Competition and Strategic Incentives in the Market for Credit Ratings: Empirics of the Financial Crisis of 2007*

Chenghuan Sean Chu† Marc Rysman‡
Sunrun Inc. Boston University

October 21, 2018

Abstract

We study the market for ratings agencies in the commercial mortgage backed securities sector leading up to and including the financial crisis of 2007-2008. Using a structural model adapted from the auctions literature, we characterize the incentives of ratings agencies to distort ratings in favor of issuers. We find important equilibrium distortion, which decreased after the crisis. We study several counterfactual experiments motivated by recent policy-making in this industry.

1 Introduction

Credit ratings agencies play a key role in the functioning of financial markets, and some observers point to the role of credit ratings agencies in exacerbating, or perhaps even causing, the financial crisis of 2007 in the United States.[1] Whereas many financial markets

*Sean Chu is the lead author of this article. We wish to thank four anonymous referees and the editor for excellent advice on how to improve the paper.

†seanchu@gmail.com. This research was conducted during my time at the Federal Reserve Board of Governors. Views expressed herein are my own and do not reflect the views of either the Federal Reserve or Sunrun Inc.

‡mrysman@bu.edu

[1]Examples of such statements in the popular press and government reports are Krugman (2010) and Financial Crisis Inquiry Commission (2011).

are characterized by enormous numbers of relatively small players, the market for ratings is dominated by only three firms, and the market faces strategic concerns such as market power and asymmetric information. Thus, the scope for inefficiency, and perhaps the potential benefits of government intervention, are particularly high. This paper analyzes credit ratings agencies in the market for commercial mortage-backed securities (CMBS), and characterizes the extent of market power and how it evolved during the financial crisis. We show that strategic incentives to increase ratings are important, but decreased substantially as a result of the financial crisis. We study several counterfactual scenarios, motivated by several prongs of regulatory reform targeting the ratings agency market.

Credit ratings agencies rate financial products for riskiness. In order to determine which agencies rate a CMBS product, the issuer observes the ratings offered by each agency and then chooses which ratings agencies to use, typically selecting two agencies. Even if ratings agencies accurately report their assessments, issuers may practice *ratings shopping*, that is, the issuer may choose only the highest ratings among agencies, which distorts ratings upwards.[2]

Furthermore, the issuer pays fees to the agency that provide the ratings, which puts the agencies in competition to attract the issuer. *Ratings catering* refers to when ratings agencies boost ratings in order to appeal to issuers, presumably at the cost of the agencies' reputations with investors and regulators, to the extent that ratings prove misleading. Ratings shopping and ratings catering may drive up ratings relative to what ratings agencies believe is appropriate. That is important because not only does the rating provide information to potential investors, but also, many institutional investors face restrictions on how much risk they may take on, so the rating determines the size of the potential investment pool for the financial product. Thus, inappropriately high ratings distort the information on the market and allow financial institutions to invest in riskier portfolios than regulators intend to allow. These issues are the focus of much criticism of the industry, and have led to calls for reform.

The CMBS market was arguably central in triggering the financial crisis, and collectively experienced $39.1 billion in foreclosure liquidations during that time. We study ratings for

[2]According to the chief operating officer of Moody's, "There is a lot of rating shopping that goes on." (from Bolton, Freixas & Shapiro, 2012).

CMBS issued between the early 2000s, when the CMBS market first emerged, until the tail end of the crisis in 2010, and continuing to 2012. We draw data on CMBS from the data vendor Morningstar. We also collect by hand data from publicly available sources on the evaluation of the CMBS by the chosen ratings agencies, a new data set to our knowledge.

We provide a model of how issuers obtain ratings. In the model, the issuer has a financial product, and each ratings agency observes a private signal of the true riskiness level of the product. Each ratings agency proposes a rating, or more specifically, what percentage of the product the agency is willing to rate with the highest rating, AAA. The issuer chooses the most favorable ratings. Thus, the game is similar to an auction in which the issuer is the buyer or the auctioneer, the agencies are the bidders, the signal is the private value that the bidder places on the object, and the rating is the bid. The difference between the bid and signal is the margin, although unlike usual auctions, the bidder prefers low margins and the buyer prefers high margins. Insights from the auction literature guide our empirical approach in several ways, as discussed below. To our knowledge, this is the first paper to apply structural estimation to the credit ratings market. We adapt arguments from Athey & Haile (2002) to show that our model is identified under general functional form assumptions. Our estimation approach adapts two-step methods from the auction literature, such as Guerre, Perrigne & Vuong (2000).

We provide a measure of the distortion in proposed ratings from agencies' true beliefs (akin to a margin in an oligopoly setting or bid-shading in an auction setting), and we validate the measure by showing that this measure well-predicts ex post performance, even conditional on the rating. We calculate how this distortion evolved as a result of the financial crisis. We find a large adjustment downwards, with the median post-crisis distribution equal to the sixth percentile before the crisis. Thus, it appears that ratings declined significantly after the crisis in a way that cannot be explained by observable variables predicting the asset riskiness or other features, such as market structure. This result is consistent with increased investor scrutiny of ratings after the crisis.

Making use of our structurally estimated parameters, we consider several counterfactual scenarios relating to recent policy proposals and policy changes, particularly the Dodd-Frank Act of 2010, which increases regulatory oversight of the credit ratings market, and is currently being contested in Congress. One concern of this oversight is that it will dissuade entry into the ratings market. We simulate distortions under reductions in the number of ratings agencies. We find that going from three to two agencies decreases distortions, both due to ratings catering and ratings shopping. However, going to monopoly (which eliminates distortions) decreases distortions to a much greater degree. Thus, only very substantial increases in reporting requirements would expect to have a large effect through reduced competition. In contrast, we find that regulation that requires agencies to disclose model inputs generally increases distortion by a small amount. By reducing differentiation between firms, ratings catering becomes more intense. Thus, this policy does not look attractive, at least within our framework.

2 Literature Review

This paper relates to several literatures. Most generally, our paper is part of a growing literature on the industrial organization of financial and banking markets. Much of the analysis of the financial crisis of 2007 has been focussed on poor regulation, rather than misguided fiscal, monetary or trade policy (e.g., Stiglitz, 2010; Financial Crisis Inquiry Commission, 2011). Issues of market power, asymmetric information, and conflicts of interest appear central, which has led to new analysis of financial markets by industrial organization economists. Examples are Hortacsu & Kastl (2012), Kastl (2016) and Dewatripont, Rochet & Tirole (2010). Within policy-making institutions, there has also been increased emphasis on industrial organization issues.[3] There is also analysis outside of the formal academic literature along these lines, much of it emphasizing the role of the CMBS market and the credit ratings market, separately or together (for example, see Lowenstein, 2011; Taibbi, 2013). For a discussion of the prevalence and problems with ratings catering, particularly driven by competition, see Securities and Exchange Commission (2008).

A substantial theoretical literature studies incentives in the ratings market. Ratings shopping is taken up in Faure-Grimaud, Peyrache & Quesada (2009), Sangiorgi, Sokobin & Spatt (2009) and Skreta & Veldkamp (2009). If information is opaque around the existence

[3]See the recent special issue in the *Review of Industrial Organization* highlighting research taking place within the Board of Governors of the Federal Reserve (Prager, 2016).

of undisclosed ratings, then issuers can disclose selectively while avoiding a completely adverse inference by the investors (Sangiorgi & Spatt, 2017). Camanho, Deb & Liu (2012) and Bolton et al. (2012) study how competition leads to ratings catering, whereas Mathis, McAndrews & Rochet (2009) show that even a monopoly rating agency has an incentive to cater, insofar as issuers can choose the outside option of not having a rating. Beyond the ratings of securities, there is also a more general issue of the incentives of reviewers, especially in situations in which reviewers are compensated by parties with a financial interest in the outcome of the review. An example of such research in the context of Yelp is Luca & Zervas (2016).

Prior empirical work specifically on the CMBS market emphasizes conflicts of interest, particulary as a result of the issuer-pays model. For example, Cohen & Manuszak (2013) show that the percentage of a CMBS deal that is rated AAA is correlated with instruments proxying for the intensity of competition among rating agencies for that deal. Flynn & Ghent (2014) find that, with the entry of several new CMBS rating agencies very recently, incumbents have responded with laxer ratings.

Similar competitive effects have been documented for ratings on other asset classes including collateralized debt obligations (CDOs) (Griffin & Tang, 2012; Griffin, Nickerson & Tang, 2013), corporate bonds (Becker & Milbourn, 2011), and residential mortgage and other asset-backed securities (Efing & Hau, 2015; He, Qian & Strahan, 2011). Becker & Milbourn (2011) find that Moody's and S&P responded to Fitch's entry by increasing their own ratings. Efing & Hau (2015) find that rating agencies give higher ratings to issuers with which they have a larger bilateral deal volume, presumably motivated by the prospect of future business. Likewise, He et al. (2011) find that, within a given rating category, securities issued by larger issuers are priced at a greater discount, consistent with investors believing that rating catering is more severe when the issuer is large.

An interesting alternative approach to ours is Griffin & Tang (2012) and Griffin et al. (2013). These authors impute agencies' ratings based on the agencies' published methodology and self-reported values for a standardized set of model inputs.[4] The papers compare these imputed ratings to observed ratings to create a measure similar to the distortion we

[4] These input values are analogous to the pre-sale report variables, but are tailored for CDOs rather than CMBS.

study here. In contrast, we use our theoretical model and observed ratings to compute an analogy to their imputed ratings. Thus, while our approach relies on our modeling, our approach is robust to adjustments in agencies' internal methodologies and to the existence of information that the agencies make use of that is not included in the published methodology. We regard their results as corroborating ours. Like us, they find distortions predict worse loan performance, and they also find some competitive effects.

