuring default risks for loans qualifying for FHA, VA, and rural housing subsidies.[4] Because Ginnie Mae guarantees loans with higher LTV ratios, borrowers tend to be riskier (both in terms of default and prepayment). FHA borrowers in particular are often first-time home-buyers who eventually transit to conventional products after building enough equity. As a result, these loans cannot be sold to the GSEs, Fannie Mae and Freddie Mac, which implies that this segment operates more or less independently of the other segments. The other reason for focusing on Ginnie Mae loans is that Ginnie Mae does not discriminate between lenders when setting its guarantee fee. It is fixed exogenously and the same for all lenders.[5]

3.1 Loan Securitization and Performance

The first data set, *eMBS*, provides detailed information all mortgage backed securities insured by one of the three agencies that were issued from January 2013 to the present and their component loans. In our analysis, we use data Ginnie Mae securities issued between October 2013 and December 2019. The characteristics of the MBS include the CUSIP (security identifier), coupon rate, issuance date, issuer and servicer identity, maturity, and par amount. The characteristics of the loans include: the CUSIP with which it is associated, the subsidizing agency, the loan type (purchase, refinance, etc.), original principal balance, note rate, loan-to-value (LTV), debt to income ratio, FICO score, number of units on the property, state, origination type (retail or non-retail), and the identity of the issuer which, in case of Ginnie Mae, is the servicer. For each component

[4]This term is the most common, accounting for 93% of the loans in our data. The remainder is divided between fixed rate mortgages with different maturities (6%) and adjustable rate mortgages (1%).

[5]The GSE's also charge a fixed, monthly g-fee, but they allow the lender to "buy down" the fee by converting the flow into an upfront payment or, alternatively, to "buy up" the fee and receive an upfront payment from the GSE (see andreas et al. (2013)).

loan, we observe the unpaid principal balance on a monthly basis until it has either been paid off or defaulted. We use the unpaid balance for each loan in March 2022 to measure the loan duration until full payment.

We use the above information to infer the service income earned each month by the servicer. In regards to the security prices, we do not observe the prices of the custom pool securities but are able to obtain the daily agency TBA MBS prices from Bloomberg.[6]

Table 1 1 provides summary statistics on the characteristics of the loans. The sample consists of 30-year, fixed rate Ginnie Mae loans issued between October 2013 to December 2019 that had at least 6 months of performance data.[7] Note rates are typically quoted on a 1/8 percentage points (p.p) grid. The note rate varies between 3% and 5 % and averages 4%. The average loan size is $210K, but it varies a lot. The loans typically have a very high LTV, with 58% of them having values between 95 and 100. Moreover, the FICO scores of the borrowers are a relatively low 687, which is just above 670, the cutoff between "fair" and "good" credit. The LTV's and FICO scores are consistent with the goal of the subsidy programs. The largest subsidy category is FHA loans. These loans increased in popularity after the financial crisis, replacing privately securitized subprime loans, because the FHA program allow borrowers with low credit-scores and/or high LTV ratios to access the mortgage market (although they do incur higher insurance payments over the life of the contract).

Coupon rates are quoted on a 1/2 p.p grid. Figure 2 gives the frequency distribution of the coupon rate of the Ginnie Mae securities in our sample. Most pools pay out a coupon that is between 3.5% to 5.0% and this accounts for over 90% of the pools in the data.

Since consumers can exert the option to pre-pay their loan early or default, very few loans last until maturity. Figure 3 quantifies this risk for the cohorts of Ginnie Mae loans originated between 2013 and 2019. The pre-2017 cohorts faced relatively stable interest rates. On average, 70% of those loans were pre-paid within the first six years, and about 10% were pre-paid within the first year. The risk of early pre-payment increased substantially for loans originated in 2018 and 2019, due to the dramatic decline in mortgage interest rates observed between 2019 and 2021. This affected especially the 2019 cohort. Nearly 50% of those loans were prepaid within the first 18 months, and 30% of loans were pre-paid during the first year. Since our data on auctioned loans covers the 2018-2020 period, we use the risk of early pre-payment as our primary measure of loan performance.

---

[6] Bloomberg sources its pricing data from Trade Reporting and Compliance Engine (TRACE), which is a database of trades maintained by Financial Industry Regulatory Authority (FINRA). This database contains the universe of TBA bond trades for which one of the parties was registered with FINRA. TBA MBS trades are typically made with a FINRA registered dealer so TRACE should contain nearly all trades (see Gao et al. (2017)). The daily price of a TBA security corresponds to the last observed trading price as of that date.

[7] We truncate the sample to avoid the Covid shock to the markets.

Table 1: Summary of Mortgages in Ginnie Mae MBS's

(a) All Loans: 2013-2019

VARIABLES	Mean	SD	P-10	P-90
Note rate	4.2	.56	3.5	4.9
Loan (x100K)	2.2	1.1	1	3.5
LTV	95	8.4	85	101
Credit Score	688	54	625	769
DTI	41	9.6	28	53
1(DTI > 40)	.58	.49	0	1
1(VA)	.34	.47	0	1
1(Second lien)	.06	.24	0	0
1(Purchase)	.76	.43	0	1
1(Retail)	.39	.49	0	1
1(Corr.)	.49	.5	0	1
Survival: 12m	89	31	0	100
Survival: 36m	57	50	0	100
Observations	751794			

(b) Matched sample: 2018-2019

VARIABLES	Mean	SD	P-10	P-90
Note rate	4.4	.61	3.5	5.1
Loan (x100K)	2.3	1	1.2	3.6
LTV	96	7	86	102
Credit Score	687	52	626	767
DTI	43	10	30	50
1(DTI > 40)	.63	.48	0	1
1(VA)	.29	.45	0	1
1(Second lien)	.019	.14	0	0
1(Purchase)	.83	.37	0	1
1(Retail)	.0023	.048	0	0
1(Corr.)	.97	.16	1	1
Survival: 12m	82	38	0	100
Survival: 36m.	30	46	0	100
Observations	53843			

Reports the summary statistics of loans that entered Ginnie Mae 30 Year Fixed Rate MBS's issued between October 2013 and December 2019 that had at least 12 months of performance data. LTV, debt to income ratio, and FICO scores are missing for some of the mortgages in our data set.

Figure 2: Density of Ginnie Mae's MBS coupons and service income between 2013 and 2019

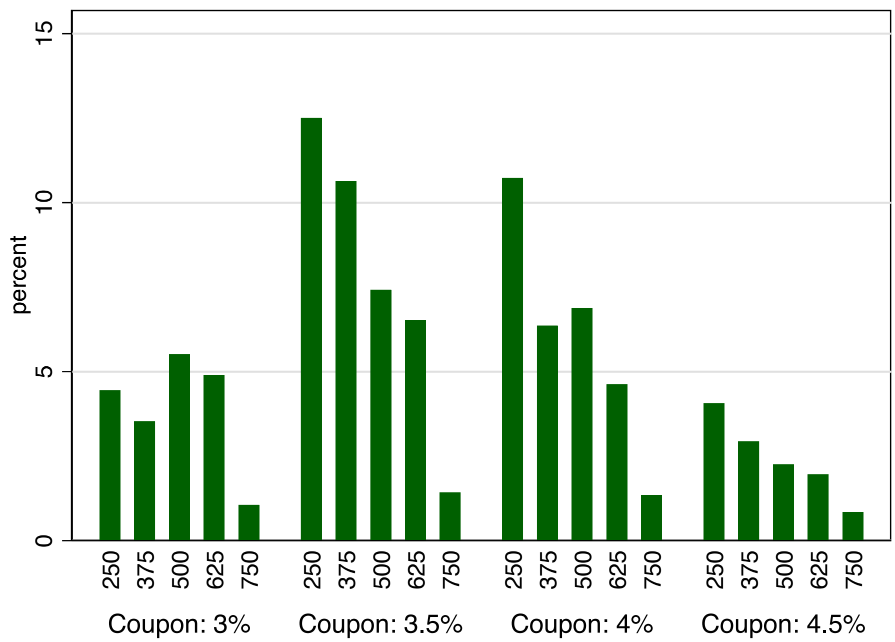

The x-axis groups loans based on the MBS coupon (bottom number, c) and service income (top number, $r - c$). The sample excludes infrequently used coupons (≤ 2.5 and ≥ 5).

Figure 3: Evolution of pre-payment risk across cohorts

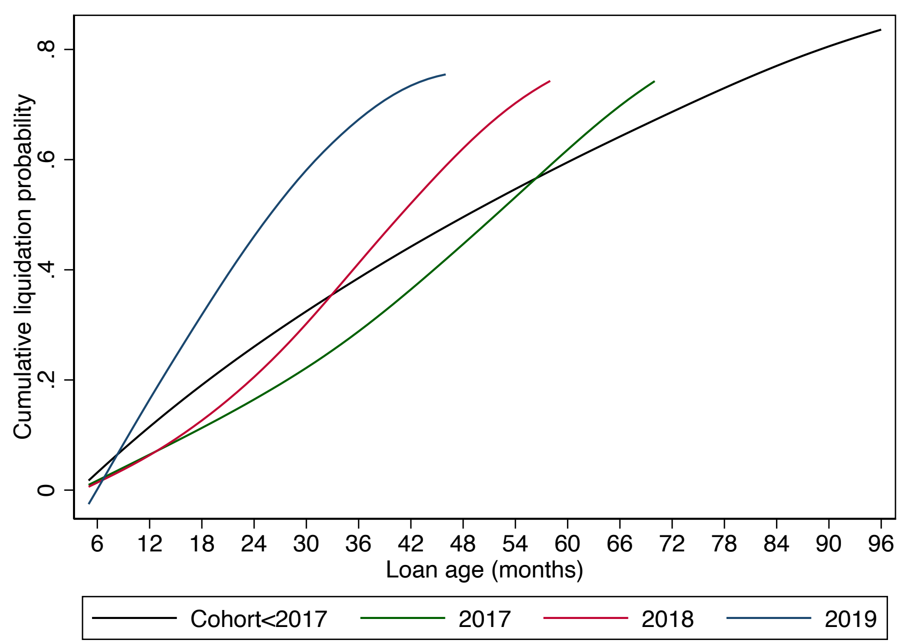

Table 2: Survival probability regression

VARIABLES	(1) $1(T > 12)$	(2) $1(T > 36)$	(3) $1(T > 12)$	(4) $1(T > 12)$
Note rate	-9.53	-15.4	-10.1	-11.4
	(0.21)	(0.31)	(0.31)	(0.60)
Loan amount (x100K)	-6.24	-18.0	-9.11	-11.4
	(0.26)	(0.45)	(0.42)	(0.76)
Loan amount squared (x100K)	0.79	3.54	1.21	1.83
	(0.066)	(0.11)	(0.10)	(0.24)
1(VA)	-6.58	-5.55	-11.7	-14.2
	(0.29)	(0.29)	(0.40)	(0.64)
1(Second lien)	5.85	10.9	6.56	1.06
	(0.37)	(0.56)	(0.43)	(1.05)
LTV	0.044	0.12	0.070	0.25
	(0.0062)	(0.0092)	(0.011)	(0.032)
Credit score groups = 2, 630-690: Fair	-0.30	-1.11	-0.72	-2.20
	(0.14)	(0.21)	(0.25)	(0.56)
Credit score groups = 3, 690-720: Good	-0.87	-1.79	-1.95	-3.75
	(0.18)	(0.27)	(0.32)	(0.71)
Credit score groups = 4, 720-850: Excellent	-0.73	-0.91	-2.32	-4.37
	(0.19)	(0.29)	(0.34)	(0.71)
Loan type = 2, Refi: Not Streamlined, Not Cash Out	-4.05	-5.05	-5.08	-3.21
	(0.23)	(0.33)	(0.47)	(0.97)
Loan type = 3, Refi: Cash out	-6.26	-8.00	-7.10	-5.95
	(0.20)	(0.26)	(0.33)	(0.75)
Loan type = 4, Refi: Streamlined	-3.82	-3.20	-8.86	1.67
	(0.36)	(0.64)	(0.81)	(2.29)
1(DTI > 40)	-1.00	-1.89	-1.37	-2.00
	(0.074)	(0.12)	(0.15)	(0.33)
1(Retail)	0.35	1.10	-1.41	6.73
	(0.18)	(0.25)	(0.33)	(5.87)
1(Correspondent)	0.53	0.27	-0.57	-1.33
	(0.15)	(0.22)	(0.28)	(1.04)
Constant	140	140	147	139
	(1.48)	(1.95)	(2.29)	(4.40)
Observations	748,612	631,127	273,303	48,551
R-squared	0.128	0.194	0.152	0.160
Pool date FE	yes	yes	yes	yes
State FE	yes	yes	yes	yes
Issuer FE	yes	yes	yes	yes
Sample	2013-18	2013-18	2018-19	Matched
Mean dep. variable	89.1	57	83.1	82.3

Robust standard errors in parentheses

The sample in columns (1-3) correspond to a 10% random sample of Ginnie Mae loans issued between 2013 and 2019. Standard errors are clustered at the note-rate/issuing date level.

