Asymmetric information in the wholesale market for mortgages: The case of Ginnie Mae loans

– Preliminary and Incomplete: Please Do Not Circulate or Cite –

Ken Hendricks[1,2], Jean-François Houde[1,2], Diwakar Raisingh[1*]

November 21, 2020

1 Introduction

In this paper, we analyze the cost of financial intermediation services provided by traditional and shadow banks in the market for conforming mortgages (over 90% of loans in the US).[1] Intermediation consists of three activities: (i) the acquisition of loans, (ii) the creation of MBS pools, and (iii) the servicing of these loans. The vast majority of conforming mortgage in the U.S. are supplied via this *origin-to-distribute* channel. The supply chain involves several inter-related markets. There is a retail (or primary) market in which mortgage specialists (brokers or correspondent lenders) and banks compete to originate and fund mortgages; a wholesale market in which downstream originators use auction mechanisms to sell their loans individually to a small number of banks, many of whom also originate loans (Stanton and Wallace (2014)); and a secondary market in which banks who have acquired loans (either through the retail market or the wholesale market) negotiate default insurance premiums with one the three Agencies (Freddie Mac, Fannie Mae, and Ginnie Mae) and pool them into mortgage-backed securities (MBS), which they then sell to a large number of price-taking investors. In selling the securities, the banks retain the servicing rights. These rights are a source of profits for the seller since it earns a fee each month for collecting and

*Affiliations: [1]Department of Economics, University of Wisconsin-Madison. [2]National Bureau of Economic Research.

[1]A loan is conforming if it meets the underwriting criteria of Agencies guaranteeing payments to MBS investors in the event of default.

distributing the monthly payments of the borrower. This fee is equal to the note rate of the loan minus the coupon rate that must be paid to investors. The wholesale market is significantly more concentrated than the retail market, although originators can enjoy some pricing power due to presence of search frictions (Allen et al. (2014), Allen et al. (2019), Agarwal et al. (2020)).

Our main focus is on the interaction between the wholesale and secondary markets. The bids that banks submit in an auction depend upon their willingness-to-pay (WTP) for the loan. This WTP consists of two components: the resale value of the loan in the MBS market and the value of future cash-flows associated with servicing the loan (net of servicing costs). These values depend on the bank's and investors' assessment of the (stochastic) duration of loans.[2] A key feature of the intermediation process is that, in acquiring loans, banks have private information about loan durations and, in selling the loans, they have more information than investors or agencies. The former generates information rents to the banks and the latter leads to an adverse selection problem. The bank can get more money up front by selling a loan in a higher coupon security (i.e., high price, low fee) or more money later by selling the loan in a lower coupon security (i.e., low price, high fee). Clearly, this choice depends on the insurer's assessment of the loan's duration. It should place the loan in a lower coupon security if it believes that the risk is low and a higher coupon security if the risk is high. Furthermore, since the price of a security is set before investors observe the characteristics of the loans, the adverse selection is based on observed and unobserved (to the econometrician) characteristics of the borrower.

The main goal of this paper is to examine and quantify the effects of adverse selection the cost of intermediation. The problem of adverse-selection in the secondary market for mortgages is well documented in the Finance literature (see for instance Agarwal, Chang, and Abdullah (2012) and Downing, Jaffee, and Wallace (2009)). Our main contribution is to quantify how adverse selection affect the wholesale price of mortgages, and ultimately the borrowing cost of consumers. To our knowledge, our project is the first attempt to study adverse selection taking into account the interaction between the wholesale and secondary market for mortgages. Most of the prior literature has either focussed on the retail segment of the market (abstracting away from wholesale and securitization decisions), or analyzed the process of securitization abstracting away from the loan acquisition process. The interaction between the two markets is important for understanding frictions in financial markets, as

[2]Because of transaction costs and/or information frictions, many borrowers refinance their mortgage inefficiently, which leads to longer loan durations, and higher revenue for banks (e.g. Keys, Pope, and Pope (2016)).

well as for the design of better policies targeted at improving the stability of the financial sector, and consumers' access to credit.

We study the problem empirically by combining data on loan securitization and performance provided by Ginnie Mae, with proprietary data on loan acquisition in the wholesale market. A unique feature of our study is the use of transaction-level data in the wholesale market from one of the largest loan exchange platform operated by the FinTech company Optimal Blue (OB). The company offers a variety of pricing services to mortgage originators, including the ability of correspondent lenders to quickly compare wholesale rate sheets (i.e. posted prices) of banks in their network (see Bhutta, Fuster, and Hizmo (2019) for an analysis of price dispersion in wholesale rate sheets). The key innovation of Optimal Blue (and other similar companies) is to provide a platform in which banks can submit bids for individual loans (often more than 100 simultaneously), rather than for a bundle of loans. This feature allows for a more accurate pricing of loans based on a rich set borrower characteristics and market conditions. Sale by auction has facilitated entry into the primary market of lenders with limited capital who specialize in mortgage origination, and is a growing segment of the market. The platform quickly grew as the main selling mechanism for OB, reducing the importance of bulk or posted-price wholesale transactions. In 2019, over 75% of transactions intermediated by OB were performed via loan-level auctions.

We use this combined data set on loan acquisition and securitization to analyze two related questions. First, we measure the importance of asymmetric information between banks and MBS investors on the prepayment risk of loans. We do so by proposing a novel test for adverse-selection in this market. We then measure the effect of private information on prepayment risk on the wholesale prices for mortgages by testing for the presence of a Winner's Curse in loan exchange auctions.

We test for the presence of adverse-selection using data on the resale value of loans in the TBA market for Ginnie Mae-backed loans. This the fastest growing segment of the mortgage market (Kim et al. (2018)), and most loans in this segment cannot be insured by Freddie Mac or Fannie Mae. In addition, the securitization process is transparent and common across banks, which simplifies our model of loan acquisition and securitization.

TBA trades are forward contracts for the future delivery of the MBS, and represent the most common method to securitize loans (over 90% according to Vickery and Wright (2013)). The parties agree on a price, and on the coupon rate, loan maturity, par amount (quantity) and settlement date. The issuer can choose any pool of loans that meets these requirements and the restrictions imposed by the agency. To assess the importance of the adverse selection problem, we construct a test that is commonly used in the insurance literature (Chiappori and Salanie 2000). In particular, we test whether repayment rates are higher for the high coupon choice loans than for the low coupon choice loans. The results confirm the hypothesis that Ginnie MBS securities are adversely selected based on private information on prepayment risk.