In this paper, we do not model explicitly how the investor market responds to choices of the ratings agencies. That issue is taken up in some theoretical papers such as Bolton et al. (2012) and Mathis et al. (2009). Modeling investor response would substantially complicate the analysis here, especially if we attempted to capture the dynamics of reputation building. Rather, we capture this issue in reduced form through an explicit loss function capturing the cost of distortion to the agencies. A theory paper that takes a similar approach is Bolton et al. (2012).

3 Institutional Background

This section describes the CMBS market and the market for ratings agencies, and highlights issues that guide our empirical modeling. CMBS is an economically significant market, with $677 billion in securities outstanding in 2011. CMBS emerged as a form of financing in the 1990s, and by 2011 accounted for 53 percent ($180B) of funding for all new commercial mortgages.[5] A CMBS deal consists of a pool of mortgages from commercial borrowers. The mortgages may be originated by the issuer or may be purchased from other originators or intermediaries.[6] Each mortgage will collect a stream of principal and interest payments as long as the borrower continues to pay back the loan. An investor in a CMBS obtains the right to collect this stream of payments. The right to collect the stream of payments is divided up into securities that are ranked by seniority. As borrowers stop making payments, the CMBS will stop paying on more junior securities first. The sizes of the various securities and the seniority relationships among them are collectively known as the *deal structure.*

[5]Federal Reserve December 2011 Flow of Funds Tables F.219 and F.220. Number includes $42B in loans backed by government-sponsored enterprises agencies, which our data exclude.

[6]Industry jargon is inconsistent but often uses the term *lead underwriter* to denote what we refer to as the issuer, and the term *issuer* to refer to any loan originator with loans in the deal.

Furfine (2014) provides a more comprehensive discussion of the deal structures.

When evaluating a CMBS deal, credit ratings agencies provide a rating indicating the riskiness of each security in the deal. Ratings take the form of letters, with AAA being the best, followed by others such as AA or Aa2, and BB or Ba2. By convention, deals are structured such that the most senior claim is always rated AAA. For a given pool, increasing the proportion of the AAA security results in the following tradeoff for a ratings agency. On the one hand, investors are willing to pay a premium for AAA insofar as a higher rating on a security signals lower risk and confers regulatory advantages to holding such a security (such as lower bank capital requirements). Thus, all else equal, the issuer earns higher profits if the share of AAA is higher, which makes the issuer likely to select the ratings agency. On the other hand, a higher proportion of AAA decreases the share of the securities subordinate to AAA, reducing the extent to which the AAA investors are protected from defaults. Therefore, assuming investors make the correct inferences, a rating agency that consistently inflates the share of AAA reduces the value signaled by its AAA rating, potentially leading the agency to be less desirable to issuers.

A rating agency has two chief sources of information about the pool. First, the issuer provides a deal prospectus. Second, each agency performs its own discovery process to assess the overall pool as well as the larger individual mortgages within the pool—a practice known as *reunderwriting*.[7] Based on these information sources, the agency models the probability distribution of losses on the entire pool and performs a cashflow analysis of the securities. The loss distribution is very sensitive to assumptions in these models, such as the default probability (see Coval, Jurek & Stafford, 2009). Thus, reunderwriting and modeling can lead ratings agencies to make different assessments of a given CMBS. These features are exacerbated by the relative complexity of CMBS deals in general, as well as the fact that CMBS are a relatively new phenomenon. These features motivate us to model rating agencies as having private information about each deal. As evidence of our private information set-up, we show below that proposed ratings differ across agencies for the same deal.

[7] The practice of reunderwriting the larger loans distinguishes the rating process for CMBS from the rating process for residential mortgage-backed securities (RMBS). Because a typical CMBS pool contains far fewer loans than an RMBS pool, the agencies are able to examine individual loans more closely and, in some cases, physically visit the property locations. See Black, Chu, Cohen & Nichols (2012) for details.

After reunderwriting, the agency then gives the issuer a rating proposal – an indication of how it would rate the securities contingent on alternative deal structures. The industry refers to these ratings proposals as shadow ratings. One element of the shadow rating is the maximum share of the deal that the agency would be willing to rate AAA, or equivalently, the minimum required amount of protection to the AAA security in the form of junior securities.[8] In order to maintain a given level of risk for the AAA security, the riskier the mortgage pool, the lower must be the proportion of AAA. The ratings proposals also include stipulated minimum protection amounts for each of the non-AAA securities in order for the respective security to receive a particular rating, and may require that some mortgages be excluded altogether to achieve a particular ratings structure. Thus, in practice, the final deal structure is determined in conjunction with the ratings process.

In our main model specification, we treat the ratings proposals as a one-dimensional bid corresponding to the maximum share of AAA that the agency would allow (thus ignoring the rating agency's stipulations regarding the junior securities). Although this approach abstracts from some of the issues mentioned above, the AAA share is the most salient feature of the deal structure according to discussions with market participants, and it is typically highly correlated with the stipulated amount of protection for each of the lower-rated securities. In the robustness checks, we consider an alternative formulation that takes into account the structure of both AAA and non-AAA securities.

We assume an exogenous bidder set comprising the agencies S&P, Moody's, and Fitch, based on the fact that all three agencies have historically almost always submitted a proposal whenever a CMBS issuer has had a deal to be rated.[9] Based on the ratings proposals, the issuer chooses one or more rating agencies to publish their ratings. According to testimony before the U.S. Senate Committee on Banking, Housing, and Urban Affairs (2008), the issuer may pay a small fee for the rating proposal, but actually delivering the rating accounts for approximately 90 percent of all fees paid by the issuer to the agency.

[8] During the peak of the real estate bubble, deals often had two AAA-rated securities marketed as *super-senior AAA* and *junior AAA*, respectively, with the latter subordinate to the former. We define the AAA security as the aggregate of *all* claims rated AAA.

[9] U.S. Senate Committee on Banking, Housing, and Urban Affairs (2008). The new entrants DBRS, Morningstar, and Kroll were fringe competitors throughout the study period, and we ignore their presence in our analysis due to their negligible market share. Morningstar, which has a subscription format to provide ratings, is also a data vendor and the source of the main data.

Most typically, the issuer chooses two winners, because the investment guidelines of many institutional investors require ratings from two agencies (see Basel Committee, 2012). The deal is structured in a manner that satisfies the requirements of every chosen ratings agency. In particular, the AAA share is set to the most conservative (lowest) winning proposal. Note that for the *non*-AAA securities, we observe instances of split ratings—that is, cases in which the final rating on a security differs across the agencies chosen to rate the deal. We take up this issue in our robustness checks using non-AAA securities.

For most of the paper, we assume that the issuer picks two winners. Nevertheless, in our data, 7 percent of deals have one winner and 10 percent have three winners. In order to address this, the Appendix studies a model that endogenizes the number of winners. This endogeneity is presumably of limited importance given that we observe two winners for more than 80% of the observations, and not surprisingly, the Appendix finds that a model with an endogenous number of winners leads to similar results.

Our paper focuses on the proposed regulatory requirement to require disclosure of any research that an agency does prior to the selection of a ratings agency. In addition, it is important to recognize that the regulatory reform of the credit agency market continues to be a hotly debated topic. While the SEC has promulgated new transparency rules under its enhanced rule-making mandate, the impact remains to be seen. Furthermore, some of the more dramatic interventions that were envisioned have not been implemented. For instance, the Dodd-Frank Act requires the SEC to issues rules making credit ratings agencies legally liable for their ratings, but the SEC has not done so yet. Rivlin & Soroushian (2017) reviews recent policy towards credit ratings agencies in light of the Dodd-Frank Act of 2010.

4 Data

We use two main data sources. First, we have detailed information from the CMBS data vendor Morningstar on each deal and on each individual mortgage in the deal pool. Second, we hand-collected *pre-sale reports* given by the winning agencies to investors at the time that the securities go on sale. We believe this is the first paper to make use of pre-sale reports. Most importantly for our purposes, the pre-sale report describes the agency's assessment of the mortgage pool based on the reunderwriting process described in the previous section.

Morningstar provides data on 613 deals, which represent 94% of total issuance in this market from 2000-2012.[10] We drop 22 deals for which Morningstar does identify the winning agencies. Thus, we have data on 591 deals. Table 1 summarizes key variables at the deal level.

As shown, the AAA security comprises, on average, 82.8 percent of the pool principal.[11] S&P, Moody's, and Fitch rated 70.1, 70.5, and 58.1 percent of the deals, respectively. Throughout the paper, we drop other ratings agencies from the data set, and proceed as if these were the only ratings agencies. Table 1 also shows deal-level weighted (by loan size) averages for some (but not all) observed characteristics of the 60,748 individual loans in the pools as recorded in the loan origination documents. The loan characteristics are related to the risk of borrower default, either through causal mechanisms—by affecting the borrowers' incentives—or through borrower self-selection into loans with certain characteristics. For a discussion of the importance of specific covariates, the reader is referred to Black et al. (2012). We also report deal-level Herfindahl-Hirschmann Indices (HHI) for various loan characteristics pertaining to the originating firm, property type, geographic distribution, and size of the loans. These concentration indices proxy for the extent of pool diversification.