3.2 Auction Market

The mortgage auction data comes from Optimal Blue, a FinTech firm that operates the largest loan exchange platform in the market. Mortgage originators use this platform to sell loans to banks to free up capital that they can use to originate more mortgages. The OB auction is a first-price, sealed bid auction that lasts one to two hours. The seller usually invites all the buyers in its network of buyers (typically 8 to 15 buyers) to submit bids, called *bulk* bids, for the loan. For more specialized loans, the seller may invite fewer buyers. The set of potential bidders varies across lenders, but vary very little across time. Forming a relationship is costly because it involves both parties conducting due diligence as to the reliability and underwriting standards of their counterparty. In most auctions, all invited bidders submit bulk bids and nearly all do, because bidding is essentially cost-less and is a way of maintaining the relationship with the seller. A bidder can always submit a low bid if it does not want to purchase the mortgage.

An unusual feature of the auction is that the seller lender can always choose to sell the loan at the buyer's lock price if this price is higher than its bulk bid. Thus, from the seller's perspective, each buyer's actual bid is the maximum of its lock price and bulk bid, and the winning bid is the maximum of the bulk bids and lock prices. The other unusual feature of the auction is that, after observing the bids, the seller can decide not to sell the loan, in which case it either resells the loan in a later auction or sells it in the secondary market. This event is common for conforming loans, but not for Ginnie Mae loans, since the seller of these loans bears the costs of securitization.

We focus on 30-year, fixed rate mortgages eligible for Ginnie Mae insurance sold by auction during the period January 2018 to January 2019..For each mortgage, we (and the bidders) observe the following characteristics: original principal balance, loan-to-value ratio (LTV), note rate, loan type (purchase, refinance, etc.), property type, number of units, and zip code.

Since mortgages are sold individually, each mortgage is associated with an auction. For each auction, we observe the following variables: auction date, a seller id, the number of invited bidders, the value of each submitted bulk bid if it is above the bidder's lock price and the lock price if it is not, and the associated bidder id. The sellers and bidders have unique identifiers, but we can infer their identities from eMBS and our data set on the origination market.

The characteristics of the Ginnie Mae loans sold at auction are reported in Table 3. The summary statistics are quite similar to those reported in Table 1 with three key differences: (i) more loans are used to purchase a property as opposed to being refinanced, (ii) the note rate is substantially higher, and (iii) more loans have an LTV near 100. The higher prevalence of loans for purchasing a home is likely driving the other two differences. The larger note rate corresponds to a more valuable loan because, all else equal, a higher interest rate leads more income.

The value of each bid corresponds to the wholesale price for a \$100 loan. A bid of \$100 corresponds to paying the par-value of a loan. The mortgages are typically securitized and sold on the TBA MBS market and the MBS's fetch prices above the par value of the loan (loans generate

Table 3: Summary of Mortgages Sold at Auction

Variable	N	Mean	SD	Pctl10	Pctl90
Note Rate	42,852	4.660	0.472	4	5.250
Original Loan Amount	42,852	217,173	103,622	108,080	343,000
LTV	42,852	94.486	8.339	85	100
1(LTV $\in$ (75, 80])	42,852	0.024	0.152	0	0
1(LTV $\in$ (95, 100])	42,852	0.715	0.451	0	1
Debt to Income Ratio	42,852	42.232	10.018	29.280	54.480
Monthly Income	40,501	6,042	3,189	2,900	10,130
1(Monthly Income $\geq$ \$20K)	40,501	0.006	0.076	0	0
FICO	42,852	681.438	52.224	622	760
1(Purchase)	42,852	0.831	0.375	0	1
1(Retail)	42,852	0	0	0	0
1(Agency = FHA)	42,852	0.620	0.485	0	1
1(Agency = VA)	42,852	0.295	0.456	0	1
1(Agency = Rural Housing)	42,852	0.085	0.279	0	0
1(Paid off within 12 months)	39,889	0.111	0.314	0	1
i_PaidOffOrDelinquentInOneYear	39,889	0.122	0.328	0	1

interest). The TBA MBS price effectively acts as a floor on the bids that the seller will accept since it knows that buyers will get at least that price when they securitize the loan. From a bidder's perspective, the TBA MBS price corresponds to an observable reserve price. There are some cases in which the seller accepts a price that is below the TBA MBS price and near the loan's par value. This can occur for loans that have an expected short duration (high prepayment risk) that will yield limited buyer interest. This is not to say that sellers lose money on those loans, since consumers typically pay upfront fees to the mortgage originators.

Table 4 summarizes the main outcome variables in the auction data. The bids are typically above the par value of the loan and the highest bid is typically well above the par value. The money left on the table (highest bid less the second highest bid) is on average 0.20. While this is a relatively small amount relative to the par value, it is quite substantial portion of the bid after we net out the MBS price.

The bids are nearly always above the MBS price which is consistent with the idea that it acts as an observable floor on the bids. In comparison, the lock price is effectively secret and it is common to see bids below the reserve.

3.3 Origination Market and Matching

The third data set is on loans originated between 2013 and 2022. It provides detailed information on borrower characteristics (including three-digit zip code and county), the identity of the origina-

Table 4: Summary of Bids

Variable	N	Mean	SD	Pctl10	Pctl90
# Bids in Auction	42,852	6.731	4.301	1	12
# Serious Bids in Auction	42,852	6.192	4.264	1	12
Bid	288,454	104.214	1.487	102.325	105.780
Serious Bid	265,333	104.345	1.331	102.614	105.826
Highest Bid	42,779	104.765	1.450	102.869	106.436
Money on Table[a]	32,518	0.237	0.467	0.020	0.534
Money on Table (Auctions $\geq$ 2 Serious Bids)	31,259	0.212	0.300	0.020	0.491
TBA MBS Price Associated with Loan[b]	42,779	103.180	1.349	101.703	104.891
Bid - TBA MBS Price	288,454	1.070	0.994	0.109	2.040
Serious Bid - TBA MBS Price	265,333	1.248	0.632	0.440	2.069
Highest Bid - TBA MBS Price	42,779	1.586	0.759	0.710	2.487
Auction Reserve Price[c]	33,571	104.356	1.373	102.505	105.890
Bid - Reserve Price	231,768	−0.222	0.890	−0.875	0.460
Serious Bid - Reserve Price	212,977	−0.115	0.547	−0.703	0.470
Highest Bid - Reserve Price	33,571	0.397	0.572	−0.059	1.049
Reserve Price - TBA MBS Price	33,571	1.179	0.850	0.177	2.149

[a] Money on Table is the highest bid less the second highest bid and is only computed for auctions with at least two bidders.
[b] If a mortgage can go into MBS's with two different coupons, we choose the price of the MBS with the higher coupon.
[c] Data on the reserve price was not available for a subset of the auctions.Serious bids are those that are above the TBA MBS Price - 1/32. The latter is a bit of a buffer.

tors, and fees. To assemble this data set, we combine data from three sources. First, CoreLogic provides data on the name of the retail originators (including brokers and correspondents) for all new mortgage transactions. Second, publicly available data from HMDA contain information on the universe of new mortgage transactions (including rejections), as well as key demographic characteristics of borrowers such has race, gender and income. Third, Optimal Blue's RateLock data set provides rich information on upfront lending fees (or points) that consumers are paying at closing.

To match HDMA loans with OB auction loans, we conduct a strict match on the following loan characteristics: sponsoring agency, loan term, property location (zip code and county), note rate, occupancy status, number of units, dummy variable for whether the loan was a home purchase, and a dummy variable for whether the loan was an adjustable rate mortgage. We then perform a fuzzy match on the continuous characteristics: loan amount, yearly income, DTI ratio and CLTV ratio. We were able to match approximately 85% of auction loans to a unique loan in HMDA. We employ similar procedures to match OB auction loans to loans in eMBS issuances and HDMA originations to loans in eMBS issuances for the period 2013 to 2022. We are able to match x% of HDMA originations with a unique loan in eMBS and match over 91% of auctioned loans with a unique loan in eMBS.

Using data from HMDA, we find that the wholesale market is highly concentrated, with the top-4 banks purchasing more than 45% of all agency loans. The retail market is less concentrated, with the top-4 originators acquiring less than 20% of the mortgages. This difference highlights the importance of scale when selling loans in the secondary market,.and the lack of scale when selling loans in the wholesale market. Large issuers can acquire a diversified pool of loans in the retail and wholesale markets, enabling them to produce higher value securities and to achieve lower servicing costs. By contrast, the barriers to entry in the retail market are much lower. Originators need only to invest in local networks of loan officers and real-estate agents, and they require relatively little liquidity to operate when they can sell loans at competitive prices in the wholesale market. Thus, the emergence of the online auction market has generated potentially important gains from trade: economies of scale upstream and geographic segmentation downstream.

We use the matched HMDA-eMBS data set to identify which originators are using which channels to sell loans, and the volumes that they are selling. We classify originators into three groups: *retail lenders* who securitize and service more than 95% of the loans that they originate, *correspondent-only lenders* who rely exclusively on the wholesale market to finance their retail operation, and *hybrid lenders* who sell loans using both the retail and the wholesale channel.

Table 5 provides summary statistics on the three lender types for Ginnie Mae loans.[8] The first group, correspondent-only lenders, is the largest in terms of numbers, but they only represent 5% of originations. These lenders tend to be small, as indicated by the low within-segment concentration level (the top-4 concentration measure is 15%).[9] The opposite is true for retail-only lenders. The

[8]The results for Fannie Mae and Freddie Mac are similar.

[9]Since we focus on loans that are matched between HMDA and eMBS, this only represents a subset of all lenders

Table 5: Firm size distribution across origination channels

Lender types	Number	C4	Market	Acquisition channel		Median days to securitize	
		%	share	Wholesale	Direct	Wholesale	Direct
Correspondent	429	0.15	0.05	0.97	0.03	24	17
Hybrid	782	0.22	0.53	0.45	0.55	26	18
Retail	120	0.56	0.41	0.02	0.98	25	19

Source: HMDA and eMBS. The sample includes all matched loans between HMDA and eMBS originated between 2018 and 2021. Correspondent lenders are defined as originators who securitize directly less than 5% of their volume. Retail lenders are defined as originators who securitize directly more than 95% of their volume. Hybrid lenders active in both secondary and wholesale markets. The first columns counts the number of lenders in HMDA that are matched with eBMS. The second column measures the market share of the top-4 lenders within each group, and the third column measures the overall market share of each lender types. The acquisition channel measures the share of loans originated by each group that are either sold to the wholesale market, or sold directly to the secondary market. The number of days to securitize is the median number of days between origination date and the issuing date of the security (once per month).

top-4 lenders originate 56% of the loans in that group, and their overall category market share is 41%. This group includes two of the largest originators: Bank of America and Quicken Loans. The middle category of hybrid lenders has the largest market share with 55%, and they tend to be much larger than the correspondent-only lenders.

On average, these lenders securitize directly 55% of the loans they originate, and sell the remainder on the wholesale market. For these lenders, as well as for correspondent-only lenders, the cost of originating loans is directly affected by the expected resale value of loans in the wholesale market. We estimate that, for roughly 65% of agency loans originated in the US, the wholesale price of loans has a direct impact of the price that consumers pay.

The last two columns of the panel highlight another important difference between the retail and wholesale acquisition channels. For loans acquired directly on the retail market, the median time between closing and MBS issuance is 14 days. In contrast, loans acquired on the wholesale market are securitized 26 days after origination. This highlights the main financial cost of selling loans on the wholesale market: lenders must keep loans on their balance sheet longer before receiving a payment.

4 A Model of Loan and Security Valuations

In this section, we describe the profits generated by the process of acquiring, securitizing, and servicing Ginnie Mae mortgages. Importantly, this process is identical for loans originated and serviced by the same financial institution (i.e., *retail*) and for loans originated by correspondent lenders or brokers (i.e., *wholesale*). We focus primarily on loans acquired via auctions because our data comes from a wholesale loan exchange platform, but our model of loan and security valuations

in the US. Our measure of the importance of the correspondent channel is therefore under-stated.

applies to both retail and wholesale loans.