Having established the importance of private information on prepayment risk, we evaluate its impact on the wholesale price for mortgages. In particular, we use data from OB loan exchange platform to construct a test for the Winner's Curse. Because of adverse selection in the TBA market, the expected duration of loans is an important component affecting banks' WTP. The key issue is whether the dispersion in bids is due to idiosyncratic shocks or to different signals about the borrower's repayment risk. If it is the former, then the auction is a private value auction (PV) and a bidder only needs to worry about the level of competition. But, if it is the latter, then the auction is a common value auction (CV) and, in addition the level of competition, a bidder needs to worry about adverse selection. A bidder tends to win only if it has the most optimistic assessment of the borrower's repayment risk, a phenomenon known as the "winner's curse". This issue has important implications for empirical work and, more generally, for auction and market design. The use of high-frequency auctions represents an efficient mechanism to exchange loans in this market, and lower the deadweight-loss of market power. However, if the common-value component of banks' valuation is important, the benefits of competition and technology innovations are attenuated by the winners' curse problem. Quantifying the importance of the winner's curse on bids is crucial to determine the impact of adverse-selection on the wholesale price of mortgage (and ultimately on borrowers).

2 Data

In this section, we describe the two main sources of data used in the empirical analysis. The first is data on loan securitization and performance and the second is data on mortgages sold at auction.

2.1 Loan Performance and the Secondary Market

We begin with a discussion of the data on loan performance and the secondary market. The details of this segment of the market will be important when we discuss the upstream auction market.

We focus on loans insured by Ginnie Mae, which is a public corporation responsible of insuring default risk for loans qualifying for subsidies from the Federal Housing Administration (FHA), Veterans Affairs (VA), Rural Housing Service (rural housing), and Public and Indian Housing (PIH). These loans cannot easily be sold to Fannie Mae or Freddie Mac (these are government sponsored enterprises or GSEs), which implies that this segment operates more or less independently of the others.

Our data set contains all mortgage backed securities insured by Ginnie Mae that were issued from January 2014 to September 2019 and their component loans. This data set does not include loans not in Ginnie Mae pools but are eligible for the insurance; these loans are kept on a bank's balance sheet. The data includes the main characteristics of the MBS's including the CUSIP (security identifier), coupon rate, issuance date, maturity, and par amount. The loan characteristics we see include: the CUSIP it is associated with, 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, the issuer, and origination type. Note that the issuer is the financial institution that is servicing the loan, which includes collecting payments from the borrower and disbursing it to the MBS investors. We classify the origination type as either retail or non-retail. Retail mortgages are those that are originated, serviced, and securitized by the same financial institution; otherwise, it is a non-retail loan.[3] For each component loan, we observe the unpaid principal balance on a monthly basis until it has either been paid off or defaulted.

For the remainder of the paper, we will focus on Ginnie Mae MBS's for 30 year fixed rate loans. This comprises the most common type of mortgage with 93% of the loans in our data meeting this criteria (6% are fixed rate mortgages with different maturities and the remaining 1% are adjustable rate). The characteristics of these loans are reported in Table 1. These mortgages typically have a note rates that are 4% and loan sizes on the order of \$210K though there is quite a bit of variation in these values. These loans typically have a very high LTV and 58% of them are between 95 and 100. Moreover, the FICO scores of the borrowers are a relatively low 687 (a score of 670 is the cutoff between "fair" and "good" credit). The LTV's and FICO scores are consistent with the goal of the subsidy programs. Consider the largest subsidy category, FHA loans. These loans have increased

[3] There are two types of non-retail originated loans: correspondent and broker. A correspondent is a lender that originates, underwrites, and funds the mortgage in their name then sells the loan to a larger mortgage lender. A broker originates the loan by matching a borrower with a lender that underwrites and funds the mortgage. In both cases, the institution that services the loan did not directly interact with the borrower. We abstract away from the nuance between correspondent and broker origination and treat both as non-retail mortgages.

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

Variable	N	Mean	SD	Pctl10	Pctl90
Note Rate	8,357,213	4.043	0.490	3.375	4.750
Original Loan Amount	8,357,213	209,854	114,803	93,000	357,000
LTV	7,867,307	94.392	8.884	84.790	101.690
$1(\text{LTV} \in (75, 80])$	7,867,307	0.024	0.154	0	0
$1(\text{LTV} \in (95, 100])$	7,867,307	0.580	0.494	0	1
Debt to Income Ratio	7,066,561	40.187	9.649	26.970	52.670
FICO	7,940,323	687.111	56.833	625	769
1(Purchase)	8,357,213	0.676	0.468	0	1
1(Retail)	8,357,213	0.437	0.496	0	1
1(Agency = FHA)	8,357,213	0.615	0.487	0	1
1(Agency = VA)	8,357,213	0.314	0.464	0	1
1(Agency = Rural Housing)	8,357,213	0.069	0.253	0	0
1(Paid off within 12 months)	8,357,213	0.095	0.293	0	0

Reports the summary statistics of loans that entered Ginnie Mae 30 Year Fixed Rate MBS's issued between January 2014 and September 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.

in popularity after the financial crisis, replacing privately securitized sub-prime loans, and enable low credit-scores and/or high LTV ratios have access to a mortgage (though these borrowers incur higher insurance payments over the life of the contract). Prepayment risk is a clear issue for these loans as 9.5% of them being paid off within twelve months. As we will discuss in more detail in the model section, this is a concern to the servicer and MBS investor because their income is derived from the interest payments, which cease when a loan is paid off.

The distribution of the coupon rate of the mortgage pool is illustrated in Figure 1. 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. This is mass is driven by the Ginnie Mae rule that the loans in a pool have to have an interest rate that are 0.25% to 0.75% (inclusive) greater than the coupon rate.

2.2 Loan Exchange Auctions Data

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 and other financial firms in order to free up capital that they can use to originate more mortgages. Banks purchase loans to pool them into MBS's, which they sell to investors, or to keep them on their balance sheet as an investment. Going forward, we will refer to the mortgage originators as sellers and the banks and financial firms as buyers.

Figure 1: Reports the coupon rate of Ginnie Mae 30 Year Fixed Rate MBS's issued from January 2014 to September 2019. The vertical dashed line is the median.