Figure 1 shows how ratings evolved over time. The share of AAA ratings climbed steadily from 1996 until the crisis, then dipped before the entire market disappeared for more than two years. The market finally emerged with a lower share of AAA ratings, which then continued to decline in our data.

The most important loan characteristics are the debt-service coverage ratio (DSCR) and loan-to-value (LTV) ratio. The DSCR is the ratio of the borrower's monthly rental and other income to the mortgage payments owed, and measures the borrower's ability to make its monthly payments. The LTV is the ratio of the loan amount to the assessed property

[10] Total issuance in Morningstar is $975.3 billion. The website Commercial Mortgage Alert reports total issuance in the US CMBS market during this period is $1037.8 billion (https://www.cmalert.com/rankings.pl?Q=91). The 6% of the market not covered in our data mostly represent distinct products, such as CMBS with bundled floating rate loans. Note that our data does not include any deals backed by the federal agencies (Fannie Mae and Freddie Mac) or CMBS deals bundled into other deals, such as CDOs.

[11] Ratings data are missing for some rated securities. For each such security, we assume that the security is AAA if the total principal amount of other securities that are junior to it, which we always observe, is weakly higher than for any other security in the deal.

value, and measures the borrower's equity in the property. Thus, a lower DSCR or a higher LTV each corresponds to greater risk of default.

Additionally, we observe ex post performance at both the loan- and deal levels (summarized at the deal level in Table 1). At the loan level, we observe the payment history of each mortgage borrower through the censoring date of September 2012. We infer from the payment history whether the loan defaults and, if so, at what point in time. At the deal level, we observe the total amount of written-down pool principal and the shortfall in interest payments from the mortgage borrowers, as of the censoring date.

The right-hand panel of Table 1 summarizes variables that vary across ratings agencies, which vary over both deals and agencies. S&P, Moody's, and Fitch produced pre-sale reports for 92.8, 72.2, and 95.5 percent of the deals that they rated, respectively.[12] These reports provide the agency's private assessment of the DSCR and LTV for the largest individual loans and for the overall pool (as a weighted average by loan size). These reunderwritten DSCR and LTV values are typically more conservative than the corresponding original values—being lower than the original DSCR and higher than the original LTV. For our reunderwritten values, we use what S&P refers to as the "actual" DSCR, what Moody's and Fitch refer to as the "stressed" DSCR, and what all three refer to as the "actual" LTV (S&P, 2009; Moody's, 2000; Fitch, 2008). The reunderwritten DSCR and LTV also differ systematically across agencies. These differences must be interpreted with caution, however. The agencies employ the values in different ways to reach their assessments, so the reunderwritten DSCR and LTV may not be directly comparable either across agencies or with the corresponding original values. In estimation, we include reunderwritten variables from the pre-sale reports as explanatory variables in the bidding functions. To justify this, we demonstrate that these variables are informative about ex post loan performance even after controlling for the original DSCR and LTV (see Table 2).

There is anecdotal evidence that rating agencies lowered their standards in response to events causing them to lose market share.[13] A number of academic studies have found that

[12]The absence of a pre-sale report may indicate that the pool is more "standard" or otherwise more transparent in quality to investors. Indeed, Table 2 shows that the absence of a pre-sale report predicts a lower probability of default on the loans in the pool—consistent with this view.

[13]Popper (2013) writes in the *New York Times*: "The Wall Street ratings game is back. Five years after inflated credit ratings helped touch off the financial crisis, the nations largest ratings agency, Standard &

agencies with low recent market share tend to produce higher ratings for CMBS (Cohen & Manuszak, 2013; Flynn & Ghent, 2014) and for subprime mortgage-backed securities (Sun, Li & Tan, 2013). In order to study this in our model, the bidder-deal-specific variables also include several alternative measures of incumbency—that is, each bidder's market share among previously issued deals. We find similar results to the previous literature in our data in Appendix A, which presents the reduced-form relationship between the AAA share for a deal and various covariates, including the recent market share of the winning bidders. Thus, we specify bidder utility to allow for such dependence. To be clear, we allow for incumbency to affect the current value of winning, but we do not model the dynamics of how ratings agencies might make proposals that account for the value of incumbency in the future. While potentially interesting, that is outside the scope of our paper and we are not aware of any papers in the literature that have explored this.

While we rely on pre-sale reports, that data is missing for a number of observations. In particular, there are 55 deals with no pre-sale reports by any of the rating agencies (52 in the estimation sample after dropping a small number of deals for reasons discussed elsewhere in the paper). We use these observations in estimation but include a control variable to account for possible nonrandomness in when the pre-sale reports are missing. Here, we briefly describe these observations. Compared with the summary statistics for the full sample of deals (Table 1), these deals tend to be smaller, with an average pool principal of $734 million (standard deviation of $1013 million) versus an average pool principal of $1591 million for the full sample. All other deal-level characteristics have means that are very similar to – or at least not significantly different from – the corresponding means for the full sample. For example, the share of the AAA security is on average 0.844 (versus 0.828); the original DSCR and original LTV are on average 1.527 and 0.672, respectively (versus 1.507 and 0.677 for the full sample). The ex post performance of these deals is also not significantly different from that of the full sample: for deals with no presale reports, the average principal writedown and interest shortfall as of September 2012 (as a proportion of the original principal) is 0.009 (standard deviation of 0.018), versus 0.0243 for the full sample.

Poors, is winning business again by offering more favorable ratings."

5 Model

In our model, ratings agencies have private signals about the riskiness of a financial product. Agencies simultaneously propose ratings to an issuer. The issuer selects the most attractive ratings, leading to ratings shopping. In response, agencies offer higher ratings than they otherwise would, leading to ratings catering. In equilibrium, we observe high ratings relative to agency signals, which we refer to as distortion. A complexity is that agencies have private information, and that information is correlated across agencies and so agencies use their signal to infer the signals of their rivals.

For now, we assume there is a single issuer. The issuer solicits ratings from ratings agencies indexed by j, $j = 1, \ldots, J$, representing S&P, Moody's, and Fitch, so $J = 3$. Agency j places a proposal (or bid) $b_j \in [0, 1]$ that states the maximum share of the AAA security (as a percentage of total pool principal) that j would allow. For example, a choice of $b_j = 0.7$ states that j would give a AAA rating to a security backed by no more than 70 percent of the pool, meaning that AAA investors would suffer losses only if more than 30 percent of the repayment were lost.

Given a bid profile $b \equiv \{b_j\}_{j=1,\ldots,J}$, the issuer chooses a subset of agencies to publish its ratings. We denote this choice by $d \equiv \{d_j\}_{j=1,\ldots,J} \in \{0, 1\}^J$, a $J \times 1$ vector of ones and zeros indicating the set of winners. Let $B(d)$ be the set of bids such that $d_j = 1$, so $B(d) = \{b_j : d_j = 1\}$. The issuer must set the AAA security's proportion of the pool to $\min(B(d))$. Thus, the lowest winning rating is the *pivotal* rating. In the base specification, we assume that the number of winning agencies is exogenously determined.[14] For an auction with K observed winners, the issuer therefore chooses the K highest bidders. Formally, the issuer solves:

$$d^*(b) \equiv \{d_j^*(b)\}_{j=1,\ldots,J} = \arg\max_{d} \left\{ \min B(d) : \sum_{j=1}^{J} d_j = K,\ d \in \{0, 1\}^J \right\}. \tag{1}$$

Agency j has type $\tilde{b}_j \in [0, 1]$, which represents the bid the agency would make if the agency was assured that it would be selected by the issuer. Thus, it represents the agency's evaluation of the risk of the issuer and the cost of serving the issuer. The type $\tilde{b}_j$ is made

[14] Appendix D presents a specification that endogenizes the number of winners. We find similar results.

up of the following elements. We assume that agencies receive a private signal s_j of the riskiness of the pool, and the agencies have a publicly observed cost shifter w_j, in particular the incumbency of the agency. Denote $w \equiv \{w_j\}_{j=1,\ldots,J}$, the $J \times 1$ vector of elements of w_j. We assume that $\tilde{b}_j = \Phi(s_j + w_j\beta_1)$. Here, Φ is the standard normal CDF, and we use it as a convenient functional form that maps the real valued term $s_j + w_j\beta_1$ into the unit interval.

The signal s_j is made up of several elements. Deal characteristics x are publicly observable. Agency j also observes z_j, which becomes public after the auction if j wins. Thus, z_j are private to the agencies at the time of bidding but z_j of winners are observed by the econometrician. These characteristics are the reunderwritten DSCR and LTV for the mortgage pool. In addition, agencies observe a private signal u_j that is not observed by the econometrician. We specify s_j as $s_j = x\beta_2 + z_j\beta_3 + u_j$. Thus:

$$\Phi^{-1}(\tilde{b}_j) = \beta_1 w_j + \beta_2 x + \beta_3 z_j + u_j. \tag{2}$$

We assume $E[u_j|w, x, z_j] = 0$, and that the joint distributions of z_j and u_j are common knowledge to the agencies. We assume further that the agencies' utilities do not depend on any common elements that are unobserved by the agencies. However, we allow u_j to be correlated across agencies. Thus, in the parlance of auctions, this is an affiliated private values (APV) setting. We discuss this assumption in Section 8.[15]

The cost to the agency of winning is $c(b_j, \tilde{b}_j)$ and is increasing in the difference between b_j and $\tilde{b}_j$. In practice, we specify cost in the real-valued normalization:

$$c(b_j, \tilde{b}_j) = \left(\Phi^{-1}(b_j) - \Phi^{-1}(\tilde{b}_j)\right)^2 = \left(\Phi^{-1}(b_j) - w_j\beta_1 - x\beta_2 - z_j\beta_3\right)^2.$$

We assume $b_j \geq \tilde{b}_j$, which as we discuss shortly, is not binding in practice. Thus, the significance of the squared term is to impose an increasing marginal cost on increasing b_j above $\tilde{b}_j$.