A mortgage generates two sources of income for a bank that retains servicing rights: (i) an upfront security price obtained from the sale of the loan in the secondary market, and (ii) monthly service income. The value of the service income does not depend upon the security in which the loan is placed, but the security price does. Thus, the expected revenue of a loan depends upon the securitization decisions of the bank, and the bank must anticipate those decisions when it bids for the loan in the wholesale market. In what follows, we describe each stage of the valuation process starting with service income, and identify a set of empirical predictions that we take to the data in Section 5.

4.1 Mortgage Service Rights

The revenues from monthly service rights (or MSR) are determined by the difference between the note rate associated with the mortgage (r), the coupon c that must be paid to investors, and the guarantee fee (or *g-fee*) paid to the agency (g). These three variables are measured in percentage points (p.p.). The excess corresponds to the gross profit margin on monthly servicing activities. Coupons are chosen from a discrete grid with 0.50 p.p. increments, while the note rates are typically selected from a finer 0.125 p.p. grid increment. When loan i is placed in an MBS with a coupon of c, the ex-post revenue to the bank for a \$100 tranche is given by the upfront payment $P(c_i)$ plus the discounted value of service income

$$\underbrace{\sum_{\tau=1}^{t_i} \delta^\tau L_{\tau,i}}_{\text{service multiple } (M_i)} \times \underbrace{\frac{r_i - c_i - g}{1200}}_{\text{service income}} \equiv M_i(t_i) \times (r_i - c_i - g) \tag{1}$$

where δ is the discount rate that the banks use to weight future cash flows, $L_{\tau,i}$ is the loan balance at the end of the month τ, and $t_i \leq 360$ is the realized duration of loan i.[10] Since payments are made monthly and the units of r_i, g, and c are in percentage points, we divide by 1200 to obtain the fraction of the monthly interest payment received. The service multiplier $M_i(t_i)$ depends on loan duration and is strictly increasing in t_i. Loans that terminate later represent higher values for the mortgage servicers and investors because they generate cash-flows for a longer period.

In practice, the duration of the loan, which we denote by T_i, is a random variable. It depends on when the borrower decides to default or prepay the mortgage. If the borrower defaults, then payments end and Ginnie Mae repays the loan to investors. If the borrower prepays the mortgage early, then she does so for one of two reasons: either she wants to refinance (cash-out or regular) the mortgage or she wants to move to another home or location. Refinancing depends on interest-rate fluctuations (i.e. interest rate drops), as well as the fixed costs of negotiating and closing a

[10]For simplicity, we assume that consumers do not make unscheduled payments, and focus on the duration of the mortgage as our measure of pre-payment risk.

new mortgage contract. The presence of fixed costs implies that borrowers with larger loans are more likely to be “in the money” to refinance their existing mortgages. The decision to move is more difficult to predict based on observed loan attributes and depends on economic and socio-demographic factors, such as unemployment, relocation or divorce.

The realized service income of a loan is common to all servicers, but uncertainty about loan duration means that banks need to form beliefs about $M_i(T_i)$ when they value the servicing rights. Those beliefs can vary across banks depending upon the information they have about the loan and on how they use this information to forecast loan duration. We discuss this issue more formally Let Z_i denote the vector of loan characteristics that is observed by each bank. In what follows, we consider two possible information structures: (i) common beliefs about $M_i(T_i)|Z_i$ and (ii) heterogenous beliefs about $M_i(T_i)|Z_i$.due to private information about loan duration.

Prediction 1. *Regardless of the information structure, our model of service income valuations generates three testable implications:*

a. *Banks should bid more for loans with higher markups (i.e., $r - c - g$).*

b. *Banks should bid less for loans with higher interest rates since they are more likely to be refinanced earlier.*

c. *Banks should also bid less for larger loans because they are also more likely to be refinanced earlier.*

The discount factor δ may vary across lenders and time because of reserve requirements and the need for liquidity. However, the Ginnie market is composed mostly of non-depositary institutions with similar liquidity needs so, in what follows, we assume that δ is the same across banks. More generally, banks put more weight on liquidity than MBS investors, especially non-banks. This generates gains from trade even in the presence of asymmetric information (Downing et al. 2009).

4.2 Security Valuations

At the securitization stage, banks must decide how to allocate loans across securities. For each Ginnie Mae loan that they have acquired, banks select the delivery month, coupon rate, and type of security. Banks typically select the earliest delivery date available (next calendar month accounting for a 1-2 weeks of delivery time), and so we focus on the latter two decisions.

4.2.1 Security customization

Given a portfolio of loans available with coupon c, each bank must decide how to optimally design mortgage-backed securities. We focus on two types of Ginnie Mae II securities: custom and multi-issuer.[11] A custom security corresponds to a "specified pool" (or spec-pool) of mortgages from a single issuer with similar attributes. The alternative is to pool loans in a multi-issuer pool, which is a fully diversified pool with loans produced by all issuers. In a given month, there is a single multi-issuer Ginnie Mae MBS for each coupon. In contrast, there are several hundred custom securities with the same coupon and issuing date, each including on average 25 loans from the same issuer.

Both types of securities are traded in over-the-counter (OTC) markets. The main distinction is that the uncertainty in the price of multi-issuer pools can be hedged by participating in a very liquid future (or TBA) market. We assume that, at the time of acquiring a loan, banks expect an upfront payment of $P_t^{tba}(c)$ for all loans pooled in multi-issuer securities. This price is common across all banks and reflects market expectations at the time of bidding of the value of loans that will be delivered in multi-issuer securities with coupon c.[12] Custom Ginnie Mae securities are traded in a more fragmented spec-pool market, and the prices for these securities are determined by soliciting bids from MBS investors. Importantly, investors can observe the characteristics of the loans in the spec pool, so prices reflect the composition of the pool and can vary across pools and issuers.

In principle, issuers can create spec-pools with several dimensions of differentiation (e.g. FICO $\geq$ 700, or New York state loans). However, in practice, the vast majority of spec-pools are segmented by loan size. The industry standard is to construct pools based on seven maximum loan size thresholds, $k \in \{85K, 110K, 125K, 150K, 175K, 200K, 225K, 250K, 250K+\}$. We use those thresholds to define seven types of spec pools: $(0, 85K]$,$(85K, 110K]$, $(110K, 125K]$, $(125K, 150K]$,$(150K, 175K]$, $(175K, 200K]$,$(200K, 225K]$, $(225K, 250K]$ and $250K+$. The last category includes loans with unrestricted loan sizes. A type k spec-pool consists mostly of loans whose size is between k and $k-1$, although an issuer may sometimes include a few smaller loans. The custom securitization rate is approximately 85% for loans less than $85K$ but declines sharply to roughly 10% for loans above $200K$. We will refer to loans below $200K$ as *small* loans and loans above $200K$ as *large* loans.

We assume that an issuer takes the market's segmentation on loan size as given when it makes securitization decisions. For each k, an issuer has Q^k loans in size bin $(k-1, k]$ with coupon c. These loans are essentially the same size, but are differentiated by other characteristics (e.g., FICO, state). We order them from highest to lowest based on their expected value to investors, which is typically equivalent to ordering loans based on their expected duration if investors have a sufficiently large discount rate relative to the coupon. Given this ordering, the issuer's optimization problem is to choose q to

$$\max\{qP^k(q;c) + (Q^k - q)P^{tba}(c) + \sum_{i=1}^{Q^k} E[M_i(T_i)(r - c - g)] - 1(q > 0)\kappa \tag{2}$$

[11] Ginnie Mae I securities represent a small share of the market and correspond to pools of loans with a common interest rate (except manufactured housing pools). Under Ginnie Mae II MBS, the interest rates may range from 25 to 75 basis points on mortgages in a pool.

[12] TBA trades are characterized by a coupon and a delivery date. We use a 2-month delivery window in our empirical analysis.

where κ is the cost of producing the security and $P^k(q;c)$ is the price that the investor pays for each loan in the security. If the spec-pool market is competitive, then this price is equal to the average expected revenues that an investor earns from the loans in the pool. More precisely, in a competitive equilibrium,

$$P^k(q;c) = \left[\frac{1}{q}\sum_{i=1}^{q} v_i(c)\right] \tag{3}$$

where

$$v_i(c) = E[cM_i(T_i) + \delta^{T_l} L_{T_i}]$$

is the expected value of loan i to investors. The first term in brackets is the coupon revenues from the loan and second term is the discounted repayment of the outstanding balance when the loan terminates. Assuming an interior solution. one can use equation (3) to show that the solution to bank $j's$ revenue maximization problem is choose q such that the expected value of the marginal loan $v^*(c)$ is equal to $P^{tba}(c)$. This leads to a threshold strategy: all loans with expected values above $v^*(c)$ are placed in a spec-pool k, and loans with expected values less than $v^*(c)$ are pooled in TBA-eligible securities. Note that the threshold does not vary across issuers or loan-size bins.

However, two other solutions are also possible. If there are too few loans that meet the threshold, then the gains from creating the spec security can be less than the costs In this case, the bank does not create a type k spec-security and allocates all of its type k loans to the TBA pool. The other possibility is that all of the issuer's loans meet the threshold. In this case, the issuer allocates all of its loans type k loans to the spec-pool and none to the TBA pool. As we document in the next section, both of these events are common.

Prediction 2. *If the price of custom securities is increasing in the average expected duration of loans included in the pool:*

a. *Loans pooled in custom securities have longer expected duration than loans pooled multi-issuer MBS.*

b. *The price of custom securities is higher than the TBA price for multi-issuer MBS.*

4.2.2 Coupon choice

Subject to restrictions on the servicing income imposed by the agency, issuers can choose the coupon that maximizes their expected revenue for loans with note rates that end in 0.25 or 0.75. In the finance literature, this problem is described as the "best execution" of an MBS. The choice of coupon depends on the security price $P(c)$, which is increasing in c, and the bank's beliefs about the loan duration (M_i). Banks face a tradeoff between increasing their expected revenue from payments by choosing a lower coupon or increasing upfront revenue from selling the security by increasing the coupon value. The rules for coupon choice are the same for each type of security.

Recall that Ginnie Mae charges a fixed g-fee of 0.06 p.p. for all banks and restricts the coupon choice such that the spread $(r - c)$ is between 0.25 and 0.75. This effectively imposes a minimum and maximum markup on banks. Since note rates are typically quoted on a 1/8 p.p. grid, it implies that for most loans, banks do not face a coupon choice. However, for note rates ending with 0.25 or 0.75, lenders can choose a high or low coupon security. For instance, a 4.25 note rate mortgage can be pooled in a 4% coupon security (high) or in a 3.5% coupon security (low). In the latter case, the bank earns a servicing income of $r - c - g = 0.69$ p.p., compared to 0.19 p.p. with the 4% coupon, but receives a lower upfront payment from selling the loan.

Define $\overline{M}_i = E[M_i(T_i)]$ as the expected duration of loan i. Since Ginnie does not collect other upfront payments from banks, the optimal coupon for eligible loans is a binary discrete choice problem:

$$c_i = \begin{cases} c_L & \text{If } \bar{M}_i > \frac{P(c_H)-P(c_L)}{(c_H-c_L)} \\ c_H & \text{If } \bar{M}_i < \frac{P(c_H)-P(c_L)}{(c_H-c_L)}. \end{cases} \tag{4}$$

where c_L and c_H denote the low and high coupons available for loan i, and $P(c_L) < P(c_H)$ are the associated security prices. Equation 4 implies that only high-duration loans are placed in the low-coupon option, with the threshold determined by the difference in the security prices relative to the difference in coupons.

Our third set of predictions is related to the adverse-selection of "high-coupon" securities. Specifically, banks will select a low coupon if they believe a loan is likely to survive for a long time, and a high coupon if they believe the loan will be pre-paid early. In practice, sellers' preferences for liquidity could also affect this decision. Sellers who value upfront cash payment (low δ) are more likely to choose the high-coupon option, which could mitigate the adverse selection problem.

4.3 Auction Valuations

Our goal is to use the bid data to test model predictions in the MBS market and to learn more about the information structure of the wholesale market. Each bidder's bid is based on its valuation of a loan, which in turn depends upon the price at which it can sell the loan in the MBS market and the value of the servicing rights. The former is known to the bidder but the latter is not since loan duration is uncertain. Thus, bidders need to form beliefs about loan duration, and those beliefs will reflect information that they have about loan characteristics and market conditions, and on how they use this information to forecast loan duration. Let Z_i denote the vector of loan and market characteristics observed by each bidder, and let S_{ij} denote a private signal that each bidder j has about the duration of loan i. This signal may reflect privately observed characteristics or how the bidder maps observed characteristics into an expectation of loan survival. Bidders are assumed to be risk-neutral. Then the willingness-to-pay of bidder j for a loan i is given by:

$$W_{ij} = P_j(c) + (r - c_i - g)E[M_i(T_{li})|Z_i, S_{ij}] - F \tag{5}$$

where $P_j(c)$ is the price at which bank j plans to sell the loan in the MBS market and F is the cost of acquiring loan i. This cost includes the financing cost of holding the loan on the bank's balance sheet until the delivery period, as well as other servicing and securitization transaction costs. For illustration purposes, we assume that this cost is common across loans and banks.