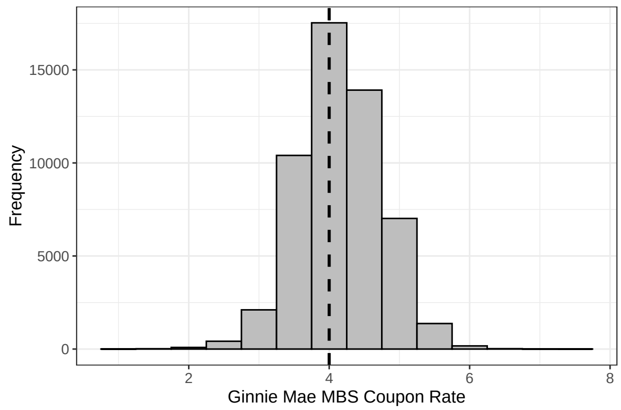

Sellers form relationships with buyers with whom they are willing to trade. Forming a relationship is costly because it involves both parties conducting due diligence as to the reliability of their counterparty and whether the originator meets the buyer's underwriting standards. Due to the cost of building these networks, relationships are sticky in the short term and re-evaluated at the quarterly level. A typical seller has anywhere from 8-15 buyers in its network.

A seller who wants to offload a mortgage on the OB platform has two options for doing so. One, she can sell the mortgage to one of the buyers in her network for a price listed on their rate sheets. A buyer's rate sheet lists the price he is willing to pay for a mortgage that satisfies specifications such as the note rate is in a certain range or the loan is for a property in certain states. Buyers typically update these rate sheets daily unless the market is very volatile, in which case they update the rate sheets multiple times per day. Two, the seller can offer the mortgage to buyers at a first price sealed bid auction. These auctions typically allow the bidders one to two hours to submit their bids. Auctions were the sale mechanism used for over 75% of the sales in 2019 and is what we have data on.

We observe mortgages eligible for Ginnie Mae insurance sold by auction during the period January 2018 to August 2018. For each mortgage we observe the following loan 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 bid and the associated bidder id, and the reserve price. We do not observe the

Figure 2: Reports the distribution of the number of bids per auction. The dashed vertical line represents the median auction.

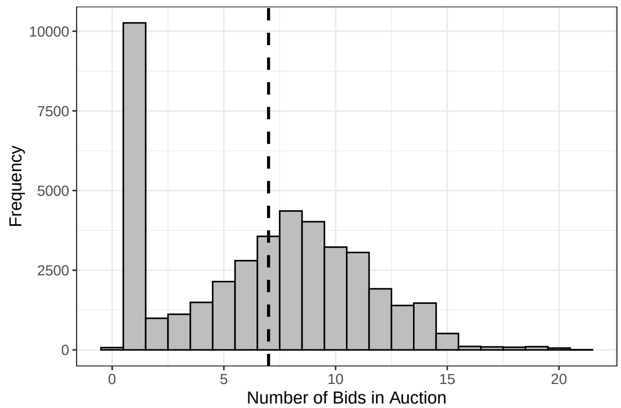

names of the sellers and bidders but our identifiers are unique. The reserve price and number of invited bidders are not observed by the bidders at the auction.

The seller usually invites all the buyers in her network to bid on a mortgage. For more specialized loans, the seller may invite fewer bidders. Invited bidders nearly always submit a bid because bidding is essentially costless and is a way of maintaining the relationship with the seller. A bidder can always submit a low bid that will be rejected if he does not want to purchase the mortgage. The distribution of the number of bids in an auction is presented in Figure 2. The specialized loans mainly correspond to the auctions that have only one bidder.

Based on the loan characteristics, we merge the mortgages in the auction data with the Ginnie Mae loan performance data described in the previous subsection; the match rate was 93%. Our discussions with industry experts suggests that the match rate is consistent with the percentage of loans that get securitized versus kept on the balance sheet.

The characteristics of the loans sold at auction are reported in Table 2. The summary statistics are quite similar to that of the mortgages in Ginnie Mae MBS's 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. These mortgages are typically securitized

Table 2: 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
$1(\text{LTV} \in (75, 80])$	42,852	0.024	0.152	0	0
$1(\text{LTV} \in (95, 100])$	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
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

and sold on the TBA MBS market and the MBS's fetch prices above the par value of the loan (loans generate interest). The TBA MBS price effectively acts as a floor on the bids that the seller will accept since she 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 will accept a price that is below the TBA MBS price and near the loan's par value. This will occur for loans that have an expected short duration (high prepayment risk) that will yield limited interest. This is not to say that sellers lose money on those loans, since consumers typically pay upfront fees to the mortgage originators. We obtain the daily Ginnie Mae TBA MBS prices from Bloomberg.[4]

Table 3 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 reserve price is effectively secret and it is common to see bids below the reserve.

[4] 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, Schultz, and Song 2017)). The daily price of a TBA security corresponds to the last observed trading price as of that date.

Table 3: Summary of Bids

Variable	N	Mean	SD	Pctl10	Pctl90
Bid	288,454	104.214	1.487	102.325	105.780
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
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
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
Highest Bid - Reserve Price	33,571	0.397	0.572	−0.059	1.049

[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.

3 A Model of Financial Intermediation for Ginnie Mae Loans

In this section, we describe a simple model of securitization and loan acquisition. We focus on loans insured by Ginnie Mae; a public corporation responsible of insuring default risks for loans qualifying for FHA, VA, and rural housing subsidies. FHA loans represent the largest category, and have increased in popularity after the crisis, replacing privately securitized subprime loans. Under this program, borrowers with low credit-scores and/or high LTV ratios have access to a mortgage, but must incur higher insurance payments over the life of the contract. Ginnie loans therefore tend to be risker (both in terms of default and prepayment), and are a common choice for first-time homebuyers who eventually transit to conventional products after building enough equity. Importantly, these loans cannot easily be sold to the GSEs, which implies that this segment operate more or less independently of the others. Another difference is that unlike Freddie Mac and Fannie Mae, Ginnie Mae is only responsible of guaranteeing payments of MBS investors, and does not acquire or securitize loans directly. Banks are responsible of assembling MBS pools, subject to Ginnie's underwriting rules. The process of securitization can be quite complex. We focus on single-issuer securities that are created by the issuer, and sold to the To-Be-Announced market (TBA) (as opposed to *specified* pools).

Security valuation. MBSs are known as *pass-through securities*. They generate two sources of income for the bank: (i) upfront security price ($P(c)$), and (ii) monthly service income (stochastic). The monthly service income is 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. For a loan i placed in an MBS with coupon choice of c, the revenue to the bank for a \$100 is given by:

$$R_i(c) = \underbrace{P(c)}_{\text{upfront payment}} + \underbrace{\sum_{\tau=1}^{T} \delta^{\tau} L_{\tau,i}}_{\text{service multiple } (M_i)} \times \underbrace{\frac{r_i - g - c}{1200}}_{\text{service income}} \tag{1}$$

where δ is bank's discount rate, $L_{\tau,i}$ is the loan balance at the end of month τ, and T is the term of the loan. The second term is the discounted value of the service fees, which is the product of a service multiple, M_i, and service income. In practice, M_i is a random variable because it depends on the stochastic process determining prepayments, repayment, or default. The service income captures how the servicer collects a fraction of the interest produced by the loan. 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 security price $P(c)$ reflects the market expectation about M_i when it is placed in a security with coupon c. Loans that are expected to generate more cash-flows than this market average are more valuable for the bank (i.e. high M_i).