When the issuer selects an agency to provide a rating, the agency receives a payoff

[15] We take the z_j to be given exogenously. We might worry that z_j are chosen strategically in practice. However, we believe that our assumption is reasonable assumption because these variables, the reunderwritten values of DSCR and LTV from the presale reports, are released after the auction is over. Their strategic value could be part of a relatively complex dynamic game, but we do not take that on here.

denoted by V. Overall, agency j's payoff given type $\tilde{b}_j$ and bid profile b is:

$$\pi_j(b, \tilde{b}_j) = d_j(b)[V - c(b_j, \tilde{b}_j)] \tag{3}$$

We normalize the value of not being chosen to 0, we assume V is high enough that bidders always participate. Agencies simultaneously choose bids, and must take expectations over rival bids b_{-j} because agencies do not observe their rivals' private information. We use Bayesian Nash Equilibrium as our solution concept. Optimal bidding by agency j conditions on j's information set, and is characterized by the following first-order condition:

$$\frac{\partial E_{b_{-j}}\left[\pi_j(b, \tilde{b}_j)|x, w, z_j, u_j\right]}{\partial b_j} = 0 \tag{4}$$

Because of the quadratic cost term, the first-order condition has two roots. We pick the root that satisfies the inequality $\Phi^{-1}(b_j) - s_j - w_j\beta_1 > 0$, which always exists in practice. Thus, we assume that ratings agencies never rate below their type. In auction parlance, this assumption is akin to assuming that sellers bid above cost.

Because we allow the private terms z_j and u_j to be correlated across agencies bidding on the same security, agencies should use their own draws to infer the draws of their rivals. This feature is implicit in Equation 4 (in the expectation over rival bids), and we account for it in our numerical solutions to equilibrium outcomes described below.

Also, it is important to recognize that our empirical method will infer $\tilde{b}_j$ from bids b_j (as well as other covariates and the model). Even if $\tilde{b}_j$ is poorly chosen, for instance, if agencies use flawed ratings models, our method is robust to the extent that we correctly model the choice of b_j conditional on $\tilde{b}_j$.

6 Identification and Estimation

This section presents our estimation technique, and specific functional forms that we impose on the model. Before doing so, we provide a discussion of the identification of the parameters, and we argue that significant elements of the model are non-parametrically identified, which suggests that functional form assumptions are not driving our results. The estimation technique is closely related to the identification argument, but these sections can be read separately.

One approach to estimation would be a "full-solution" method where we solve for the equilibrium outcome at a given set of parameters by finding the choices that simultaneously solve all of the best-response functions. However, these types of techniques are numerically challenging to the point of being infeasible, both in our setting and many other auction and oligopoly settings. Instead, for both identification and estimation, we utilize a two-step method as in Guerre et al. (2000), Bajari, Hong, Krainer & Nekipelov (2010) and Bajari, Benkard & Levin (2007). In the first step, we use reduced-form estimation to find the distribution of rival choices that agencies expect to face. In the second stage, we structurally estimate the best response function, using the first-stage results to proxy for rival choices. This approach is attractive because it circumvents having to solve for an equilibrium. The approach requires that the first stage provides a good approximation for the beliefs that agencies have over their rivals. That is, outcomes in the data should be a good approximation for equilibrium rational expectations. Thus, showing that the approach is non-parametrically identified is particularly attractive.

6.1 Identification

In order to identify the structural parameters, we first construct a reduced-form model of the bid distribution that the agencies expect to face. Then, based on the expected bid distribution, we can use the first-order condition implied by the assumption of a Bayesian Nash Equilibrium (Equation 4) to identify the structural parameters and the unobserved term that rationalizes each observed bid. Thus, with sufficient variation in bids and explanatory variables, we can identify the structural parameters and the non-parametric distribution of the unobserved terms.

6.1.1 First-stage identification

A major challenge in this approach is to identify the distribution of bids that the agencies expect. In order to proceed, we assume that the agencies' expected distribution of bids equals the observed distribution of bids, which is a form of a rational expectations assumption. By itself, this assumption is not enough, because we observe only the winning bid and we need to know the distribution of all bids the agencies face.

Our approach is to adapt Theorem 5 in Athey & Haile (2002) from the case of identifying the distribution of values in a second-price auction to the case of the distribution of bids in our setting. They show how to identify the distribution of values in a second price auction when the researcher observes only the transaction price and sufficient variation in bidder characteristics. Their approach relies on the result that agents bid their values in a second price auction. Being able to identify the distribution of values when the researcher observes only the second highest value is similar to identifying the distribution of bids when the researcher observes only the second highest bid.

We first specify a non-parametric reduced-form equation of the bids:

$$b_j = \Phi\left(m_j(x, w, z_j) + \varepsilon_j\right), \quad j = 1, \ldots, J. \tag{5}$$

The unobserved terms ε_j are distributed according to some unknown joint CDF $H(\varepsilon_1, \varepsilon_2, \varepsilon_3)$. The function Φ is a known monotone transformation that we choose, without loss of generality, to be the standard normal CDF.[16] The normal CDF is a convenient functional form that constrains the bids to be between 0 and 1. Equation 5 is non-parametric in the sense that $m_j(\cdot)$ and $H(\dot{)}$ are unconstrained, although we assume $H(\dot{)}$ does not depend on other observed or unobserved variables (that is, the vector ε is distributed *iid*).

Let the vector of dummies $d^* \equiv \{d_j^*\}_{j=1,\ldots,J}$ indicate the set of winners and let b^* denote the pivotal bid, both of which are observed. The function $m_j(\cdot)$ is identified by the observed relationship between b^* and the covariates (x, w, z_j) for winning bidders j.

The intuition for how $H(\cdot)$ is identified when there are two winners is as follows. Suppose, to begin with, that we observe all bidder-specific covariates (including z_j) for all bidders j. The distribution of $(\varepsilon_1, \ldots, \varepsilon_J)$ is traced out by the empirical probability of observing each potential combination of winning bidders d^*, conditional on the other observables $\{b^*$, x, w, $z\}$. For instance, suppose we wish to identify $H(\cdot)$ at particular points $\{a_1, a_2, a_3\}$. We find observations where the values of $\{b^*, x, w, z\}$ are such that a draw of $\{\varepsilon_1, \varepsilon_2, \varepsilon_3\}$ equal to $\{a_1, a_2, a_3\}$ would cause each bidder to bid exactly b^*, according to Equation 5. Then,

[16] We believe the additive separability of ε_j is not a critical assumption.

the empirical probability that bidders 1 and 2 win and 3 loses provides the probability that $\varepsilon_1 \geq a_1$, $\varepsilon_2 \geq a_2$ and $\varepsilon_3 < a_3$.

As an example, suppose we wanted to know $H(0,0,0)$, and we had a set of observations with values of x, w, z such that $m_j(x, w, z_j) = 0$ for all j. For these values and draws of $\varepsilon_1 = 0, \varepsilon_2 = 0, \varepsilon_3 = 0$, Equation 5 would predict each bidder bids 0.5, so would use $b = 0.5$. At this point in our data, we check the probability of seeing bidder 1 and bidder 2 win and bidder 3 lose, which gives us the probability that $\varepsilon_1 \geq 0$, $\varepsilon_2 \geq 0$ and $\varepsilon_3 < 0$. Similarly, checking for the probability of other combinations of winners provides the probability of other outcomes evaluated at $\{0,0,0\}$. Repeating this for other combinations of $\{b^*$, x, w, $z\}$ allows us to evaluate the CDF at other points.

More formally, the probability of bidders 1 and 2 winning and 3 losing evaluated at $\{b^*$, x, w, $z\}$ is:

$$\begin{aligned} \Pr(d_1^* = d_2^* = 1,\ d_3^* = 0 | b^*, x, w, z) &= \Pr(\varepsilon_1 \geq a_1, \varepsilon_2 \geq a_2, \varepsilon_3 < a_3). \qquad (6) \\ \texttt{where } a_1 &= \Phi^{-1}(b^*) - m_1(x, w, z_1), \\ a_2 &= \Phi^{-1}(b^*) - m_2(x, w, z_2), \\ a_3 &= \Phi^{-1}(b^*) - m_3(x, w, z_3). \end{aligned}$$

which is the CDF of $(-\varepsilon_1, -\varepsilon_2, \varepsilon_3)$ evaluated at the point $(-a_1, -a_2, a_3)$. With sufficient variation in d^*, x, z and w, the entire joint distribution of $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ is identified. To evaluate the CDF at any point a_1, a_2, and a_3, we find values of b, x, w and z in our data such that they satisfy the condition upon which the conditional probability in Equation 6 depends, and then we check in our data the probability of each outcome d. Identification based on auctions with a single winner is similar. The case with three winners also contributes to identification because we observe all the bids in that case, so $\Pr(d^*, b^* | x, w, z)$ is the probability of observing $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ that generates the three bids.