Our model of auction loan valuations has potentially two sources of heterogeneity, and both are private information. One is the private signal on loan duration that induces unobserved differences in bidder valuations of the servicing rights. The second is the resale price of a loan in the MBS market. Bidders may not have enough loans to create a custom security of type k and, even if they do, the price will depend on the composition of loans in the pool. Since a bidder's portfolio of loans is private information, this means that bidder's resale price of a loan is also private information.

4.3.1 Common-value: Mortgage servicing

To illustrate the empirical implications of private information on servicing revenue, we consider first the case of loans without a coupon-choice option and limited customization: $P(c_i, S_{ij}) = P^{tba}(c)$. For instance, loans with interest rates ending in digits other than 0.25 and 0.75 do not have a coupon choice, and loans above \$200 are unlikely to be pooled in custom securities. Further, consider auctions for loans with the same servicing fee. The value in this case is given by:

$$W_{ij} = s \times \bar{M}(S_{ij}) + P^{tba}(c) - F$$

where $s = r - c - g$ is a constant. This corresponds to the pure common-value environment. To see this, recall that the common component is associated with predictions about loan duration. Bidders may use different methods to process the data or have different incentives to invest in information acquisition, resulting in heterogeneous interpretations of the data. The value of the service income is a random variable from all bidders' points of view, and bidders' signals can be interpreted as noisy estimates of the true but known common value of the MSR. We allow the distribution of the signals to vary across bidders to reflect possible differences in the informativeness of the signal.

We assume that the auction has a pure strategy, monotone increasing equilibrium. Let $\beta_j(S_{ij})$ denote bidder j's equilibrium bid function expressed as a "pay-up" above the TBA price (or net bid):

$$\mathrm{B}_{ij}^{bulk} = \mathrm{Price}_{ij} - P^{tba}(c) = \beta_j(S_{ij}) \tag{6}$$

where Price_{ij} is the price quote offered by bank j. Since the resale value is known and common to all players, it is competed away at the auction (i.e. full pass-through).

We consider two empirical implications of the common-value assumption for the ex-post performance of loans. Consider the distribution of ex-post loan duration conditional on bids and observed characteristics of available to all bidders, denoted by Z_i. If banks are symmetrically informed, the realized repayment decision should be independent of bids. More formally, since the bid strategy is monotone in S_{ij}, the conditional expectation of T_i

$$E[T_i|S_{ij} = s; Z_i, \iota_i] = E[T_i|B_{ij}^{bulk} = b; Z_i, \iota_i]$$

in strictly increasing in b if bidders are privately informed about loan duration (independent otherwise). To allow for asymmetries in banks' strategies, we condition on the identity of the bidder ι_i when constructing the conditional expectation. In the pure common-value model, the ex-post measures of loan performance should be strictly increasing in b. In contrast, if bidders rely on the same estimates of $\bar{M}_i$, the distribution of bids is generated by a private value model (due to differences in F for instance), and the ex-post duration of loans is independent of bids conditional on Z_i. We refer to this test as the **monotonicity** test.

Another bid-level test is to condition expected survival on the bidder's bid b and the maximum rival bid. Ex-post measures of early prepayment should not vary with the maximum rival bid if the auction is PV, but be strictly increasing in this bid if the auction is CV. In fact, the slope should be higher than the slope for b since the maximum rival bid is a summary statistic of multiple signals and therefore more informative.

A related empirical implication uses information on winning bids only and is known as the **Winner's Curse** test. Bidder j wins the auction if it submits the highest bid. Let $B_{i,-j}^{bulk}$ denote the vector of bids submitted by j's rivals for loan i. Then, under the pure common-value assumption,

$$E[T_i|B_{ij}^{bulk} = b, \max\{B_{i,-j}^{bulk}\} < b; Z_i, \iota_i] < E[T_i|B_{ij}^{bulk} = b; Z_i, \iota_i].$$

This equation highlights the adverse-selection effect of winning in a pure common-value auction environment. Winning is "bad news" because it means that rivals have lower signals, which implies lower survival rates.

A challenge in implementing these two tests of the common-value hypothesis is that bids are truncated by the presence of the lock price. Since bidders are committed to buying loan i at price B_{ij}^{lock}, the platform does not allow firms to submit bulk bids below B_{ij}^{lock}. The bid that we observe is therefore:

$$\text{Bid}_{ij} = \max\{B_{ij}^{lock}, \beta_j(S_{ij})\}. \tag{7}$$

The lock price varies across lenders and loans, and reflects the value that bank j assigns to loans with similar attributes to i. Since B_{ij}^{lock} is unobserved, we assume that the bulk bid strategy takes the following form:

$$B_{ij}^{bulk} = \begin{cases} B_{ij}^{lock} + \beta_j(S_{ij}) & \text{If } S_{ij} \geq s_j^*(Z_i) \\ B_{ij}^{lock} & \text{If } S_{ij} < s_j^*(Z_i) \end{cases} \tag{8}$$

where $s_j^*(Z_i)$ is a threshold strategy that depends only on the identity of the bidder and loan characteristics. Under this assumption, the decision to submit a bulk price reflects the private signal observed at the time of bidding (e.g. additional loan characteristics or updated market condition). We can therefore adapt the monotonicity and Winner's Curse prediction using the "participation" signal associated with the decision of bidders to submit a bulk bid (rather than the actual bid signal).

The monotonicity property can be written as:

$$E[T_i|S_{ij} > s_j^*(Z_i); Z_i, \iota_i] = E[T_i|1(\text{Bulk}_{ij} = 1); Z_i, \iota_i] > E[T_i|Z_i, \iota_i$$

where $1(\text{Bulk}_{ij} = 1)$ is an indicator variable equal to one if bidder j submitted a bulk bid. Similarly, the Winner's Curse property implies that the conditional expectation of loan duration is increasing in the number of rival bidders submitting a bulk bid:

$$\begin{aligned} E[T_i|1(\text{Bulk}_{ij} = 1), \sum_{j' \neq j} 1(\text{Bulk}_{i,j'} = 1) = n; Z_i, \iota_i] > E[T_i|Z_i, \iota_i] \\ > E[T_i|1(\text{Bulk}_{ij} = 1), \sum_{j' \neq j} 1(\text{Bulk}_{i,j'} = 1) = n'; Z_i, \iota_i], \quad \forall n' < n. \end{aligned}$$

Prediction 3. *If banks are asymmetrically informed about loan duration, the ex-post duration of loans is monotonically increasing in bidders' own and rivals' participation decisions.*

4.3.2 Private-value: Security customization

The securitization process induces unobserved differences across banks in the resale value of loans in the secondary market. We focus in particular on private-value differences due to the ability to produce customized securities.

The model described above implies that the resale value depends on the quality of the pool in which a loan is securitized, as well as the decision by the bank to create a custom security or not. Both are likely to differ across banks due to differences in customization cost (κ) and differences in the volume of loans by size segments. For instance, if a bank does not have a large enough volume of small loans (less than \$85K), it will either pool loans less than \$85K in pools of size \$110K or \$125K, or place small loans in mutli-issuer pools. Both strategies will result in lower expected resale value. To the extent that the distribution of loan size in bank's j portfolio is private information, the customization process induces private-value differences in willingness-to-pay.

Since we do not observe the price of securities for loans sold in the spec-pool market, we cannot measure directly the effect of custom security prices on bids. For instance, small loans are more likely to be pooled in custom securities (higher resale value), but also generate higher expected cash-flows (higher MSR). Both factors generate a negative relationship between bids and loan size. To get around this problem, we can exploit the discreteness of the menu of spec-pool securities.

To the extent that loans of similar sizes are pooled together, the fact that MBS issuers use a discrete grid of loan size to define spec-pools implies that loans to the right of a given size cutoff are more likely pooled with larger loans, which lead to a lower expected resale value. This leads to a standard regression discontinuity design. The value of customization can be measured by comparing the expected bids for loans to the left and right of each size cutoffs:

$$\beta_k = \lim_{\epsilon \to 0} E[\text{Bid}_i | l_i = \bar{l}_k + \epsilon] - \lim_{\epsilon \to 0} E[\text{Bid}_i | l_i = \bar{l}_{k+1} - \epsilon] \tag{9}$$

The conditional mean difference β_k, is under the null hypothesis that the resale value of custom securities is independent of loan size, and positive if spec-pools with small loan sizes are more valuable than pools with of larger loans.

Prediction 4. *Bids are discontinuous in loan size around the spec-pool loan size cutoffs, and loans pooled in "small" spec-pools have higher prices.*

5 Empirical analysis

In this section, we test the four predictions of the loan valuation model described in the previous section. We start by testing Predictions 1a and 2 by analyzing the relationship between loan survival and securitization decisions. We then use bids to measure the value of customization using a regression-discontinuity approach (Prediction 2b and 4). Finally, we construct our test of common-value by analyzing the relationship between loan survival and bids.

5.1 Adverse selection at the securitization stage

5.1.1 Coupon choice

The adverse selection test for the coupon choice is based on the following empirical discrete choice model. Recall that banks select the low-coupon security if:

$$\bar{M}(S_{ij}) > \frac{P(c_H) - P(c_L)}{(c_H - c_L)} = 2\Delta P_i. \tag{10}$$

where $c_H - c_L = 0.5$ by construction.

We test the hypothesis that banks have private information about loan duration by estimating the following linear regression model of early pre-payment risk, $Y_i = 100(T_i > 12)$,

$$Y_i = Z_i\beta + \lambda 1(c_i = c_h) + \text{Fixed effects} + e_i. \tag{11}$$

where $\lambda = E[Y_i|Z_i, \bar{M}(S_{ij}) < 2\Delta P_i] - E[Y_i|Z_i, \bar{M}(S_{ij}) \geq 2\Delta P_i]$ measures the conditional expectation of the difference in survival probability for loans pooled in low vs high coupon securities. Since

$\bar{M}(\cdot)$ is a monotonic transformation of the expected duration, this difference is negative if banks have private information (beyond Z_i) about the distribution of T_i.

The ability of banks to select the coupon leads to adverse selection if the selection is driven by private information about loan duration. Selection can be *advantageous* if the factors entering the prices for high and low coupon securities are negatively correlated with duration. This can arise for instance if the security prices are set after the MBS pool is assembled, or if a bank's private benefit from assembling a diversified pool outweighs the benefit of placing a longer-duration loan in a low coupon. To minimize the importance of these factors, we focus on loans that are pooled in TBA-eligible multi-issuer securities. Recall that the TBA price reflects the expected composition of fully diversified pools, and is therefore unaffected by the decision to include or not an individual loan.

We estimate the model on loans that are eligible for a coupon choice (i.e., note rates ending in .75 or .25 digits) and that are sold in multi-issuer pools. To identify the source of selection, we estimate λ varying the set of controls Z_i. The price of the TBA MBS securities is a function of only aggregate information available to investors at the trading date (typically a few weeks before the pool issuance date). Our baseline specification therefore compares the ex-post performance of eligible loans that are securitized on the same date. We do so by controlling for note-rate/issuance date fixed-effects (i.e., $r \times t$).

Next we also condition on the identity of the issuer (i.e., aggregator or retail originator) selling the loan. We do so by augmenting the fixed-effects to compare loans securitized on the same date, with the same note rate, and the same issuer. Finally, we also control for characteristics of the loans that are observed by the lender/issuer, but not by the investor buying the TBA security. This includes the financial attributes of the borrower, contract characteristics, as well as the origination channel. Note that the investor also does not observe the identity of the issuer.

Panel A in Table 6 summarizes the results. The first row presents the estimate of λ across different sets of controls. The last two columns compare the estimate for loans acquired through the retail and wholesale channels. The results confirm that the coupon choice reflects issuer's beliefs about loan duration. Loans placed in high-coupons are 5% less likely to survive the first year, compared to loans sold on the same date and with the same note rate. This difference is reduced to 3.38% when we condition on loan attributes, and is further reduced to 1.8% when we condition on the identity of the issuer. We therefore conclude that more than half of the adverse-selection in the MBS market is due to observed characteristics of the loans. The identity of the issuer in particular is an important predictor of the coupon choice and loan survival. We find that a large fraction of issuers never select the low coupon option when it is available, and those lenders on average sell loans that are more likely to be pre-paid early.