The discount factor δ measures the weight that lenders place on future cash flows (relative to upfront cash payment). This weight may vary across lenders and time because of reserve requirements and the need for liquidity. In general, banks put more weight on liquidity than MBS investors, especially shadow-banks. This generates gains from trade even in the presence of asymmetric information (Downing et al. (2009)). Since the Ginnie market is composed mostly of shadow banks with similar liquidity needs, in what follows we assume that δ is the same across banks. We plan to allow δ to vary across lenders when we consider extensions of this model, especially to loans insured by Fannie Mae and Freddie Mac.

Revenue Optimization. In the Ginnie TBA market, banks face the same security prices and, for a given loan, can select only two characteristics: (i) the delivery month, (ii) the coupon rate. In general, banks select the earliest delivery date available (next calendar month accounting for a 1-2 weeks of delivery time), and so we focus on the coupon choice. Subject to restrictions on the servicing income imposed by agencies, banks can choose the coupon that maximize their expected revenue, taking the resale price as given. In the finance literature, this problem is described as the "best execution" of an MBS.

This choice depends on the relative resale price, as well firms' beliefs about multiplier (M_i). Sellers 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 (i.e. $P(c)$ is increasing in c). In general, banks select low coupons when they believe a loan will last longer than the market average. In contrast, for loans that are expected to be pre-paid early, banks will choose a higher coupon and earn a large profit from the resale price.

Ginnie Mae charge 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, this 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 MBS. For instance, a 4.25 note rate mortgage can be pooled in a 4% coupon MBS (high) or in a 3.5% coupon MBS (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.

Since Ginnie does not collect other upfront payments from banks, the optimal coupon for eligible loans is defined as a binary discrete choice problem:

$$c_i^* = \begin{cases} c_L & \text{If } m_i > \frac{P_H - P_L}{(c_H - c_L)/1200}, \\ c_H & \text{If } m_i < \frac{P_H - P_L}{(c_H - c_L)/1200}. \end{cases} \tag{2}$$

where m_i is the expected value of M_i, c_L and c_H denotes the low/high coupons available for loan i, and $P_L < P_H$ are the associated TBA prices.

Equation 2 makes it clear that only high duration loans are placed in the low-coupon option. 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. Note that securities traded on the TBA market can be adversely selected based on observed and unobserved characteristics (to the econometrician) of the loans, since the upfront payment received by issuers does not reflect the composition of MBS pools.

Following Chiappori and Salanie (2000), this prediction can be tested empirically by measuring the correlation between the choice of a high coupon ("low deductible"), and the probability that a loan gets prepaid early ("accident"). This test can be implemented either via regressions by measuring differences in the prepayment probability across the two contract types, or by measuring the correlation in two binary outcomes using a bivariate probit model.

Loan Acquisition. The wholesale channel is split between brokers and correspondent lenders. In 2017, 35% of all loans were originated by correspondent lenders, and 12% were originated by a broker. The remainder (53%) correspond to *retail* loans originated by vertically integrated banks. The fraction of correspondent loans is above 40% for Ginnie Mae MBSs. Correspondent loans are originated and funded by mortgage specialists or small/medium size deposit institutions who sell the loan at auction to banks, typically a few days after the closing date. Although the exact selling mechanism varies slightly across firms, our discussions with industry experts confirm that for the most part loans are sold via auction. In the retail and broker channels, banks purchase loans directly from borrowers. Lenders pay for the loan (*par* value), in exchange for a future stream of interest payments, and an upfront payment (or rebate) from the borrower. The retail price is given by lenders' rate sheets. In addition, rate sheets also depend on consumer characteristics such as FICO, location or income. This defines the boundaries of the markets in which lenders compete.[5]

Banks differ in their portfolio of retail and wholesale loans. Quicken Loans or Bank of America, for instance, are vertically integrated, and rely almost exclusively on the retail channel to acquire loans. In contrast several large shadow-banks like Pennymac do not originate loans in the retail markets, and rely exclusively on the correspondent channel.[6] For the most part however, banks and shadow banks manage a diversified portfolio of retail and correspondent loans. The exact nature of banks' portfolio is not measured accurately from public data sources such HMDA (see Stanton and Wallace (2014) for a discussion). To analyze the role of vertical integration, we will subscribe to the survey conducted by *Inside Mortgage Finance* to measure the mixed of retail and wholesale loan securitized by each banks in a given period.

Our auction data comes from one of the largest loan exchange platform Optimal Blue (OB). The OB platform allows correspondent lenders to invite bidders to submit separate bids for each loan that they want to sell. This leads to a large number of simultaneous first-price sealed-bid auctions The set of potential bidders varies across correspondent lenders. For each lender, this set is fixed in the short-run (between 8 to 15 lenders), and expanding

[5]See Fuster et al. (2013) for an analysis of lenders' rate sheets.

[6]In the finance literature, institutions buying loans from correspondent lenders are often called "aggregator" or "sponsor" banks (Fuster et al. 2017)

the network involves sunk costs associated with meeting the underwriting standards of new upstream banks. In most auctions, all potential bidders are invited to bid (except in the case of more specialized loans), and the participation is nearly 100%.

Each bank j observes a private signal S_{ij} about its willingness-to-pay W_{ij} for loan i. Let Z_i denote a vector of borrower and loan characteristics. We consider two models of the information structure. In the **private value** (PV) model, bank j's willingness-to-pay is given by

$$W_{ij} = \bar{R}_i + S_{ij} \tag{3}$$

where $\bar{R}_i$ is the expected revenue

$$\bar{R}_i = P(c) + \left(\frac{r_i - c - g}{1200}\right) E[M_i|Z_i]$$

and S_{ij} is a private idiosyncratic value shock (e.g., costs or liquidity). Here r_i is common to each bidder since they face the same security price and service fee, and have the same expectation of M_i. Bidders compete away this value, so the dispersion in bids reflects the dispersion in the value shock. In the **common value** (CV) model, the (unknown) value of the loan is the same for all bidders

$$W_{ij} = R_i,$$

and each bidder j gets an informative signal S_{ij} about R_i before bidding. In this case, the dispersion in bids reflects the dispersion in signals.