As in Athey & Haile (2002), this approach relies crucially on the presence of variation across bidders in the covariates, which comes from w and z. A complication in the data is that z is unobserved for the losing bidders, as well as for a small number of winning bidders that do not produce a pre-sale report. Instead, we integrate out elements of z that are unobserved. Let $G_z(\cdot)$ represent the distribution of z. We estimate $G_z(\cdot)$ from the subset of auctions in which all three agencies win, in which case each element of z is observed. Thus, we can rewrite Equation 6 to integrate out unobserved elements of z using $G_z(\cdot)$. In this case, Equation 6 becomes:

$$
\begin{aligned}
& \Pr(d_1^* = d_2^* = 1,\ d_3^* = 0 | b^*, x, w, z_1, z_2) \\
= & \int \Pr(\varepsilon_1 \geq a_1, \varepsilon_2 \geq a_2, \varepsilon_3 < \Phi^{-1}(b^*) - m_3(x, w, z_3))\, dG_z(z_3|z_1, z_2) \\
= & \int \tilde{H}(-a_1, -a_2, \Phi^{-1}(b^*) - m_3(x, w, z_3))\, dG_z(z_3|z_1, z_2),
\end{aligned}
$$

such that a_1 and a_2 are defined as in Equation 6, where $\tilde{H}$ is the CDF of $(-\varepsilon_1, -\varepsilon_2, \varepsilon_3)$. Thus, this is a condition on the integral over the third element of $H(\cdot)$ rather than a condition on $\tilde{H}(\cdot)$ evaluated at a particular point. This expression has the form of the non-parametric conditional mean model of Newey & Powell (2003). As shown in that paper, $\tilde{H}$ is identified provided the conditional distribution $G_z(z_3|z_1, z_2)$ satisfies the "completeness" property. Lehmann & Romano (2005) shows that a sufficient condition for completeness is for $G_z(z_3|z_1, z_2)$ to belong to the very general exponential family of distributions (which includes the normal distribution, our assumption in estimation). Based on variation in w, which is always observed for all bidders, the distribution of ε is still nonparametrically identified in the way described above.

The estimation of $m_j(x, w, z_j)$ and the distribution of ε generates the distribution of bids $G_b(\cdot|x, w, z)$. We let $\theta = \{G_z(\cdot), G_b(\cdot|x, w, z)\}$ refer to the first-stage parameters. The parameters θ allow us to characterize the distribution of rival bids expected by the agencies.

6.1.2 Second-stage identification

We now turn to identifying structural parameters, which consist of $\beta = (\beta_1, \beta_2, \beta_3)$ and the distribution of u. Formally, $G_z(\cdot)$ is also a structural parameter, but that has been identified in the first stage. The parameters β are identified by the empirical relationship between the AAA share and x, z_j and w_j for the winning bidders. In particular, with the distribution of rival bids previously identified, our model and our assumption of a Bayesian Nash Equilibrium creates a one-to-one mapping between the observed bid b_j and the type

$\tilde{b}_j$. Then, the parameters β are identified from Equation 2.

We follow Guerre et al. (2000) and Athey & Haile (2007) in order to identify the distribution of u. Bidder j's first-order condition (4) implicitly defines bidder j's unobserved term u_j as a function of the bid and other elements observable to j, which we denote as follows:

$$u_j = \xi_j(b_j, x, w, z_j, \theta, \beta). \tag{7}$$

Let the function $\xi(b, x, w, z, \theta, \beta)$ return the $J \times 1$ vector u with elements $\{u_j\}_{j=1,2,3}$ as a function of the vector of bids and other variables that are observable to at least one agency. Denote the joint distribution of u by $F(u)$. For a particular vector of values $u' \in R^J$:

$$\begin{aligned} F(u'|x, w, z; \theta, \beta) &= Pr(u < u' \mid x, w, z, \theta, \beta) = Pr(\xi(b, x, w, z, \theta, \beta) < u') \\ &= Pr(b < \xi^{-1}(u', x, w, z, \theta, \beta)) = G_b(\xi^{-1}(u', x, w, z, \theta, \beta)) \end{aligned}$$

where $\xi^{-1}(u', x, w, z, \theta, \beta)$ is the inverse function of $\xi(b, x, w, z, \theta, \beta)$ with respect to the first argument.

This completes our discussion of identification. For completeness, note that our model does not have unobserved auction-level heterogeneity—that is, a component of u_{ij} commonly observed by market participants but not by the researcher. This simplification is necessary because, without further assumptions, the observed bid distribution can be rationalized by either our model's assumption of affiliated private values (APV) or by independent private values (IPV) with unobserved heterogeneity (see Athey & Haile, 2007), or by intermediate cases between these two extremes. Krasnokutskaya (2011) explores these assumptions further. Rather than allow commonly observed elements of u_{ij}, our APV approach allows u_{ij} to be correlated across bidders, and bidders condition on their draw of u_{ij} in predicting the bids of rivals.

6.2 Estimation

The estimation procedure closely mirrors the identification arguments, and falls under the general class of two-step estimators (e.g., Bajari et al., 2007, 2010). In the first step, we estimate $\theta \equiv \{G_b(\cdot), G_z(\cdot)\}$, the joint distribution of the equilibrium bids and the joint distribution of z_j (the reunderwritten DSCR and LTV). In the second step, we recover the distribution of types. An important assumption for these types of methods is that if there are multiple equilibria in the structural model, only a single equilibrium is being played in the data. Otherwise, we cannot identify structural estimates (see Bajari et al., 2007).

In what follows, we now add an index i to indicate observations $i = 1, \ldots, n$, with each CMBS deal representing a separate observation. We first specify several functional forms, and then estimate θ by Maximum Likelihood. For estimation, we modify expression (5) by assuming the following quasilinear form:

$$b_{ij} = \Phi\left(\{x_i, w_i, z_{ij}\} \cdot \gamma_j + \varepsilon_{ij}\right) \tag{8}$$

This rather parsimonious specification is motivated by the modest data size. We further restrict the coefficients on x_i to be the same across bidders, and assume $(\varepsilon_{i1}, \varepsilon_{i2}, \varepsilon_{i3}) \sim N(0, \Omega_\varepsilon)$. We assume that $\{z_{ij}\}_{j=1,2,3}$ is jointly normally distributed with $(z_{i1}, z_{i2}, z_{i3}) \sim N(\mu_z, \Omega_z)$. Thus, we can reexpress θ as the set of parameters $\theta \equiv (\{\gamma_j\}_{j=1,2,3}, \Omega_\varepsilon, \mu_z, \Omega_z)$. In the results, we discuss how we are able to fit the data rather well even with these parametric restrictions.

Let b_i^* denote the observed pivotal bid for observation i; let z_i^* and $\tilde{z}_i$ denote the collections of all observed and unobserved pre-sale report variables, respectively, such that $z_i = \{z_i^*, \tilde{z}_i\}$.[17] Let $g_{b_j}(\cdot|x_i, w_i, z_i, \gamma_j, \Omega_\varepsilon) \equiv \frac{d}{db_j} G_b(\cdot\ |x_i, w_i, z_i, \gamma_j, \Omega_\varepsilon)$ denote the marginal density of bidder j's bid; let $\frac{d}{dz_i^*} G_z(z_i^*|\mu_z, \Omega_z)$ denote the marginal density of $G_z(\cdot\ |\mu_z, \Omega_z)$ with respect to z_i^*; and let $\frac{d}{d\tilde{z}_i} G_z(\tilde{z}_i|z_i^*, \mu_z, \Omega_z)$ denote the conditional marginal density of

[17] Recall that, in the data, some winning bidders do not produce a pre-sale report. Thus, z_{ij} is a latent variable that is observed only for winning bidders who produce a report. We must control for possible nonrandomness of when a winning bidder produces a report. For the case of a winning bidder j that does not produce a pre-sale report, we set $z_{ij} = 0$ in (8) and allow for a separate, bidder-specific, intercept term expressly for such cases. In the parameter estimates shown in Table 3, this intercept term is listed as "Bidder produced no pre-sale report."

$\tilde{z}_i$ given z_i^*. We estimate θ by maximizing the likelihood:

$$
\begin{aligned}
L(\theta|d^*, b^*, x, w) \;\; = \;\; \sum_i \log \Bigg[\sum_{j: d^*_{ij}=1} \int_{z_i^*} \;\; & g_{b_j}(b_i^* | x_i, w_i, z_i, \gamma, \Omega_\varepsilon) \\
& \Pr(b_{ij'} \geq b_i^* \;\; \forall j' \text{ s.t. } j' \neq j \;\&\; d^*_{ij'} = 1 \, |x_i, w_i, z_i, \gamma, \Omega_\varepsilon) \\
& \Pr(b_{ij'} < b_i^* \;\; \forall j' \text{ s.t. } d^*_{ij'} = 0 \, |x_i, w_i, z_i, \gamma, \Omega_\varepsilon) \cdot \\
& \quad \tfrac{d}{d\tilde{z}_i} G_z(\tilde{z}_i | z_i^*; \mu_z, \Omega_z) \tfrac{d}{dz_i^*} G_z(z_i^* \; \mu_z, \Omega_z) dz_i^* \Bigg]
\end{aligned}
\tag{9}
$$

Above, the inner sum is over the set of all possible pivotal bidders j. Within the inner sum, we evaluate the marginal density of possibly pivotal bidder j's bid distribution evaluated at b_i^*, the joint probability of the bids for all winning but non-pivotal bidders, and the joint probability of the bids for all losing bidders, all integrated over the conditional distribution of $\tilde{z}_i$ given z_i^*. We then integrate over z_i^* to create a likelihood function. Because the likelihood function depends on random variables that are jointly normally distributed (ε and z, with independence between ε and z), we can compute the likelihood function analytically. Thus, we use traditional Maximum Likelihood to compute θ.