Table 6: Adverse-selection test results

	(1)	(2)	(3)	(4)
Panel A: Coupon choice				
1(High coupon)	-5.09	-3.38	-1.82	-1.60
	(0.36)	(0.28)	(0.34)	(0.37)
Observations	2,870,163	2,867,143	2,844,005	2,831,919
Fraction High Coupon	0.87	0.87	0.87	0.87
Panel B: Security customization				
1(Multi-issuer pool)	-9.37	-4.61	-3.23	-2.97
	(0.24)	(0.17)	(0.19)	(0.19)
Observations	14,102,465	14,076,442	13,970,263	13,922,924
Fraction multi-issuer pool	0.84	0.84	0.84	0.84
Loan characteristics	no	yes	yes	yes
Fixed effects	$t \times r$	$t \times r$	$t \times r + I$	$t \times r \times I$
Mean dep. var.	89.6	89.6	89.6	89.6

Panel A: All loans without a coupon-choice option. **Panel B:** All loans above 200 in TBA pools with a coupon-choice option. Period: August 2013 to December 2019. Control variables: Loan size, State, FICO score, LTV, DTI, loan purpose, FHA/VA.

5.1.2 Security customization

In Panel B, we estimate a similar survival model to test for adverse selection in the TBA market by comparing the performance of loans placed in multi-issuer and custom pools. In particular, if issuers securitize high-performing loans in custom pools (because the security price is increasing in loan duration), loans placed in multi-issuer pools should have lower survival probability:

$$Y_i = Z_i\beta + \lambda 1\{\text{Multi-issuer}_i\} + \text{Fixed effects} + e_i. \tag{12}$$

The results strongly confirm this hypothesis. As before our baseline specification includes note-rate x issuance date fixed-effects. Without conditioning on other loan or issuer characteristics, we find that loans placed in multi-issuer pools are 9.37% more likely to be pre-paid within the first year. Given that the average 12-month survival probability during our sample period was 89%, this difference suggests a very severe adverse selection problem. Once again, conditioning on loan and issuer characteristics attenuates substantially this difference. In column (4), we find that the difference in survival between multi-issuer and custom securities is 2.87%.

Next, we measure the importance of adverse selection across different loans of different sizes. In particular we estimate equation (12) separately by loan size categories: less than 85, 85-110, 125-150, 150-175, and 175-200, 200-300, and greater than 300. Recall that the first 8 correspond to the

Figure 4: Adverse-selection and security customization by loan size bins

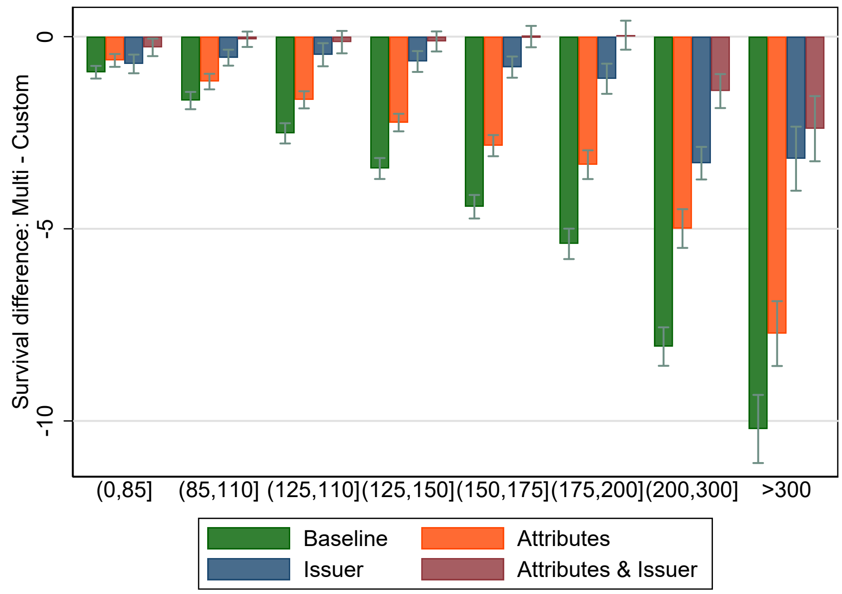

Period: August 2013 to December 2019. Control variables: FICO score, Loan size, single borrower, DTI, loan purpose, FHA/VA. Fixed-effects: (green) note-rate x month, (orange) note-rate x month x issuer, state. Standard errors clustered at the rate-date level.

standard cutoffs used to customize securities. Figure 4 plots the estimated coefficient $\hat{\lambda}$ obtained with and without loan and issuer characteristics. The "Baseline" estimates from specification (1), "Attributes" corresponds to specification (2), and "Attributes and issuer" corresponds to specification (4). The estimates labeled "Issuer" correspond to Speciation (2) using time x coupon and issuer fixed-effects. We interpret the fourth estimate as measuring selection on "unobservables", and the first three as combining selection on observables and unobservables. As before the difference between the second and third measures the importance of issuer unobserved heterogeneity. The figure illustrates two important results.

First, the scope for adverse-selection in multi-issuer pools is monotonically increasing with loan size. The survival probability difference between loans in custom and multi-issuer pools is 10% for loans above \$300K, compared to roughly 1% for loans less than \$85K. Therefore, not only are smaller loans less likely to be prepaid early than large loans, but smaller loans are also significantly more homogenous in terms of expected durations. This is consistent with the idea that borrowers with small loans have less to gain from refinancing when interest rate falls, which makes it more difficult for banks to predict the prepayment risk of individual borrowers. The repayment decisions of borrowers with small loans depend primarily on factors that are unobservable to banks, such as job relocation events to divorce. In contrast, borrowers with large loans have more to gain from refinancing their mortgages, and banks can more easily distinguish between loans with more or less prepayment risk.

The second takeaway relates to the type of information that banks use to allocate loans across securities. For loans below \$200K, the results show that the selection of custom securities is driven entirely by observable differences in loans and issuers characteristics. For large loans, roughly 1/3 of the difference is assigned to private information that issuers have about the quality of individual loans, and most of the remaining differences are due to issuer unobserved heterogeneity.

Although it might be tempting to conclude that private information about loan duration is not the main driver of selection. We cannot conclude this from Figure 4 because most issuers put all or none of their loans in custom pools (i.e. corner solutions), especially for loans less than 200K. The regression specifications that control for coupon x date x issuer fixed-effects therefore identifies λ_k by comparing a relatively small number of issuers.

To see this, Figure 5 illustrates the prevalence of corner solutions by loan size segment in the full sample of MBS issuers. For each segment, coupon and date, we calculate the fraction of loans in custom pools within an issuer porfolio (x-axis). This corresponds to $q^*_k(c)$ in the model. The figure plots the histogram of this variable weighted by the number of loans. Figure 5a plots the distribution of q^*_k for coupons loans less or equal to \$150K, and 5b plots the distribution for coupons for larger loans. The two figures plot the distribution of q^* only for coupons above or equal to the median, since most loans with low coupons are pooled in multi-issuer securities they are associated with low refinancing risk. We observe a clear bimodal distribution. For the smallest group, roughly 25% of loans are serviced by issuers who do not produce a small spec-pool security (i.e. $q^*_{85} = 0$), and slightly more than 30% are serviced by issuers who pooled 100% of their small loans in spec-pools. Therefore for slightly less than 50% of small loans the optimal pooling strategy is in the interior. For large loans, the fraction of securities with $q^*_c = 0$ is much higher, while the fraction of interior solutions is more stable. The "right corner" solution is less prevalent for loans above \$150K. The main takeaway from both figures is that the change in the probability of customization across loan sizes is due to extensive margin changes in the probability of producing or not custom pools.

To better understand the role of issuer unobserved heterogeneity, recall that an important driver of the decision to place a loan in custom security is the number and quality of loans with the same coupon and size. Since custom securities include loans with similar sizes and identical coupons, banks need a large volume of loans to justify the fixed cost of participating in the spec-pool market. Banks with insufficient volume will either decide to put all of their loans in multi-issuer pools (to save on the fixed cost), or place all of their loans in a spec-pool regardless of their underlying quality (to ensure a large enough pool). In contrast, banks with a large volume of loans can create spec-pools by screening individual loans, and as a result, assemble higher-quality pools.

Our results suggest that this type of asymmetry between banks is correlated with loan quality. A major source of adverse selection in this market is that issuers selling under-performing loans in

Figure 5: Fraction of loans in custom pools by size segment

(a) Loans $\leq$ \$150K (b) Loans $>$ \$150K

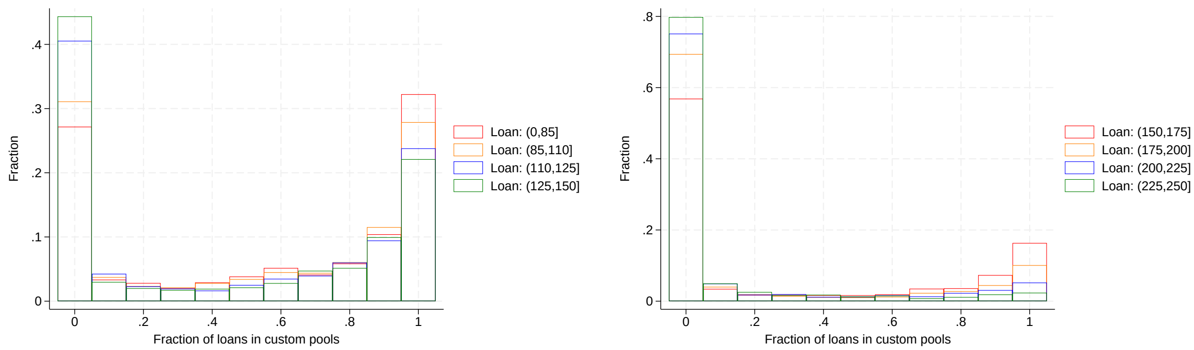

The x-axis measures the fraction of loans in a given size segment pooled in custom security within an issuer, coupon, and date portfolio. The height of the bars in the fraction of observations within in bin weighted by the number of loans in each portfolio. Width = 0.1.

the secondary market are more likely to use multi-issuer pool securities. Since these lenders tend to be smaller on average, this indicates that larger banks also tend to acquire higher-quality loans in the retail or wholesale markets.

5.1.3 Moral Hazard hypothesis

The previous two tests of adverse selection are based solely on the correlation between ex-post performance and the security choices. However, the insurance literature has long recognized that a positive correlation between prices and performance can be due either to adverse-selection or moral hazard. In our context, moral hazard can be caused by the ability of lenders to convince borrowers to repay their loans early. For instance, lenders may encourage a borrower to refinance its loan so that they can earn higher service income on the new loan.

We can test this hypothesis using the sample of loans that are not eligible for a coupon choice that are pooled in multi-issuer TBA-eligible securities. These are loans with service income of 0.375, 0.5 or 0.625 before applying the guarantee fee (i.e., $r-c$). Since we focus only on single-unit mortgages, everything else being equal, loans with higher service income are strictly more profitable for banks. We test the moral hazard hypothesis by estimating the following linear probability model:

$$Y_i = \lambda_1 1(r_i - c_i = .5) + \lambda_2 1(r_i - c_i = 0.625) + \theta r_i + Z_i\beta + \text{Fixed-effects} + u_i. \tag{13}$$

This equation describes the relationship between interest rates and loan attributes and the propensity of consumers to re-pay their loan. Since high-rate loans are more likely to be pre-paid default, we expect $\theta < 0$. The coefficients λ_1 and λ_2 allow for this relationship to be discontinuous around

Table 7: Moral-hazard test results

VARIABLES	(1)	(2)	(3)	(4)
Spread ($r-c$): 500 bbs	0.73		0.34	-0.17
	(0.15)		(0.15)	(0.13)
Spread ($r-c$): 625 bbs	1.12		0.60	-0.0033
	(0.15)		(0.16)	(0.13)
Observations	6,567,611	6,518,151	6,556,655	6,507,211
R-squared	0.107	0.129	0.136	0.168
Loan characteristics	Rate and Loan	Rate and Loan	all	all
Issuer FE	no	yes	no	yes
Latest issuance date	2021/2	2021/2	2021/2	2021/2
Fixed effects	Month	Month	Month	Month
Mean dep. var.	88.2	88.2	88.2	88.2

the digits determining the coupon choice. We infer that lenders engage in strategic pre-payment if $\lambda_2 > \lambda_1 > 0$.