Our tests assume that the auction has a pure strategy, monotone equilibrium. Let $B_{ij} = \beta_{ij}(S_{ij}; Z_i)$ denote bidder i's bid function for loan i. Then, under PV,

$$E[M_i|S_{ij} = s, Z_i] = E[M_i|B_{ij} = b, Z_i] = E[M_i|Z_i].$$

The first equality follows from monotonicity and the second from independence of S_{ij} and M_i. By contrast, under CV,

$$E[M_i|S_{ij} = s, Z_i] = E[M_i|B_{ij} = b, Z_i]$$

in strictly increasing in b. Thus, ex post measures of loan performance do not vary with b if the auction is PV and they are strictly increasing in b if the auction is CV. We refer to this test as the **monotonicity** test.

We refer to our second test as the **Winner's Curse** test. Bidder j wins the auction when it submits the highest bid. Let $B_{i,-j}$ denote the vector of bids submited by j's rivals for loan i. Then, under PV,

$$E[M_i|B_{ij} = b, \max\{B_{i,-j}\} < b, Z_i] = E[M_i|Z_i]$$

and under CV,

$$E[M_i|B_{ij} = b, \max\{B_{i,-j}\} < b, Z_i] < E[M_i|B_{ij} = b, Z_i]$$

If the auction is PV, there is no selection effect: since ex post measures of loan performance not vary with b, they also do not vary when b is the winning bid. By contrast, if the auction is CV, there is a selection effect: winning is "bad news". It means that rivals have lower signals, which implies lower survival rates.

We run these tests on the sample of Ginnie Mae loans where bidders cannot choose the coupon rate - i.e., the note rate uniquely determines the coupon rate. For these loans, the security price and service fee are the same for each bidder. The only source of variation in the willingness-to-pay across bidders is their private signal, which can reflect differences in information about prepayment risk or costs. An extension that we intend to explore is a richer model of bidder heterogeneity in which bidders can differ in both liquidity needs and information about prepayment risk.

4 Testing for adverse-selection in pre-payment risk

In this section we present the results of an analysis of the importance of adverse selection in the market for Ginnie Mae securities. We exploit the fact that lenders face a coupon choice for mortgages with note-rates ending in 0.25 and 0.75. For those loans, the optimal coupon choice is the result of a tradeoff between future cash flows, and resale price. Since loans are sold on the TBA market (future contract), the resale value is not function of the loan characteristics (including the note rate). Everything else being equal, banks are more likely to deliver loans with speedier expected payments in high-coupon MBS (lemons), and place loans with a longer expected duration in low-coupon MBS.

We test this prediction of the model using publicly available data on loan performance and securitization provided by Ginnie Mae for the period between 2014 and 2018. We measure the activity status of loans in the Fall of 2019. This data-set includes an identifier for the MBS security attached to each loan (CUSIP), as well as a large set of borrower characteristics (i.e. note rate, risk score, loan size, loan purpose, region, etc.). We supplement this data with the price of TBA trades prior to the delivery date of each security (from TRACE). We use this to measure the difference in upfront payments ($P(c)$) between high and low coupon securities.

Our first set of results are obtained by analyzing the reduced-form relationship between the choice of a high coupon, and the probability of early pre-payment. This is similar to the analysis in Agarwal et al. (2012). Importantly the popularity of the low-coupon option has been increasing over time, since the risk of prepayment decreased between 2013 and 2018 (i.e. decreasing yield curve). For much of this period, the default option for most lenders is the high coupon MBS. By the end of the period, however, nearly 40% of banks were mixing between the two options.

To account for this underlying trend in prepayment and coupon choice, we compare loans with the same note-rates that are securitized in the same period (month/year). This leads to the following linear probability model with note-rate/period fixed effects ($\mu_{r,t}$):

$$Y_{it} = \alpha 1\{\text{Coupon eligible}\}_{it} \times 1\{\text{High coupon}\}_{it} + \lambda 1\{\text{Coupon eligible}\}_{it} + X_{it}\beta + \mu_{r,t} + e_{it}$$

where X_i are loan characteristics. The "Coupon eligible" indicator variable is equal to one if the note-rate ends with 0.25 or 0.75. The dependent variable is an indicator equal to one if the loan pre-paid prior to month $T \in \{12, 24, 36\}$. We multiple the dependent variable by 100 to facilitate the interpretation of α (i.e. percentage difference in pre-payment rate).

Preliminary results are presented in Table 4. We estimate equation (4) separately for loans acquired in the retail (odd columns) and wholesale (even columns) channels. Since we condition on borrower and loan characteristics, the coefficient measures the extent of adverse-selection from private information learned by lenders through the loan acquisition process. In both channels, we find that loans placed in high-coupon securities are adversely selected, relative to loans with the same note-rate and delivery date pooled in low coupon MBS. Interestingly, high-coupon loans acquired through the retail channel are between 9% and 10% more likely to pre-pay early (compared to the same average), while high-coupon loans originated by brokers and correspondent are between 5% and 7% more likely to prepay early. This difference is stable across specifications (i.e. duration cutoff) and is statistically significant from zero.

Next we measure directly the correlation between the choice of using a high-coupon and the pre-payment probability using a bivariate probit model similar to one used in Chiappori

Table 4: Estimated differences in pre-payment probability between high and low coupon MBS

VARIABLES	(1) $T < 12$	(2) $T < 12$	(3) $T < 24$	(4) $T < 24$	(5) $T < 36$	(6) $T < 36$
High coupon x Coupon eligible	0.94***	0.678***	1.79***	1.141***	2.11***	1.200***
	(0.11)	(0.0922)	(0.13)	(0.111)	(0.14)	(0.117)
Constant	-2.20***	1.105***	3.97***	5.930***	7.02***	5.261***
	(0.47)	(0.394)	(0.61)	(0.509)	(0.67)	(0.553)
Observations	3,369,327	4,741,907	3,369,327	4,741,907	3,369,327	4,741,907
R-squared	0.109	0.110	0.134	0.144	0.164	0.177
Channel	retail	wholesale	retail	wholesale	retail	wholesale
Borrower characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	9.41	9.405	18.5	18.47	24.2	24.17

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

and Salanie (2000):

$$\mathbf{1}(\text{High Coupon})_i = Z_i = \mathbf{1}(X_t\beta + \varepsilon_t > 0),$$
$$\mathbf{1}(\text{Prepaid in 1 Year})_i = Y_i = \mathbf{1}(X_t\gamma + \eta_t > 0).$$

Here ε_t is an unobserved shock determining lender's coupon choice, and η_t is an unobserved shock determining the borrower's decision to prepay the loan in the first year (i.e. $T = 12$). The test focuses on the correlation between ε and η (ρ). If ε includes private information on the likelihood of prepayment in first year, it should also help predict this event. It implies that ε and η are positively correlated, and that more information leads to higher correlation.