We now turn to estimating the structural parameters, the second stage of estimation. Observed pivotal bid b_i^* by bidder j is rationalized by a shock $u_{ij}^* = \xi_{ij}(b_i^*, x_i, w_i, z_{ij}, \theta, \beta)$. In practice, the function $\xi_{ij}(b_i^*, x_i, w_i, z_{ij}, \theta, \beta)$ is calculated by solving Equation 4. Based on this solution, if all auctions had one winner, it would be straightforward to calculate u_{ij}^* for each auction and then to construct moments. However, for auctions with more than one winner, we do not observe which winner made the pivotal bid. For auctions with two or three winners, we construct the expectation of the pivotal bidder's shock given observed auction outcomes. Specifically, we define:

$$
\begin{aligned}
\hat{u}_i^* \;\; & \equiv \sum_{j:\, d^*_{ij}=1} \xi_{ij}(b_i^*, x_i, w_i, z_{ij}, \theta, \beta) \Pr(j \text{ pivotal in } i | \; d_i^*, b_i^*, \theta, \beta, x_i, w_i, z_i^*) \\
& = \sum_{j:\, d^*_{ij}=1} \xi_{ij}(b_i^*, x_i, w_i, z_{ij}, \theta, \beta) \Pr(b_{ij} = b_i^*, b_{ij'} > \geq b_i^* \; \forall j' \text{ s.t. } j' \neq j \;\&\; d^*_{ij'} = 1 \mid d_i^*, b_i^*, \theta, \beta, x_i, w_i, z_i^*),
\end{aligned}
\tag{10}
$$

where the probability inside the summation is the probability of j being the pivotal bidder, conditional on the observed set of winning bidders, as implied by the first-stage estimation. Specifically, $\Pr(b_{ij} = b_i^*, b_{ij'} \geq b_i^* \; \forall j' \text{ s.t. } j' \neq j \;\&\; d_{ij'}^* = 1 \mid d_i^*, b_i^*, \theta, \beta, x_i, w_i, z_i^*)$ equals:

$$\frac{\displaystyle\int_{\tilde{z}_i} g_{b_j}(b_i^*|x_i, w_i, z_i, \gamma, \Omega_\varepsilon) Pr(b_{ij''} \geq b_i^* \; \forall j'' \text{ s.t. } j'' \neq j \;\&\; d_{ij''}^* = 1 \, |x_i, w_i, z_i, \gamma, \Omega_\varepsilon) dG_z(\tilde{z}_i|z_i^*)}{\displaystyle\sum_{j' \text{ s.t. } d_{ij'}^* = 1} \left[\int_{\tilde{z}_i} g_{b_{j'}}(b_i^*|x_i, w_i, z_i, \gamma, \Omega_\varepsilon) Pr(b_{ij''} \geq b_i^* \; \forall j'' \text{ s.t. } j'' \neq j' \;\&\; d_{ij''}^* = 1 \, |x_i, w_i, z_i, \gamma, \Omega_\varepsilon) dG_z(\tilde{z}_i|z_i^*) \right]}$$

We determine this probability by simulating over $\tilde{z}_i$ and ε.

Recall that we assume $E[u_i|x_i, w_i, z_{ij}] = 0$, which implies $E[u_i^*|x_i, w_i, z_{ij}] = 0$. We use this condition to estimate β by means of the generalized method of moments. We construct standard errors using the bootstrap to take into account the first-step estimation error. In our bootstrap, we select auctions at random with replacement to construct bootstrap samples. Thus, we use the block bootstrap to account for within-auction correlation.

Having identified the β parameters, we can obtain the entire joint distribution of agencies' types by evaluating the expression (7) while simulating over the estimated equilibrium bid distribution. More specifically, for each auction i, we condition on common information in that auction (x_i and w_i) and draw from the joint distribution of the private information terms z_i and bids b_i. The distribution of z_i given x_i and w_i is determined by the parameters μ_z and Ω_z, and the conditional distribution of b_i given x_i, w_i, z_i is determined by the parameters γ and Ω_ε, according to expression 8.

Lastly, we discuss the measurement of incumbency. We consider three alternative measures for a bidder's incumbency status, as summarized in Table 1. The first is the share of deals that a bidder j rated among the ten most recent (by issuance date) deals. The second is the share of deals that a bidder rated among the three most recent deals issued by the same bank as the one issuing deal i (e.g., if the bank issuing deal i is JPMC, this measure is bidder j's share of the three most recent deals issued by JPMC). The latter measure captures incumbency effects that are relationship-specific as opposed to general.

A potential concern is whether bidder j's type is serially correlated over auctions i through the residual u_{ij}, which would violate the orthogonality conditions on w_{ij} that we use in estimation. This concern is at least partially mitigated by the fact that we control for year effects in our deal covariates x_i. Furthermore, Appendix A discusses several tests in which we fail to reject the null hypothesis of zero autocorrelation.

Nonetheless, we consider a third measure of incumbency as a robustness check. Intuitively, it seems more plausible for the bidder's beliefs to be correlated over deals issued by a particular bank (for example, due to persistence over time in an individual bank's reputation) than for those beliefs to be correlated across deals by unrelated issuers. We therefore define the third measure of incumbency to be the share of deals that bidder j rated from among the ten most recent deals issued by banks *other than* the one issuing the current deal.

7 Results

We discuss the first-step reduced-form estimates, followed by the second-step structural estimates. In the third subsection, we use the results to construct a measure of the distortion, or the difference between the bidder's belief and their bid, and we compare how that evolved with the financial crisis.

7.1 First-step estimates

Tables B.1 and B.2 in Appendix B report first-step estimates for the base specification. Table B.1 shows the joint distribution of the weighted average reunderwritten DSCR and LTV for each deal (z_{ij}).

The estimates of the means (μ_z) show that S&P's mean reunderwritten DSCR and LTV are each typically higher than those of the other agencies, consistent with the differences across agencies in the raw sample means (Table 1). The covariance (Ω_z) estimates indicate that the reunderwritten DSCR and LTV are strongly correlated across agencies, especially the reunderwritten LTV. We also allow for correlation between DSCR and LTV.

Table B.2 shows estimates for the joint distribution of equilibrium bids, parameterized by $\{\gamma_j\}_{j=1,2,3}$ and Ω_ε. As expected, when DSCR is significant (Moody's), it is positive, and when LTV is significant (S&P and Fitch), it is negative. The insignificant results could be driven by the strong positive correlation between LTV and DSCR in the data. The reunderwritten DSCR and LTV net of the original DSCR and LTV have a correlation of 0.7 for Moody's and Fitch and 0.3 for S&P (not shown in tables).

For parsimony, we let bidder j's bid depend on its own incumbency measures and the *average* of j's competitors' incumbency measures. Interpretation of the coefficients is complicated by the fact that a bidder's own incumbency is, by construction, negatively correlated with that of its competitors. The overall implication of the estimates is that bidders submit higher bids when their recent market share is high relative to that of their competitors. The effect is weaker when the most recent deals were issued by the same bank as the issuer of the current deal.[18]

The estimated effects of the deal covariates x_i show that bids are lower for pools containing more loans with balloon payments and, conditioning on the size of the loan, when many properties serve as collateral (cross-collateralization). Higher pool concentration by originator, property type and region are each associated with lower bids. The vintage fixed effects indicate an increasing trend in the bids from the year 2000 until around 2006. Over this time period, the trend implies an average increase of 13.7 percentage points in the AAA share prescribed by the bids.[19] However, this trend reversed after 2006 with the onset of the financial crisis. Finally, the estimates of the covariance matrix Ω_ε indicate a high degree of correlation across bidders in the residual of the equilibrium bid function (ε_{ij}).

Identifying the structural parameters requires having first-step estimates that fit the data well. Tables B.3 and B.4 in Appendix B demonstrate this to be the case by presenting measures of within- and out-of-sample fit. As shown, our estimates predict winners at about the correct rate, and we are relatively close on the winning bid. This is reassuring, especially given that we have specified the first stage somewhat parsimoniously. The out-of-sample fit is only somewhat worse when we estimate using only half the sample, using the other half for model validation. This gives us reassurance that we are not overfitting the data.

[18] These variables are not individually significant in our results, although their joint significance tends to be satisfied at least at 85% confidence levels across specifications. Dropping bidder fixed effects leads the incumbency measures to be very highly significant, but overall, we felt that fixed effects were appropriate.

[19] This number comes from evaluating the impact of the 2006 dummy on the bid, evaluated at the sample mean for the AAA proportion (0.828). More formally, $13.7\% = \Phi(\Phi^{-1}(0.828) + 0.657) - 0.828$, where 0.657 is the estimated coefficient for the 2006 vintage (relative to 2000).

7.2 Structural estimates

Table 3 reports estimates of β, the coefficients for the bidder's type $\tilde{b}_{ij}$. The table reports estimates for the base specification and for an alternative specification, which uses a different measure of incumbency (the vector w_{ij}). The first-step estimates corresponding to the alternative specification are reported in Appendix Tables D.1 and D.2.