Table 7 summarizes the results. The top panel measures the effect of service income on performance for retail loans. Without conditioning on the identity of the lender and loan characteristics, we find that loans with higher service income are more likely to survive (consistent with the moral-hazard hypothesis). As with the selection test, this difference is largely due to unobserved heterogeneity across lenders. Once we condition on the identity of sellers, the highest margin loans are more likely to survive (relative to loans with the lowest service income group), but the difference is zero for loans placed in 50 bps coupons. Once we also condition on loan attributes, this difference is eliminated. We therefore cannot reject the hypothesis that banks engage is strategic prepayment based on observed attributes of the loans. However, this effect is economically much smaller than the selection effect, and could be driven by selection of which loans are securitized vs kept on the banks' balanced sheet. Therefore, we conclude that the correlation between security characteristics and loan performance is most likely due to adverse selection based on observed and unobserved loan attributes.

5.2 The value of security customization

In the previous section, we establish that lenders select the characteristics of securities, coupon and pool attributes, using private information about loan duration. The previous literature has documented that custom securities are transacted at a premium over multi-issuer securities (cite papers here). In this section, use our auction data to measure the effect of customization on

Figure 6: Distribution of loan size across spec-pools

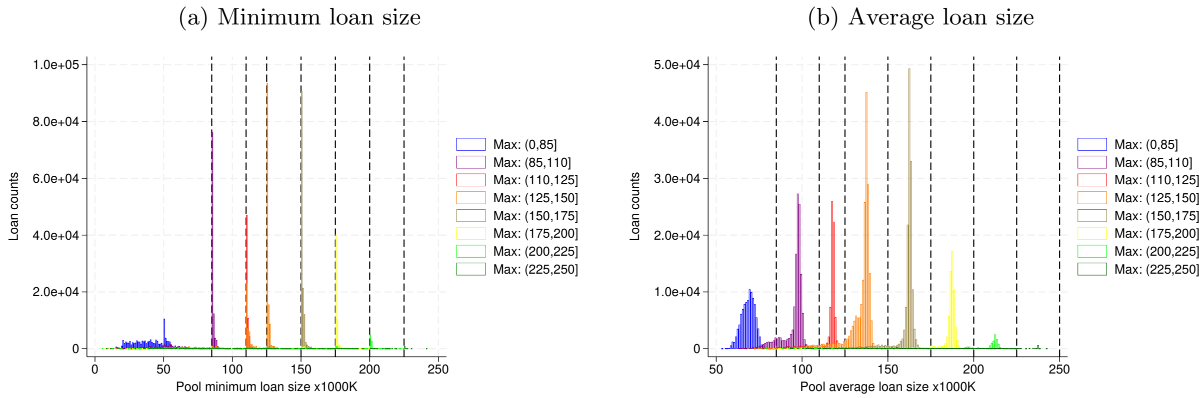

the value of loans in the wholesale market. Since custom securities are associated with higher transaction costs, it is an empirical question whether or not customization leads to higher loan prices. Moreover, since banks have heterogeneous access to the spec-pool market, the pass-through of spec-pool pay-up is not necessarily complete.

As we discussed above, we cannot measure directly the pas-through of custom securities prices on bids since we do not observe the transaction of the security associated with each loan being auctioned. To identify the value of customization we instead rely on discontinuities in the probability of customization across loan sizes.

To describe the pooling strategy of MBS issuers, we need to characterize spec-pool securities. Recall that spec-pool security is defined by a coupon, issuer, and date (month). Large issuers produce several hundred custom securities with different attributes, and the CUSIP of each security does not specify a unique market segment. In the Ginnie market, spec-pools are mostly differentiated by loan size, and we use the maximum loan size included in each security (or CUSIP) to define 9 types of spec-pools: $(0, 85K]$, $(85K, 110K]$, $(110K, 125K]$, $(125K, 150K]$, $(150K, 175K]$, $(175K, 200K]$, $(200K, 225K]$,$(225K, 250K]$, and $205K+$. The last category includes securities with larger loans that might be differentiated along other dimensions (e.g. FICO or agency). In our sample, 15% of loans in custom pools belong to this "other" category. The two categories above 200K are also much smaller in terms of volume, and we omit those from part of our analysis.

Figure 6 summarizes the loan size distribution within custom securities of different types. Each histogram is weighted by the number of loans in each pool, so the y-axis is the number of loans in each size and security-type bin. The different colors correspond to the types of spec-pools. The "blue" securities are pools with maximum loan size less than 85K. The most common category includes the pools with max loan size between 125 and 150 (22%).

Figure 7: Fraction of loans in multi-issuer pools by loan size

(a) Full sample: 2013-2021 (b) Auction sample: 2018-2019

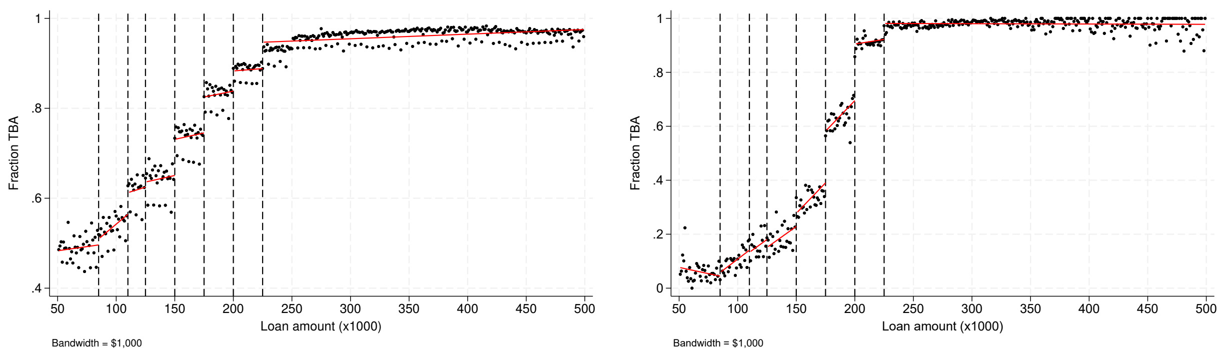

Figure 6 plots the distribution of minimum loans in each category (excluding the "other pool" category). If all spec pools with maximum loan size less than \$250K were perfectly segmented the distribution of the minimum would correspond to the maximum of the next lowest category. For instance, the yellow category shows that the vast majority of pools with max loan size between 175 and 200 have a minimum at 175, which means that the typical range is 25K. This corresponds to a pool exclusively with loans between 175 and 200. Although pools are not perfectly segmented, we see that most custom securities have a loan size range of 25K or 15K (except for the \$85K pool). The segmentation is less clear for the 125K cutoff. This is because some issuers a single pool category from 110-150, while others create two types of pools 110-125 and 125-150.

Figure 6 plots the distribution in the average loan size per security. As we can see, securities are highly differentiated by size, and the average size distributions reveal very little overlap across categories. Importantly, securities are almost perfectly differentiated along the average loan size dimension, despite the fact that the distribution of the minimum loan size per pool exhibits a long left tail. Therefore, as first-order approximation, the data suggest that firms bundle loans with loan sizes on a 25K grid (15K for loans between 110 and 125), consistent with our modeling assumption that loans "fit" in a single spec-pool. This suggests that spec-pools containing loans with similar sizes are more valuable than diversified pools with the same average loan size.

An implication of the high level of segmentation among custom securities is that the expected resale value of mortgages in the secondary market is discontinuous in loan size. Our model of security pooling predicts that a discontinuity can arise from two sources. First, loans on the right side of a customization threshold can be less likely to be sold in the spec-pool market. Second, the average size of pools in which a loan is securitized is discontinuous in loan size.

To see why the probability of customization can be discontinuous in loan size, recall that our model of pool selection under perfect competition implies a common cutoff for all pool types, defined in terms of expected loan value: $v(c, S_i) > P^{tba}(c)$ for all security types k. This implies that when

Figure 8: Security average loan size by loan amount

(a) Full sample: 2013-2021

(b) Auction sample: 2018-2019

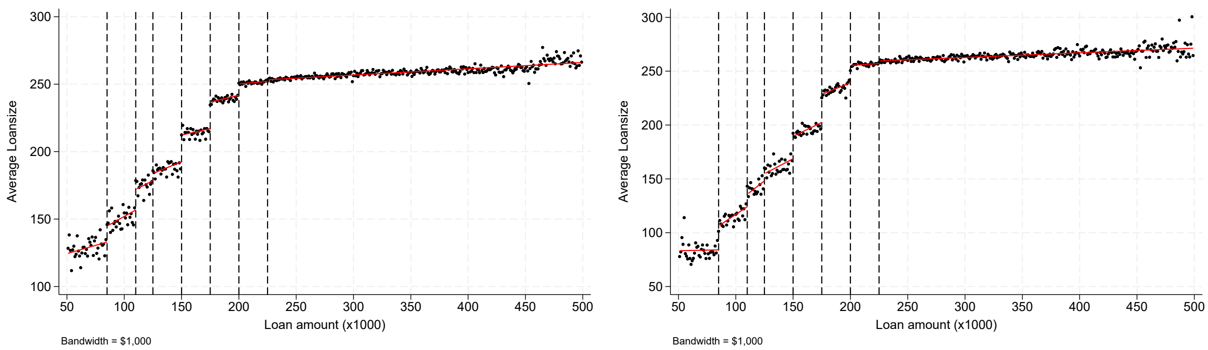

issuers select a fraction between zero and one to pool in spec-pool k for all k (i.e. $q_k^* \in (0,1)$), the probability customization should be a smooth function of loan size in the $(0, 250]$ range, and should not exhibit any discontinuity. Moreover, the probability of placing a loan in a TBA-eligible pool should decrease in loan size because $v(c, S)$ is decreasing in l. When the optimal pooling strategy exhibits a corner solution (i.e. $q_k^* = 0$ or $q_k^* = 1$), the probability of placing a loan in a custom pool can be discontinuous across loan size categories. This is especially true if the quantity of loans with a value above $P^{tba}(c)$ is too small to justify paying the fixed cost of securitization.

Figures **??** plot the fraction of loans in multi-issuer pools by loan size ($1K bin). Figure 7b is consistent with the presence of an interior solution for loans less than $150K for most banks in our auction sample. The probability is monotonically increasing in loan size in this range and exhibits a small discontinuous jump at $150K, and more pronounced increases at $175K, $200K, and $225K. A likely interpretation is that fewer loans in the larger size categories have values above the threshold, and therefore the securitization cost is more likely to exceed the value of creating large-loan spec pools leading to corner solution (e.g. $q_{175}^* = 0$). Note that we observe more discontinuities and an overall higher fraction of loans in TBA-eligible pool in the full MBS sample (Figure 7a). This is partly because the period 2018-2019 corresponds to an episode of declining interest rates during which the gain for customization is larger. The other factor is that the auction sample includes only large MBS issuers who are more likely to produce custom securities across loan size and coupon segments.

The second source of discontinuity in values is due to the segmentation of custom securities. As we saw earlier, most securities of type k include loans of sizes ranging from $\bar{l}_{k-1}$ and $\bar{l}_k$. Therefore the type of custom securities changes discretely when loans cross the custom size cutoffs. Figures 8a and 8b illustrate this point for the full and auction samples, respectively. Both figures highlight the presence of discontinuities in the average loan per security across the 8 cutoffs. The discontinuity is less marked for the 125 category because banks use different pooling strategies for loans between 110 and 150.

Table 8a quantifies the size of the discontinuities in pool characteristics (average loan size and multi vs custom pool) for each of the loan size cutoffs. Each entry corresponds to the regression coefficient β_k obtained by estimating the following semi-parametric regression:

$$Y_i = \beta_k 1(l_i > l_k) + g_k(l_i) + \epsilon_i, \quad \text{If } l_i \in (l_{k-1}, l_{k+2}). \tag{14}$$

We estimate $g_k(l_i)$ using a triangular Kernel.[13] The results are consistent with the figures above. There exist large discontinuities in the average loan size across all size cutoffs, roughly equal to the range of each security segment (between \$15K and \$30K). The second panel also confirms that we fail to detect discontinuities in the probability of placing a loan in a TBA-eligible pool for loans less than 150K. The estimated jump is very large (24%) in the \$175K group.

Next we look at the effect of security customization on loan prices using the auction data. Figure 9 plots the relationship between winning bids and loan size. As before, each dot measures the average highest-bid received across auctions within the same loan size bin (\$1K width). The figure reveals a sharp declining relationship between bids and loan size. This is because smaller loans are associated with larger servicing revenue and higher spec-pool prices. Similar to Figures **??** and 8, the data reveals discrete jumps in this relationship at the 6 most common spec-pool cutoffs. We use the size of these jumps to measure the value of customization.