As with the regression analysis, we account for differences across note rates by estimating the model separately for the most common note-rates available during our sample period: 3.75, 4.25, 4.75, and 5.25. Recall that roughly half of lenders almost never exercise the low-coupon option when available. To estimate the bivariate probit models, we focus on the sub-sample of banks who exercise this option more than 10% of the time. This represents a sub-set of the full sample used in the regression analysis above.

Table 5 presents the main results. The top panel reports the estimation results for retail loans. Consistent with the regression results, we find that the propensity to use a high-coupon is positively correlated with early pre-payment. The correlation is different from zero for all note-rates expect for the 5.25 contract (which is the less traded loan in our sample).

Table 5: Bivariate Probit results

(a) Retail loans

NoteRate	N	N: High Coupon $Y=0$	$Y=1$	N: Low Coupon $Y=0$	$Y=1$	Bivariate Probit $\hat{\rho}$	se($\hat{\rho}$)	95% CI
3.750	134,483	115,693	9,150	9,116	524	0.039	0.012	(0.015, 0.063)
4.250	144,962	102,574	22,352	18,025	2,011	0.111	0.009	(0.094, 0.128)
4.750	59,564	28,099	10,493	17,324	3,648	0.094	0.010	(0.075, 0.113)
5.250	17,520	3,783	1,301	9,025	3,411	0.009	0.017	(-0.024, 0.043)

(b) Wholesale loans

NoteRate	N	N: High Coupon $Y=0$	$Y=1$	N: Low Coupon $Y=0$	$Y=1$	Bivariate Probit $\hat{\rho}$	se($\hat{\rho}$)	95% CI
3.750	232,468	202,574	13,515	15,553	826	0.018	0.011	(-0.003, 0.039)
4.250	194,705	149,402	20,240	22,357	2,706	0.037	0.008	(0.022, 0.053)
4.750	72,875	47,318	6,512	16,421	2,624	0.029	0.010	(0.009, 0.05)
5.250	22,654	10,258	1,570	9,177	1,649	0.065	0.014	(0.037, 0.093)

The bottom panel presents the same results for loans acquired in the wholesale market. Altough the correlation coefficients are statistically different from zero, the magnitude of the correlation is smaller in the wholesale channel. This is consistent with the idea that vertically integrated banks are more informed about prepayment risks than banks who rely on third-party originators.

These two set of results confirm that the TBA market is adversely selected based on private information that banks have about pre-payment risk. This is consistent with the results obtained previously by Agarwal et al. (2012) and Downing et al. (2009), although our identification strategy is different.

5 An analysis of the Winner's Curse in the wholesale market

In this section we present the results of our monotonicity and winner's curse tests of the PV model against the alternative of CV model. We use the valuation model describe in Section 3 to conduct our analysis. In general, the willingness to pay of banks is a non-linear function of their beliefs about prepayment and preference for liquidity, because of the coupon optimization problem. However, the Ginnie Mae securitization rules mean that there are loans that can only enter one coupon (i.e. loans with a note rate that does not end in .25 or .75). We focus on these loans without a coupon option for the winner's curse test.

Before discussing the results of this test, we first present a descriptive analysis of the distribution of bids in the loan-exchange platform. To describe the bidding patterns we see in the data, we estimate the following OLS regression of bids ($b_{i,t}$):

$$\ln(b_{it} - P(c)) = x_t'\beta + \delta_s + \delta_i + \delta_m + \delta_{\mathcal{N}} + \epsilon_{it}, \tag{4}$$

where i denotes an investor (bidder), t denotes a loan (auction), s denotes a seller, $P(c)$ is the TBA MBS price that the loan can enter, x_t are loan characteristics, δ_s are seller fixed effects, δ_i are investor fixed effects, δ_m are month fixed effects, and $\delta_{\mathcal{N}}$ are number of bidder fixed effects. Note that the natural logarithm of the service income, $\ln(\frac{r_{it}-c-g}{1200})$, is included in the loan characteristics and that when creating a fixed effect for number of bidders, we pool all loans with 16 or more bidders. Since the dependent variable is a logarithm, we only focus on bids that are above the MBS price. This specification of the regression allows us to test the validity of our valuation model. Algebraically manipulating equation (1) results in the logarithm of the WTP less the MBS price is linear in the logarithm of the service income, which has a coefficient of one. This implies that the logarithm of the service income in our specification should have a coefficient of one. This is because the service income is a common valuation component for all bidders so they will all adjust their bids by this amount.

We run the regression on two samples and report the results in Table 6. Column (1) shows the estimates when we used all bids that were above the MBS price. Column (2) is similar but focuses on bids in auctions that received at least two bids above the MBS price; this is the sample on which we run the winner's curse test. Both specifications produce similar estimates. The estimates show that logarithm of the service income has a coefficient that is close to one, which is consistent with our valuation model's prediction and is evidence of its validity.

To understand which covariates account for the variation in the bids, we use specification

Table 6: Bid OLS Regression

	ln(Bid - TBA MBS Price)	
	(1)	(2)
$\ln((r-g-c)\ /\ 1200)$	1.008*** (0.054)	1.016*** (0.054)
NoteRate	0.762*** (0.226)	0.702*** (0.226)
NoteRate^2	−0.106*** (0.023)	−0.099*** (0.023)
ln(LoanAmount)	−0.128*** (0.049)	−0.120** (0.049)
LTV	−0.005* (0.003)	−0.006** (0.003)
$1(\text{LTV} \in (75, 80])$	0.005 (0.009)	0.004 (0.009)
$1(\text{LTV} \in (95, 100])$	0.022*** (0.004)	0.021*** (0.004)
Debt to Income Ratio	0.001 (0.0005)	0.001* (0.0005)
FICO	0.001 (0.001)	0.0002 (0.001)
Monthly Income ($K)	−0.003*** (0.001)	−0.003*** (0.001)
$1(\text{Monthly Income} \geq \$20\text{K})$	−0.320*** (0.019)	−0.328*** (0.019)
1(LP = Purchase)	−2.502*** (0.549)	−2.929*** (0.549)
1(LP = Refi Cashout)	−2.827*** (0.559)	−3.132*** (0.559)
1(LP = Refi Rate-and-Term)	−1.967*** (0.556)	−2.450*** (0.556)
1(LP = VA Rate Reduction)	−2.624*** (0.606)	−3.090*** (0.606)
1(Occupancy = Primary)	−0.142** (0.064)	−0.183*** (0.064)
1(Occupancy = 2nd Home)	−0.082 (0.152)	−0.103 (0.152)
1(Agency = VA)	−0.128** (0.052)	−0.130** (0.052)
1(Agency = FHA)	0.097* (0.051)	0.096* (0.051)
1(Units = 2)	−0.027** (0.012)	−0.029** (0.012)
1(Units = 3)	−0.037* (0.020)	−0.040** (0.020)
1(Units = 4)	−0.018 (0.048)	−0.018 (0.048)
LTV×1(LP = Purchase)	0.006* (0.003)	0.007** (0.003)
LTV×1(LP = Refi Cashout)	0.006** (0.003)	0.007** (0.003)
LTV×1(LP = Refi Rate-and-Term)	0.005 (0.003)	0.006** (0.003)
LTV×1(LP = VA Rate Reduction)	0.008*** (0.003)	0.009*** (0.003)
FICO×1(LP = Purchase)	0.003*** (0.001)	0.003*** (0.001)
FICO×1(LP = Refi Cashout)	0.003*** (0.001)	0.004*** (0.001)
FICO×1(LP = Refi Rate-and-Term)	0.002*** (0.001)	0.003*** (0.001)
FICO×1(LP = VA Rate Reduction)	0.003*** (0.001)	0.003*** (0.001)
Observations	182,638	174,847
R^2	0.376	0.371
Adjusted R^2	0.376	0.370