The base specification proxies for bidder j's incumbency using j's share of the ten most recent deals and the share of deals that j rated from among the three most recent deals issued by the same bank as the issuer of deal i. The alternative specification proxies for bidder j's incumbency using the share of deals that bidder j rated from among the ten most recent deals issued by banks *other than* the one issuing the current deal. As we argue in subsection 6.2, the latter definition is robust to serial correlation in the unexplained component of bidders' beliefs (i.e., that which remains after controlling for year effects) about deals issued by a particular bank. These different specifications for the variables associated with β_1 still lead to similar estimates of β_2 and β_3. Importantly, the correlation of the privately observed element of bidder type ($\beta_3 z_j + u_j$) between specifications is 0.8622 (not shown in table). Thus, the choice of specifications does not substantially affect the post-estimation exercises that we shall discuss below.

Some of the individual parameters among the vectors β_1, β_2, and β_3 are imprecisely estimated, but explanatory variables as a group are jointly significant. The vintage dummies indicate that bidders' types peaked in 2006, just before the financial crisis. The bidder-deal-specific covariates are jointly significant, but the individual coefficients must be interpreted cautiously given the high degree of correlation among incumbency measures and between the reunderwritten DSCR and LTV.

Table 4 reports the covariance matrix for the estimated joint distribution of u_i, the idiosyncratic components of $\tilde{b}_{ij}$. As expected, the idiosyncratic component u_{ij} is positively correlated across bidders.

7.3 Distortions

We now construct a measure of the amount that bids are distorted away from the agencies' beliefs. Ideally, we would compare the pivotal bid to the pivotal agency's belief about the ideal bid. However, we cannot observe the identity of the pivotal agency when there is more than one winner, so we instead construct the *expectation* of the pivotal agency's belief given observable variables. Specifically, we constructed the expected z_{ij} of the pivotal bidder, similar to Equation 10. Let $\hat{z}_i^*$ and s_i^* be:

$$\begin{aligned}
\hat{z}_i^* &\equiv \sum_{j:\, d_j^*=1} z_{ij} \Pr(j \text{ pivotal in } i |\, d_i^*, b_i^*, \theta, \beta, x_i, w_i, z_i^*), \\
s_i^* &= x_i\beta_2 + \hat{z}_i^*\beta_3 + \hat{u}_i^*.
\end{aligned}$$

We define the *distortion* for a deal i as:

$$\lambda_i = \Phi^{-1}(b_i^*) - s_i^*,$$

i.e., the normalized difference between the pivotal bid and the expectation of the pivotal bidder's type given observed auction outcomes. This measure captures the extent of distortion embodied in the equilibrium share of AAA for deal i. We construct λ_i ignoring the role of w_i in driving bids away from beliefs. We do so because w_i differs across agencies and auctions, and would make λ_i more difficult to interpret. The term λ_i can be thought of as the distortion holding w_i at zero.

Figure 2 plots the distortion against the date of issuance for each deal. The gap between 2009 and 2010 reflects the absence of new CMBS issuance during the crisis period. As shown, distortion amounts decreased following the crisis, with the median distortion during the post-crisis period corresponding to the sixth percentile of the pre-crisis distribution.

To get a sense of the magnitude of this change, consider that the median distortion (λ_i) before the crisis is 0.926, whereas the median distortion after the crisis is 0.759, a decrease of 0.167. Evaluated at the sample-mean AAA share, 0.828, the change in the median corresponds to a decrease in AAA share of 4.6 percentage points. These results are qualitatively and quantitatively similar across robustness checks.

The decrease in distortions that we find is driven by the fall in ratings after the crisis, with no equivalent reduction in observable variables explaining ratings. Two additional factors at work are that we observe an increase in the dispersion of private information after the crisis, and there is a large change in incumbency measures after the crisis. The first may reflect changes in how ratings agencies collect information due to increased investor scrutiny. Because greater dispersion of private information means agencies are more differentiated from each other, distortions fall. The second is due to switching by issuers among agencies, in part due to an accounting scandal at S&P in 2011 (Baghai & Becker, 2018), but perhaps also a response to the crisis. However, neither the increase in dispersion of private information nor the changes in incumbency are nearly large enough to explain the overall decrease in distortion. Rather, it appears ratings decreased substantially relative to fundamentals after the crisis.[20] This result is consistent with increased investor scrutiny after the crisis.

A back-of-the-envelope calculation yields a rough sense of the extent to which CMBS issuers benefit from strategic bidding behavior. Specifically, we can approximate the impact of the distortions on issuers' funding costs, holding the bond yield associated with each rating fixed at historical values (i.e., ignoring equilibrium effects taking into account the investor response) and varying the share of AAA. The mean of λ_i went from 0.91 in 2006 to 0.76 in 2009, a change of 0.15. That implies a share of AAA that is 4.14 percentage points lower, for the average deal. To be conservative, suppose the change in AAA share were displaced by a corresponding adjustment in the share of AA (as opposed to some even lower rating such as BBB). Based on an 11.67 basis point historical premium on AAA relative to AA, a rough lower bound on the effect of the decrease in AAA share on annual issuer costs is +0.483 bp, or +$77 thousand annually for the average deal principal amount of $1,591.5 million.[21]

[20] Our baseline model presumes the same data-generating process for z_j before and after the crisis, which contradicts the finding that the variance of z_j increases post-crisis. As a robustness check, we estimated an alternate specification in which z_j is drawn from distinct distributions before and after the crisis. However, doing so does not change the finding that distortions decreased, with the median distortion during the post-crisis period corresponding to the 5th percentile of the pre-crisis distribution.

[21] The 11.67 bp premium is based on our calculation for the rating-implied spread on AA relative to AAA during the sample period, described in Appendix C. See footnote 32. We view this as a conservative estimate.

8 Robustness

The foregoing model specification is a stylized representation of the actual rating process. In this section, we consider the impact of perturbing some of the assumptions about the auction mechanism and agents' payoffs.

8.1 Alternative bidder preferences

First, we consider the possibility that winning bidders do not care about how much their own bid deviates from their belief, but rather about how much the pivotal bid deviates from their beliefs. This makes sense if investors can discover, with some probability, that the AAA share was distorted, and subsequently blame all of the winning agencies without distinguishing among individual agencies based on their bids. In order to capture this set-up, we modify the agencies' preferences such that the reputational loss function depends on the pivotal bid, rather than the agency's own bid. Specifically, we replace b_j with b^* in Equation 3:

$$\pi_j(b, \tilde{b}_j) \quad = \quad d_j(b) \left(V - c(b^*, \tilde{b}_j) \right), \tag{11}$$

In other words, bidders care about the equilibrium rating on the deal (b_i^*) and prefer for it to be close to their own type.

The difference between this alternative specification and the baseline one is analogous to the difference between uniform-price versus pay-as-bid multi-unit auctions: under uniform prices, all winning bidders pay the same price (as determined by the pivotal bid), whereas under pay-as-bid rules, the winning bidders' net utility depends on their own bids.

We report structural estimates in the left panel of Table 5. The scatterplot in Figure 3 compares the distortion λ_i implied by the "pay"-as-bid (baseline) versus uniform-"price" specifications (obviously, bidders do not literally pay a monetary price in our setting). Under the uniform-"price" specification, bidders have an incentive to choose slightly higher distortions than in the baseline case, because there is some probability the bidder can win without having to "pay" its own bid. However, the high correlation between the distortion measures implied by the two specifications ($\rho = 0.962$) is comforting because it suggests that our findings are robust to the choice of specification.

8.2 Ratings of both AAA and non-AAA securities

For reasons discussed in Section 3, the baseline model makes the simplifying assumption that a bid specifies the share of AAA, rather than a multidimensional set of conditions for all of the securities. To check the impact of this simplification, we consider an alternative assumption where the ratings proposals stipulate treatment of all of the securities, both AAA and non-AAA. This section provides a high-level overview of this approach, with details relegated to Appendix C.

For identifiability, we must define what a "bid" is in terms of something observable. In the baseline model, defining b_{ij} to be the specified share of AAA served this purpose; for this robustness check, we instead define b_{ij} to be the most favorable weighted-average rating for the deal that the rating agency would permit under an asserted loss distribution.[22] This specification is particularly sensible if the distribution of losses on the pool principal comes from some stochastically ordered family of distributions.[23] With or without stochastic ordering, the weighted-average rating represents a sensible way to capture the entire risk profile of a deal. Naturally, it is highly correlated with the percentage of the deal that is AAA, because AAA represents on average more than 80% of each deal.

Structural estimates are reported in the right panel of Table 5, with the correlation in unobserved terms appearing in Table 6. The scatterplot in Figure 4 compares the distortion λ_i implied by this alternative specification with the baseline specification. As in the previous figure, the distortions are correlated between specifications ($\rho = 0.372$), though to a lesser extent.

[22]The rating agencies have published statements purporting to describe the relative risk associated with different ratings—the so-called idealized loss rates." For example, according to Moody's, "AAA", "'A1" and "B1" have idealized loss rates over four years of 0.001 percent, 0.104 percent, and 7.6175 percent, respectively. See "Probability of Default Ratings and Loss Given Default Assessments for Non-Financial Speculative-Grade Corporate Obligors in the United States and Canada," Moody's Investors Service (August 2006), Appendix 1 (available online).

[23]Under stochastic ordering, given any two bids b and c where $c > b$, the higher bid c allows for a deal structure that "dominates" any deal structure allowed by the lower bid b, in the sense of having a higher share of the deal at or above each rating. If stochastic ordering holds, one dimensional bids would be appropriate.