To interpret these differences as causal, our main identification assumption is that loan attributes (observed and unobserved) affecting the willingness-to-pay of banks evolve smoothly with loan size. In Table 8b, we evaluate this assumption by testing the null hypothesis that observed loan attributes are not affected by the discreteness of the customization process. Each entry corresponds to an estimate of β_k for key borrower characteristics. The results are consistent with our main identification assumption. The characteristics of securities are the only observable covariates that change discontinuously with loan size. Perhaps surprisingly, even the interest rate is continuous around the loan size cutoffs. In other words, consumers with loans that can be securitized in more valuable spec-pools do not pay lower interest rates, suggesting a zero pass-through of spec-pool prices in the retail market.

Table 9 summarizes our main regression discontinuity results. The reduced-form (or RF) estimates correspond to the estimate of β_k from equation (14) using either the highest bid or the winning bid as dependent variable. Our regression corresponds to a fuzzy regression-discontinuity design since the assignment of loans to spec-pools (treatment) is not deterministic. To measure

[13] We consider two sub-samples: the sample of winning bids accepted by the seller, and the sample of highest bid. In about 10% of auctions, the seller does not select the highest bidder. This happens mostly when the highest bid is a lock price with delivery methods that differ across buyers.

Table 8: Evidence of discontinuity in securities and loan characteristics across common spec-pools cutoffs

(a) First-stage test results

	85k	110k	125k	150k	175k
Panel A: Average loan size					
Winning Bid	22.30	17.16	14.15	27.07	30.34
	(3.484)	(3.559)	(3.401)	(2.683)	(2.143)
Highest Bid	21.87	19.38	13.37	27.70	30.00
	(3.294)	(3.488)	(3.142)	(2.299)	(2.196)
Panel B: Multi-issuer pool (TBA)					
Winning Bid	-0.000833	0.0163	-0.00775	0.0832	0.237
	(0.0161)	(0.0170)	(0.0192)	(0.0214)	(0.0228)
Highest Bid	0.00102	0.0199	-0.00958	0.0949	0.240
	(0.0152)	(0.0166)	(0.0186)	(0.0202)	(0.0222)

(b) Balance test results

	85k	110k	125k	150k	175k
Income	-0.0783	-0.246	0.0959	-0.00836	-0.0473
	(0.118)	(0.115)	(0.103)	(0.0744)	(0.0659)
FICO	3.633	2.150	1.266	0.850	-0.592
	(2.548)	(2.639)	(1.776)	(1.385)	(1.343)
LTV	0.0389	0.343	0.165	-0.456	-0.231
	(0.430)	(0.448)	(0.285)	(0.265)	(0.196)
Note Rate	-0.00112	0.00502	0.0390	0.00110	0.00712
	(0.0353)	(0.0326)	(0.0279)	(0.0230)	(0.0204)
1(FHA Loan)	-0.0254	-0.0145	0.00820	-0.000974	-0.0113
	(0.0218)	(0.0204)	(0.0207)	(0.0129)	(0.0111)

Figure 9: Evidence of discontinuity in the relationship between highest bid per auctions and loan size

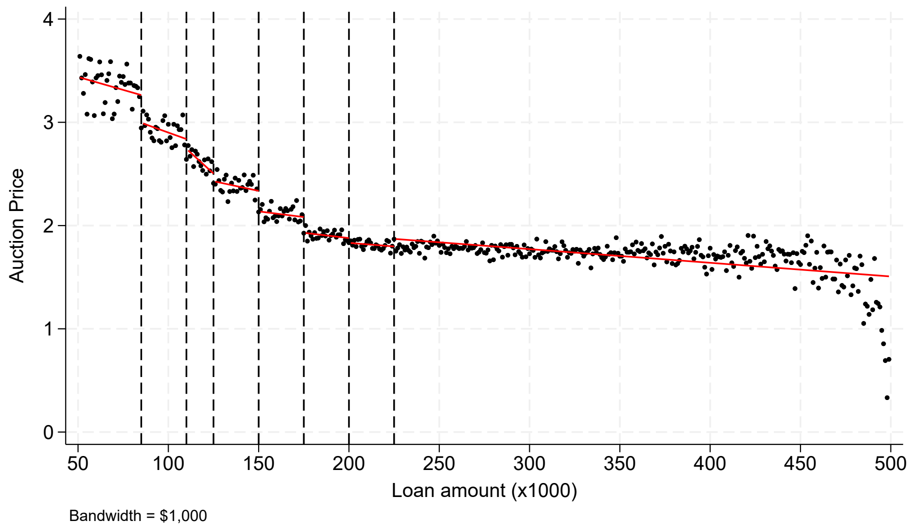

Table 9: Regression discontinuity results: Winning and Highest bids

	85k		110k		125k		150k		175k	
	RF	IV	RF	IV	RF	IV	RF	IV	RF	IV
Winning bid	-0.341	-0.0151	-0.224	-0.0133	-0.105	-0.00842	-0.255	-0.00956	-0.189	-0.00585
	(0.0709)	(0.00308)	(0.0570)	(0.00336)	(0.0441)	(0.00330)	(0.0346)	(0.00122)	(0.0251)	(0.000874)
Highest bid	-0.301	-0.0153	-0.221	-0.0121	-0.0980	-0.00833	-0.255	-0.00943	-0.166	-0.00517
	(0.0777)	(0.00283)	(0.0532)	(0.00284)	(0.0401)	(0.00333)	(0.0343)	(0.00113)	(0.0260)	(0.000909)

the treatment effect of customization we also present, in the IV columns, the results of a 2SLS regression where we use the size category indicator variable as an instrument for the security average loan size.

The reduced-form estimates confirm the presence of sizable discontinuities in the average winning bid around each of the spec-pool size cutoffs. For instance, we estimate that loans located to the right of the 85K cutoff receive bids that are $0.34 lower than loans located to the left. This corresponds to an 11% increase in the above pay-up above the TBA price for loans above the $85K cutoff (see Figure 9). The other cutoffs are associated with a similar decrease in loan prices ranging from -0.10 to -0.25.

The IV columns rescale the reduced-form estimates by the predicted change in average loan size. The point estimate corresponds to the causal effect of a $1,000 increase in the pool average loan size on the highest (or winning) bid. We can use this estimate to measure the effect of placing a loan in a pool with an average size \$25 larger (the most common range). In the largest spec-pool category, this would lead to a \$0.146 decrease in price. In the smallest spec-pool, this would lead to a \$0.377 decrease in price. We therefore find that the marginal value of customization is decreasing with loan size. Assuming that the difference is zero for loans above \$200K, this implies that under perfect segmentation, loans smaller than \$85K are worth \$1.21 more than loans above \$200 because of the ability of an issuer to create custom securities. The cumulative difference is 0.85, 0.50, 038, and 0.146 respectively for the other loan size categories.

It is difficult to interpret the magnitude of these differences in terms of pass-through since we do not have data on the value of custom securities in the spec-pool market and on the transaction cost of producing and selling custom securities. Our estimates represent a lower bound on the increase in willingness to pay due to customization. That being said, the results confirm that the ability to produce custom securities increases the willingness-to-pay of banks acquiring loans in the wholesale market. For small loans, we estimate that the customization option alone nearly doubles the average price that originators receive at auctions (relative to loans above \$200).

Do buyers value securitization equally? To answer this question we analyze the dispersion of bids within auctions across the different loan size bins. Figure 10a plots the average inter-quartile range of bids within auctions across a \$1,000 loan amount grid. Figure 10b focuses instead of the difference between the first and second highest bid. In both cases, we observe a steep decline in bid dispersion with loan size, and a stable distribution for loans that are unlikely to be pooled in custom securities. The winning average winning margin is twice as big for loans less than 100K compared to loans above 200. Dispersion across all bids decreases by a factor of three. This is consistent with the findings above that banks use different customization strategies, and as a result have different willingness-to-pay for loans that can enter spec-pools. This is an important source of market-power, which likely attenuates the pass-through of spec-pool prices on wholesale prices.

5.3 Adverse-selection at the loan acquisition stage

The last prediction of our model of loan valuation in that banks' willingness-to-pay for loans include a common-value component associated with the valuation of servicing rights. If banks observed different signals over the expected duration of loans, the winner of the auction has the most optimistic estimate of the mortgage servicing rights. This leads to an adverse-selection problem.

The source of Winner's Curse is that buyers value the mortgage servicing rights (MSR); the common-value component of the willingness-to-pay. To illustrate the importance of the MSR in explaining the dispersion of bids we exploit the fact that the profitability of mortgage servicing is increasing in the difference between the interest rate and coupon:

$$WTP_{ij} = (r_i - c_i - g)\bar{M}(S_{ij}) + P(c). \tag{15}$$

Figure 10: Dispersion of bids within auctions across loan size bins

(a) Within auction interquartile range

(b) Winning margin

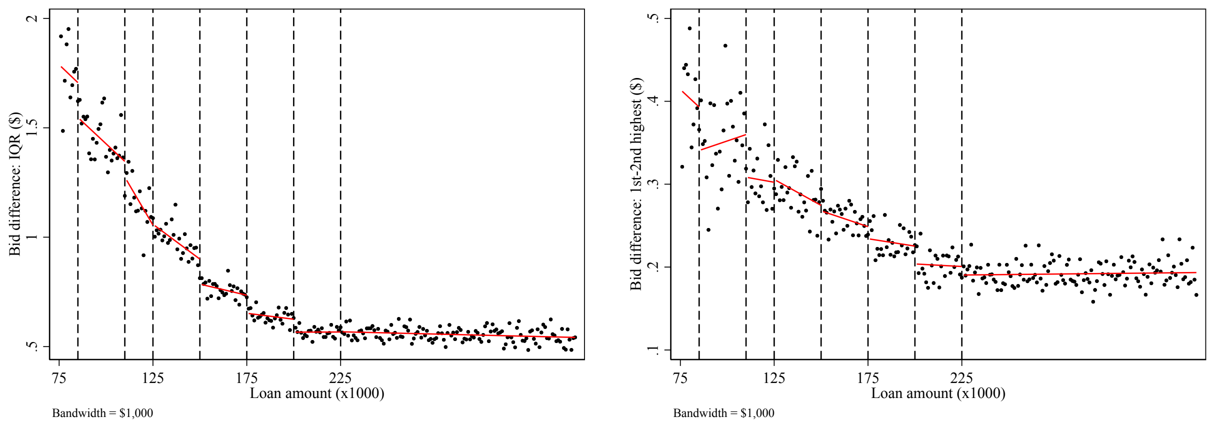

Variation in r_i across loans within the same coupon security increases the value of the MSR. To measure the effect of the service income on bids we estimate the following linear regression model:

$$\text{Highest Bid}_i - P_t^{tba}(c_i) = \sum_{r=2.75}^{6} \sim_{c\in\{c_L,c_H\}} 1(c_i = c) \times 1(r_i = r)\beta_{r,c} + Z_i\gamma + \text{Fixed-effects} + \epsilon_{ij} \quad (16)$$

The fixed-effects include coupon $\times$ auction-date, state and seller. The control variables include borrower and loan characteristics.

Figure 11 plots the predicted bid evaluated at the average of the fixed-effects and loan characteristics. Each dot correspond to the average net-bid for each values of the interest rates. For interest rates ending in .25 and .75, the two dots correspond to the average bid for banks selecting the low-coupon (top) and high-coupon (bottom). Loans with an interest rates of 3% are pooled in a 2.5% coupon, and earn a service income of 44 bps. In contrast, loans with an interest rate of 2.75 pooled the same 2.5% coupon security earn a service income of only 19 bps, while loans with an interest rate of 3.25 pooled in a low coupon security earn 69 bps.

The figure shows an almost linear relationship between bids and interest rate within a given coupon. If the highest bid reflects the WTP of the winner, the slope of this relationship measures the expected service multiple for loans placed in the same coupon-securities. The slope of the relationship indicates the service multiple is 3.58 for loans in the 2.5 coupon, but only 1.66 for the highest coupon. The first takeaway is that variation in service income across loans lead to large differences in bids, consistent with our hypothesis value the MSR. The second takeaway is that the service multiple is decreasing in the level of the interest rate (holding fix the service income). Recall that the service multiple is a measure of expected loan duration. Bids therefore reflect the negative relationship between loan duration and expected profit from mortgage servicing.