*p<0.1; **p<0.05; ***p<0.01. Robust standard errors reported and clustered at the bidder level.

- Reports the regression estimates of the regression specified in (4). Column (1) includes all bids that were at least the TBA MBS price. Column (2) is similar but only focuses on auctions that had at least two bids above the TBA MBS price. Both regressions include seller, bidder, U.S. state, month, and number of bidder fixed effects (not reported).
- The left out `Agency` is Rural Housing.
- `LP` stands for loan purpose and the left out group is "FHA Streamline Refi".
- Left out group for `Occupancy` is "Investment Property".
- Left out group for `Units` is 1.

Table 7: Variance Decomposition of Bids

Variable	Value	Percentage
$Var(ln(b_{it} - P(c)))$	0.349	100.000
$Var(\delta_s)$	0.015	4.231
$Var(\delta_i)$	0.022	6.204
$Var(\delta_m)$	0.0004	0.109
$Var(\delta_{\mathcal{N}})$	0.017	4.999
$Var(x_t'\beta)$	0.091	26.110
$Var(\epsilon_{it})$	0.220	62.884
$2Cov(\delta_s, \delta_i)$	0.0002	0.046
$2Cov(\delta_s, \delta_m)$	-0.0001	-0.020
$2Cov(\delta_s, \delta_{\mathcal{N}})$	-0.019	-5.420
$2Cov(\delta_s, x_t'\beta)$	0.001	0.347
$2Cov(\delta_i, \delta_m)$	0.00001	0.004
$2Cov(\delta_i, \delta_{\mathcal{N}})$	-0.002	-0.624
$2Cov(\delta_i, x_t'\beta)$	-0.002	-0.590
$2Cov(\delta_m, \delta_{\mathcal{N}})$	-0.0003	-0.078
$2Cov(\delta_m, x_t'\beta)$	0.001	0.147
$2Cov(\delta_{\mathcal{N}}, x_t'\beta)$	0.006	1.651

Reports the variance decomposition of ln(Bid - TBA MBS Price) based on regression specification (2) in Table 6.

(2) in Table 6 to do the following variance decomposition:

$$\begin{aligned} Var(ln(b_{it} - P(C))) = {} & Var(\delta_s) + Var(\delta_i) + Var(\delta_m) + Var(\delta_{\mathcal{N}}) + Var(x_t'\beta) + Var(\epsilon_{it}) \\ & + 2Cov(\delta_s, \delta_i) + 2Cov(\delta_s, \delta_m) + 2Cov(\delta_s, \delta_{\mathcal{N}}) + 2Cov(\delta_s, x_t'\beta) \\ & + 2Cov(\delta_i, \delta_m) + 2Cov(\delta_i, \delta_{\mathcal{N}}) + 2Cov(\delta_i, x_t'\beta) \\ & + 2Cov(\delta_m, \delta_{\mathcal{N}}) + 2Cov(\delta_m, x_t'\beta) \\ & + 2Cov(\delta_{\mathcal{N}}, x_t'\beta). \end{aligned}$$

We report the variance decomposition in Table 7. Notice that the loan characteristics drive accounts for the 26% of the variance in the bids. This suggests that these characteristics are informative about the duration of the loan. However, even after controlling for many observable characteristics, the private signals about the loan's value as represented by the epsilon shock accounts 62% of the variation in the bids. We further decompose the variance of the residuals (ϵ_{it}) by regressing them against auction fixed effects and seller-bidder fixed effects. The latter fixed effect is meant to capture any relationship effects on the bids. This variance decomposition is reported in Table 8. While the auction fixed effects explain 37% of the variance in the residuals, 63% of the variance in the residuals cannot be explained by the fixed effects. This suggests that there is a bidder specific private signal that drives a lot of the variation in the bids.

Table 8: Variance Decomposition of Bid Regression Residuals

Variable	Value	Percentage
$Var(\epsilon_{it})$	0.220	100.000
$Var(\delta_t)$	0.080	36.512
$Var(\delta_{i,s})$	0.011	4.823
$2Cov(\delta_t, \delta_{i,s})$	-0.009	-3.997
$Var(\eta_{it})$	0.138	62.662

Reports the variance decomposition of the residuals from specification (2) in Table 6. The residuals were regressed against auction fixed effects and seller-bidder fixed effects; η_{it} is the residual of this latter regression.

To implement our common-value test we use the Ginnie Mae loan performance data (from 2013 to 2019) to estimate a single-index that proxies for banks' expectation about future cash-flows. We use a Logit model to predict the probability of survival for more than one year. Let $\lambda(Z_i)$ denotes estimated odds-ratio associated with a loan with characteristics Z_i. In a revision, we plan on enriching this model of cash flow using a more sophisticated model of prepayment risk, and use machine-learning techniques to improve the predictions.

We then combine the loan performance and bid information data-sets by matching the OB data from 2018 with Ginnie's loan performance data. Since the platform collects a very rich set of consumer characteristics (including zip code), we are able to create a near perfect match. This allows us to evaluate the correlation between bids and the probability that the loan survives for more than a year. Specifically, we estimate the following Probit model:

$$\Pr\left(T_i > 12 | Z_i, \tilde{b}\right) = \Phi\left(\beta_0 + \beta_1 \tilde{b} + \beta_2 \lambda(Z_i) + \beta_3 \lambda(Z_i) \cdot \tilde{b}\right) \tag{5}$$

where

$$\tilde{b} = \frac{b - P(c)}{(r - c - g)/1200}$$

is a transformed bid (homogenizing the auctions) and $\Phi(\cdot)$ is the normal CDF. The coefficient on $\widetilde{b}$ captures any private information that a bidder has about prepayment risk. We estimate this model separately on the full sample and on the subsample of winning bids only. The difference between the two survival functions evaluated at the same bid value is a measure of the winner's curse.