Figure 11: Relationship between bids and service income

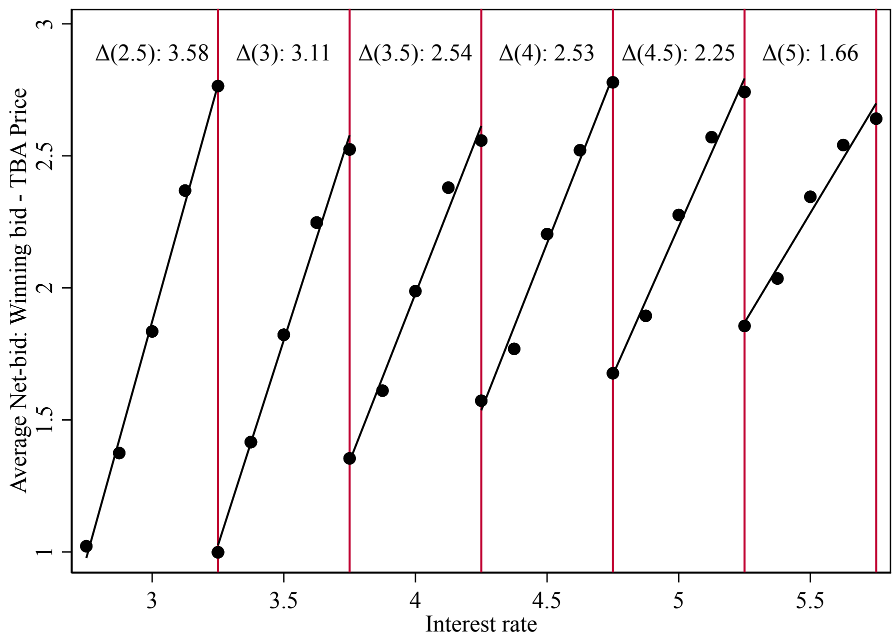

Table 10: Effect of service income on bid dispersion

VARIABLES	(1) IQR	(2) IQR	(3) Winning margin	(4) Winning margin
(r-c) = 50 bps	0.082*		-0.021	
	(0.018)		(0.012)	
(r-c) = 62.5 bps	0.18*		-0.039	
	(0.036)		(0.024)	
N. bidders (log)	-0.32*	-0.32*	-0.20*	-0.21*
	(0.017)	(0.014)	(0.0090)	(0.0079)
Coupon choice eligible		0.068*		0.023*
		(0.0044)		(0.0030)
Observations	36,689	50,183	36,689	50,183
R-squared	0.363	0.370	0.266	0.259
Sample	L>200	L>200	L>200	L>200
RMSE	0.29	0.30	0.19	0.20

Robust standard errors in parentheses

* $p<0.01$

Do banks value MSR equally? To get at this question we study the dispersion of bids within auctions with varying levels of service income. If banks use different service multiples $\bar{M}$ when placing their bids, equation 15 suggest that bids should be more dispersed in auctions with larger service income. We test this hypothesis focussing on auctions for loans above \$200K for which bid differences mostly reflect the MSR. We estimate the following regression:

$$IQR_i = \beta_{50}1((r_i - c) = .50) + \beta_{62.5}1((r_i - c) = .625) + Z_i\gamma + \text{Fixed-effects} + \epsilon_{ij} \tag{17}$$

were the fixed-effects includes coupon x auction month, issuer and state. We drop loans with a coupon choice option, and the reference category includes loans with service income of 37.5 bps. Table 10 summarize the estimates. An increase in service income of 12.5 bps is associated with a 0.082 increase in the inter-quartile range of bids. This is quite large relative to the standard-deviation of the residual of the regression (0.29). This confirms a large fraction of the dispersion in bids is due to differences in the value of the MSR across bidders. This can be due to differences in beliefs about loan duration, or differences in the discount factor.

Column (3) estimates the same regression using the difference between the top two bids (winning margin). There we find that the dispersion is independent of the service income. This suggests that the gap between the first two bids is not driven by differences in beliefs about $\bar{M}$, but other differences in WTP such as securitization cost or resale value.

To test the hypothesis that firms have different beliefs about loan duration, following Henricks et al. (2003), we test the Winner's Curse hypothesis by measuring the correlation between bids and participation in the OB auctions and ex-post loan performance, measured again using the 12-month survival probability of FHA and VA loans. We drop auctions for loans for which banks have a coupon choice, which reduces the importance of unobserved heterogeneity in resale values.

We test the Winner's Curse hypothesis by measuring the correlation between the fraction of bidders submitting bulk bids (positive signal of loan value) and ex-post loan duration. We test the common-value hypothesis by estimating the following linear survival model:

$$Y_i = \lambda_1 \text{Fraction bulk}_i + \lambda_2(B_i^{(1)} - B_i^{(2)}) + \text{Fixed-effects} + Z_i\beta + \epsilon_i \tag{18}$$

where $Y_i = 100 \times 1(T_i > 12)$, Fraction bulk$_i$ is the fraction of invited bidders who submit a bulk bid in auction i, $(B_i^{(1)} - B_i^{(2)})$ is the difference between the highest and second highest id, and Z_i includes loan and originator characteristics. We use the same set of control variables included in the adverse-selection test, augmented with information regarding the auction date (instead of MBS issuance month), and originator and borrower state fixed-effects. When controlling for loan characteristics, we also include auction date $\times$ note rate fixed-effects, which absorbs the common believes that banks have about the resale price in the secondary market.

Recall that bidders are privately informed about loan duration if the decision to submit a bulk

Table 11: Winner's Curse Test: Auction level analysis

VARIABLES	(1)	(2)	(3)	(4)
Fraction bulk bids	3.168*	8.923*	0.337	4.348*
	(1.070)	(4.155)	(1.517)	(1.446)
Difference: First and second bid	-3.407*	-6.865*	-0.973	-5.183*
	(0.704)	(3.108)	(0.795)	(1.189)
Observations	61,262	3,469	21,083	36,676
R-squared	0.156	0.179	0.102	0.158
Sample	All loans	Loan<200 & Low coupon	Loan<200 & High coupon	Loan>200
Mean dependent variable	80.57	87.02	89.19	75.01
Mean bid difference	0.231	0.207	0.292	0.198
Mean fraction bulk	0.777	0.750	0.817	0.757

bid is positively correlated with survival. This correspond to the monotonicity hypothesis. The winner suffers from an adverse-selection problem if it wins the auction by over estimating the value of the serving rights relative to other banks. To capture this, we also include the winning margin in the regression to measure the correlation between the size disagreement in valuations and loan performance.

Table 11 summarizes the key results. We estimate the regression separately for loans above and below \$200K. We expect the participation strategy of bidders to be increasing in the beliefs loan duration only for loans with an expected resale price equal to the TBA price. Column (4) is consistent with this interpretation. Auctions with 100% bulk bids (full participation) are 4.34% more likely to survive one year. On average, the participation rate is 75%, and so a .25 point increase in participation is associated an increase in survival slightly above 1%. This is comparable to the magnitude our estimates of the effect of selection on unobservables in the adverse-selection section.

The correlation between winning margin and survival is also negative as expected. The standard deviation of the winning margin is .3. Therefore, a one standard-deviation increase in the size of the disagreement between the first and second highest bidders is associated with 1.55 decrease in short-term survival. This confirms that the magnitude of the adverse-selection problem is larger when firms receive very different signals of loan quality.

Columns (2) and (3) repeat the same exercise for loans that can be sold in the spec-pool market. Column (2) estimate the model for small loans with relatively low interest rates; defined as having an MBS coupon below the median in a given issuance month. Recall that banks are unlikely to produce low-coupon custom securities, because those loans are not expose to important prepayment risk. In contrast, in Columns (3), the majority of small loans with coupons above or equal

Table 12: Winner's Curse Test: Bidder level analysis

VARIABLES	(1)	(2)	(3)
1(Bulk bid)	0.770*	0.588*	0.530*
	(0.210)	(0.157)	(0.163)
Fraction bulk bids (rivals)		4.810*	5.164*
		(1.349)	(1.354)
Net bid			-0.117
			(0.164)
Max rival bid			-3.548*
			(0.838)
Observations	364,520	364,487	364,487
R-squared	0.153	0.154	0.154
Sample	L>200	L>200	L>200
Mean dependent variable	74.63	74.63	74.63
Mean 1(Bulk bid)	0.761	0.761	0.761
Mean Fraction rival bulk	0.761	0.761	0.761
Mean Net bid	1.429	1.429	1.429
Mean Max rival bid	1.974	1.974	1.974

to the median are sold in the spec-pool market. For the first category of loans, difference in bids and participation most likely reflect differences in MSR, while the second category of loans, differences reflect both private information about duration (common value) and resale value (private value).

The results clearly show that auctions for small loans with low coupons suffer from a Winner's Curse problem. The participation fraction and winning margin are uncorrelated with the short-term survival of small loans with high coupons, but are strongly predictive of survival for small loans with low coupons.

A drawback of the previous analysis is that correlation between participation and survival could be caused by unobserved auction heterogeneity, rather than private signals about loan duration. Although we cannot fully rule out this possibility, we can estimate the previous regression at the bidder level to measure the information content of the "participation signal" of bidders, conditional on the participation of rival bids. This corresponds to our second Winner's Curse prediction. We implement this test by estimating the following regression on the sample of loans above $200K:

$$Y_i = \lambda_1 1(\text{Bulk})_{ij} + \lambda_2 \text{Fraction rival bulk}_{ij} + \text{Fixed-effects} + Z_i\beta + \epsilon_i \tag{19}$$

We use the same set of controls as before, plus bidder fixed-effect. The monotonicity hypothesis corresponds to $\lambda_1 > 0$, while the Winner's Curse hypothesis corresponds to $\lambda_2 > 0$. Table 12 summarizes the results. In column (1), we find that the decision to submit a bulk bid is positively correlated with survival. Loans 0.75% more likely to survive one year when the average bidder submits a bulk bid. This correlation weakens somewhat in Column (2) when we control for the fraction of rival bidders submitting a bulk bid (0.588), but remains statistically significant. The estimate of β_2 is also positive and precisely estimated. Together these two sets of results confirm our hypothesis that the auctions have a strong common-value component, and that bidders have private information about loan duration when deciding to participate in the auctions.

References

Agarwal, S., Chang, Y. & Abdullah, Y. (2012), 'Adverse selection in mortgage securitization', *Journal of Financial Economics* **105**, 640–660.

Allen, J., Clark, R. & Houde, J.-F. (2019), 'Market power and search frictions in negotiated-price markets', *Journal of Political Economy* **127**(4), 1550–1598.

andreas, F., Goodman, L., Lucca, D., Madar, L., Molley, L. & Willen, P. (2013), 'The rising gap between primary and secondary mortgage rates', *FRBNY Economic Policy Review* pp. 17–39.

Buchak, G., Matvos, G., Piskorski, T. & Seru, A. (2020), Beyond the balance sheet model of banking: Implications for bank regulation and monetary policy. NBER working paper 25149.

Cassola, N., Hortaçsu, A. & Kastl, J. (2013), 'The 2007 subprime market crisis in the EURO area through the lens of ECB repo auctions', *Econometrica* **81**(4), pp. 1309–1345.

Chiappori, P.-A. & Salanie, B. (2000), 'Testing for asymmetric information in insurance markets', *Journal of Political Economy* **108**(1).

Downing, C., Jaffee, D. & Wallace, N. (2009), 'Is the market for mortgage-backed securities a market for lemons?', *Review of Financial Studies* **22**(7).

Einav, L., Finkelstein, A. & Levin, J. (2010), 'Beyond testing: Empirical models of insurance markets', *Annual Review of Economics* **2**, pp.311–36.

Gao, P., Schultz, P. & Song, Z. (2017), 'Liquidity in a market for unique assets: Specified pool and to-be-announced trading in the mortgage-backed securities market', *Journal of Finance* **LXXII**(3).

Haile, P. (2002), 'Auctions with resale markets: An application to u.s. forest service timber sales', *American Economic Review* **91**(3).

Hendricks, K., Pinkse, J. & Porter, R. H. (2003), 'Empirical implications of equilibrium bidding in first-price, symmetric, common-value auctions', *Review of Economic Studies* **70**, 115–145.

Henricks, K., Pinkse, J. & Porter, R. H. (2003), 'Empirical implications of equilibrium bidding in first-price, symmetric, common value auctions', *Review of Economic Studies* .

Hortaçsu, A. & Kastl, J. (2012), 'Valuing dealers' informational advantage: A study of Canadian treasury auctions', *Econometrica* **80**(6), pp.2511–2542.

Hortaçsu, A. & McAdams, D. (2010), 'Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market', *Journal of Political Economy* **118**(5), pp. 833–865.

Robles-Garcia, C. (2019), Competition and incentives in mortgage markets: The role of brokers. Working paper, Staford GSB.

Stanton, R., Walden, J. & Wallace, N. (2014), 'The industrial organization of the us residential mortgage market', *Annual Review of Financial Economics* **6**, 259–288.

Vickery, J. & Wright, J. (2013), 'Tba trading and liquidity in the agency mbs market', *Federal Reserve Bank of New York Policy Review* .