Figure 3 presents the results. The left panel reports the Probit estimates on the full sample. The probability that the loan survives at least one year is increasing in the predicted odds-ratio and in $\widetilde{b}$ (except at very small values of λ). The latter result indicates that bids

Figure 3: Estimates of the Winner's Curse problem in the sample of Ginnie Mae loans

(a) Probit estimates

	$1(T_i > 12)$	
	All	Winners
	(1)	(2)
Constant	-1.672^{***}	-1.180^{***}
	(0.077)	(0.237)
$\tilde{b}$	-0.253^{***}	-0.265^{***}
	(0.022)	(0.056)
$\lambda(Z)$	3.274^{***}	2.582^{***}
	(0.091)	(0.283)
$\tilde{b} \times \lambda(Z)$	0.307^{***}	0.343^{***}
	(0.026)	(0.064)
Observations	165,209	19,698
Log Likelihood	−50,940.620	−6,130.793

*p<0.1; **p<0.05; ***p<0.01

(b) Predicted survival probabilities

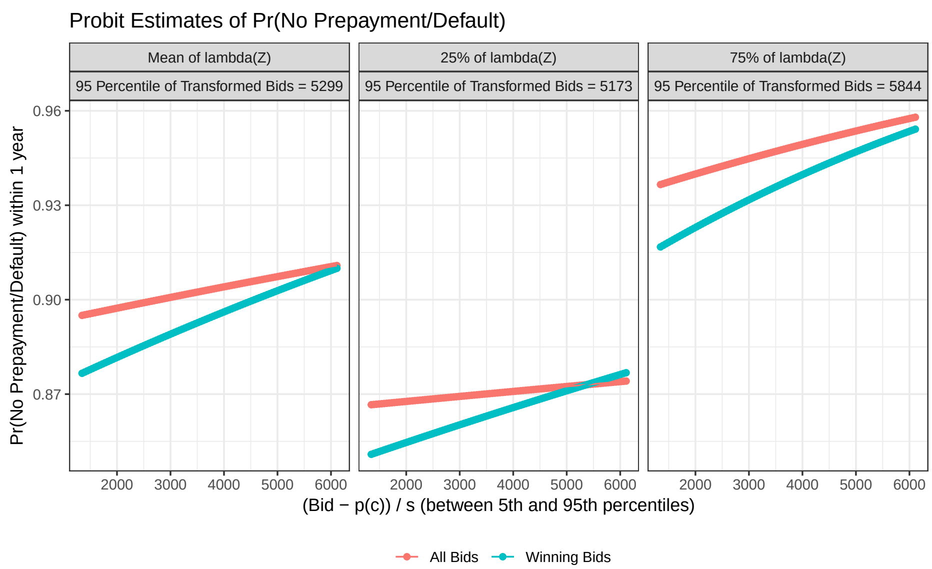

possess information about the duration of loans, which is not consistent with the PV model but is consistent with the CV model.

The difference between the Probit estimated on the sample of winning bids and on the full sample reveals the magnitude of the Winner's Curse. The right panel illustrates this difference for a grid of values for the transformed bid $\tilde{b}$. The left graph shows that, at the mean of $\lambda(Z_i)$, the survival probability is always lower for the winning bidder. This result is not consistent with the PV model and strongly supports the CV model. Note that the gap between the two curves decreases with $\widetilde{b}$. This result is also consistent with the CV model since the difference between the two conditional expectations converges to zero as the bidder's signal approaches the upper bound of the support of distribution of signals.

These results clearly demonstrate that the willingness to pay of banks includes a large common-value component. This is important for understanding the cost of financial intermediation in this market. Banks behaving rationally in the wholesale market should avoid bidding too aggressively, especially when the number of competitors is expected to be large. In equilibrium, the presence of common-values affects the profitability of the market, and implies that an increase in competition may not lead to lower wholesale prices. This is consistent with the fact that the intermediation margins have been increasing over the period (see Fuster et al. (2013) and Fuster, Lo, and Willen (2017)), despite the growth and entry of new shadow banks, and important technology innovations (Buchak et al. (2018), Fuster et al. (2019)). This increase in competition, combined with a winners' curse problem, could also explain why traditional banks have largely exited the Ginnie Mae segment following the entry of shadow banks (assuming common support).

The importance of the winner's curse is also impacted by the quality of the information available to bidders. The fact that intermediaries like Optimal Blue introduced loan-level auctions, combined with improvements in information technology and machine learning models, have certainly contributed to alleviate the Winner's Curse problem, and lower the cost of intermediation.

References

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

Agarwal, S., J. Grigsby, A. Hortaçsu, G. Matvos, A. Seru, and V. Yao (2020, June). Searching for approval. NBER working paper, 27341.

Allen, J., R. Clark, and J. F. Houde (2014). The effect of mergers in search markets: Evidence from the Canadian mortgage industry. *American Economic Review 104*, 3365–3396.

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

Bhutta, N., A. Fuster, and A. Hizmo (2019, July). Paying too much? price dispersion in the us mortgage market. Working paper, Federal Reserve Board.

Buchak, G., G. Matvos, T. Piskorski, and A. Seru (2018, December). Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics 130*(3), 453–483.

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

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

Fuster, A., L. Goodman, D. Lucca, L. Madar, L. Molley, and P. Willen (2013, Dcember). The rising gap between primary and secondary mortgage rates. *FRBNY Economic Policy Review*, 17–39.

Fuster, A., S. H. Lo, and P. Willen (2017, January). The time-varying price of financial intermediation in the mortgage market. Working paper, NYFED.

Fuster, A., M. Plosser, P. Schnabl, and J. Vickery (2019). The role of technology in mortgage lending. *The Review of Financial Studies 32*(5).

Gao, P., P. Schultz, and Z. Song (2017, June). 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).

Keys, B. J., D. G. Pope, and J. C. Pope (2016). Failure to refinance. *Journf of Financial Economics 122*, 482–499.

Kim, S. Y., S. Laufer, K. Pence, R. Stanton, and N. Wallace (2018, Spring). Liquidity crises in the mortgage market. *Brookings Papers on Economic Activity,*.

Stanton, Richard, W. J. and N. Wallace (2014). The industrial organization of the us residential mortgage market. *Annual Review of Financial Economics 6*, 259–288.

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