8.3 Informal discussion of alternative auction formats

We model the industry as a simultaneous, multi-unit sealed bid auction because we believe that best describes the industry. But in this section, we informally evaluate several modifications to our assumptions. The goal is not to exhaustively consider every conceivable auction format, but to make intuitive arguments for why our approach is either robust or superior to alternatives.

First, we consider a rating process held as an open descending rather than a sealed-bid auction. Assuming a clock-auction format, the issuer would initially ask for a high AAA share and then gradually lower it until K bidders have accepted. With a single winner ($K = 1$), this mechanism is strategically equivalent to the baseline model. The case of $K = 2$ and ex ante symmetric bidders is equivalent to a case described by McCabe, Rassenti, and Smith (AER,1990), who find that, when there is a single unit left to be won, the remaining bidders behave the same as in our model, a sealed-bid "pay"-as-bid auction with K winners. Thus, the implied distortions are exactly the same as in our baseline model. Moreover, this equivalence holds regardless of whether bidders' preferences are as specified by the baseline model or by the uniform-"price" robustness check in Section 8.1: in both cases, the bidders that remain when all but one of the winners has been chosen know that they will be the pivotal bidder if they become the last winner.

When bidders are asymmetric, McCabe, Rassenti, and Smith's equivalence result does not hold exactly, because the identities of the first $K-1$ bidders are potentially informative to the remaining bidders. To our knowledge, there are no unambiguous theoretical results for this case.

Another possibility would be a multi-unit ascending auction, but this appears to be a bad fit for our industry. For example, consider an ascending auction with three players and two winners. After the first player dropped out, the remaining two would immediately drop out knowing that they were the winners. Thus, we would observe all three with the same bids, and the winning bid set by the lowest type. In contrast, in our industry, all three bidders make heterogenous bids, with the winning bid determined by the second highest bidder.[24]

[24] To be clear, we observe only the winning bid, so we cannot verify that bids are heterogeneous in our

An important assumption in our empirical approach is that the auction is characterized by affiliated private values rather than common values. One source of common value might be the ex post performance of the pool, but this is a noisy signal that is revealed only with substantial delay, and so is unlikely to influence bidding in the time frame we consider. It is difficult to distinguish between these models, but in support, Becker & Milbourn (2011) observe that, with Fitch's entry into rating corporate bonds, the incumbents significantly *raised* their ratings. Sangiorgi et al. (2009) interpret this phenomenon as being inconsistent with common values.[25]

Another possibility would be to consider a mechanism that includes an element of negotiation. For instance, suppose that after the issuer solicits an initial round of bids, the issuer can then confront each agency with claims about that agency's rivals' offers in order to elicit a more favorable bid. The issuer plays the agencies off one another through multiple rounds until the issuer either accepts one of the offers or breaks off negotiations. Thomas and Wilson (2002) consider such a model, and argue that after the issuer has revealed a bid, the rival agencies respond by either matching the rival bid or dropping out. The situation is essentially an ascending auction, which we just ruled out as being a poor match for this industry. Further, Thomas and Wilson (2002) make the point that the outcome depends on whether the issuer can credibly reveal bids. Without credible revelation, the mechanism reduces to our baseline mechanism: knowing that an issuer's claim about a rival's offer is cheap talk, an agency would not respond to claims of higher offers, and would bid the same way initially as if there were no follow-on rounds of negotiation.

9 Ex post security performance

As a form of model validation, we can assess testable implications using independent data on the ex post performance of the deals. In particular, we test whether a greater distortion

data. However, it is commonly known that the share approved for AAA in a security varies across agencies. Indeed, we do observe the shares of non-AAA in the security and those do vary across agency bids.

[25] Based on their assumption that the agencies are nonstrategic and bid their true expected value conditional on winning, one would expect Fitch's entry into rating corporate bonds to cause the incumbents to *lower* their ratings if they were adjusting for the winner's curse. A weaker conclusion still holds if we allow for rating catering: namely, that the impact of Fitch's entry on rating catering incentives (which should tend to increase the incumbents' ratings) outweighed whatever lowering impact it had through bidders adjusting for the winner's curse.

λ_i predicts worse ex post performance for deal i after flexibly controlling for the realized AAA share for the deal. Conceptually, the distortion should be informative about the ex post performance of the deal, even after controlling for the AAA share, because the latter is not fully revelatory about the bidders' private information: the AAA share is also affected by ratings shopping and ratings catering.

Table 7 shows tobit regressions where the dependent variable is the ex post losses on the deal, defined as the sum of principal writedowns and interest payment shortfalls as a share of the original pool principal, as of September 2012.[26] The key explanatory variable is the distortion, λ_i. In the first panel, we flexibly control for the equilibrium AAA share via a cubic spline with 10 knots. In addition, we control for the number of winning bidders. We also run the regression separately for subsamples of deals depending on the identities of the auction winners, presented in the remaining panels. Due to small sample size, we report winner-specific results only for auctions with two winners. We find that the distortions are a significant predictor of ex post losses. The pooled-sample estimates (in the first panel) imply that for an increase in the distortion of one standard-deviation (over deals i)—or 0.1099—there is an associated 34 basis point increase in the realized loss. In dollar terms, this difference amounts to on average \$5.44 million in losses per deal (multiplying the percentage loss increase by the average deal size). For comparison, the sample average for the ex post loss on the pool is 243 basis points. The remaining columns of the table show that the effect is driven primarily by deals rated by S&P and Moody's and deals rated by Moody's and Fitch. All standard errors in the table are bootstrapped and take into account the effect of estimation error on the distortion measure.

10 Counterfactuals

In this section, we use the estimated structural model to examine counterfactual outcomes under potential policy reforms. Regulatory rulemakings following the financial crisis have emphasized several key areas. We focus on enhanced transparency rules.[27] Such rules

[26]The corresponding exercise using estimates for the robustness checks are reported in Appendix Tables D.5 and D.6.

[27]Other areas of reform meant to reduce ratings catering, but that do not fit naturally into our model, are reducing regulatory reliance on ratings and stricter governance (e.g., prohibiting rating agencies' sales

include requirements for rating agencies and security issuers to disclose due diligence reports and other rating inputs—either to the public or to competing rating agencies (see Securities and Exchange Commission (SEC), 2014).

The regulatory debate considered two key aspects of enhanced transparency. First, it has the potential to reinforce market discipline by reducing asymmetric information among issuers, rating agencies, and investors. Second, transparency rules could have competitive implications through effects on rating agencies' incentives to enter or exit the market. We explore both of these in our counterfactual experiments.

10.1 Market Concentration

Focusing on the second issue first, the effect on entry and exit is theoretically ambiguous. On the one hand, an additional regulatory burden due to reporting requirements could create barriers to entry. On the other hand, making rating inputs available to all could also reduce the barriers to entry for would-be new rating agencies (see discussion in Securities and Exchange Commission (SEC), 2014). That is, the cost of forming z_{ij}, which we do not model, could decrease.

Without attempting to resolve the theoretical ambiguity of the effect on market entry and exit, our first counterfactual considers how the magnitude of rating distortions would be affected supposing there were an exogenous reduction in the number of competitors. We implement this change by considering a single security in a representative (in a sense that we explain more precisely below) auction with our three bidders and one winning bidder. For our baseline scenario, we compute the distortion in the equilibrium of this game. For the counterfactual scenario, we recompute the equilibrium when dropping one bidder. We do this for dropping each of the firms as a bidder and report the average across dropping each of firms.

We are interested not only in the change in the distortion when going from three to two bidders but also in decomposing that change into a ratings catering and ratings shopping effect. In order to separate these effects, we also compute a counterfactual scenario without allowing bids to adjust from the baseline scenario. That is, after dropping a bidder from

teams from being involved in rating securities).

the market, we calculate the maximum of the remaining two bids without allowing them to adjust their bids. Compared to the baseline, this intermediate scenario has no change in ratings catering, and thus isolates the ratings shopping effect. The ratings catering effect is then implied by the difference in distortion relative to the duopoly counterfactual with equilibirum bids. Reduced competition could have additional effects not taken into account by this model, so this is admittedly a partial analysis.[28]

For simplicity, our counterfactual analysis considers only pure-strategy Bayesian Nash equilibria (PSNE). Using arguments from Athey (2001), Appendix E shows that a PSNE exists for the bidding game based on the fact that the players' objective functions satisfy a single-crossing condition and on the private values assumption. We solve for a PSNE for each auction by approximating each bidder's bid function using a high-order polynomial approximation. The method is also described in Appendix E, and is based on Hubbard & Paarsch (2009) and Bajari (2001).

To state more precisely how we implement the first counterfactual, we simultaneously solve each bidder j's equilibrium bid function conditional on what the bidder observes, which can be denoted $b_j(x_i, w_i, z_{ij}, u_{ij})$. This function accounts for the fact that bidder j will use z_{ij} and u_{ij} to infer values of $z_{ij'}$ and $u_{ij'}$ of each competitor j', because private information is correlated across rivals in the same auction. We assume that bidder j's beliefs about its competitors' private information, conditional on j's own private information, are consistent with the empirical joint distribution of these random variables. For each possible value of bidder j's private information, we specify that bidder j's bid function is optimal with respect to the conditional distribution of bidder j's. Computation requires integrating out over u_{ij}, and for each bidder j and draw u_{ij}, integrating over rival's values of u_{ij} and z_{ij}, which we do by simulation. A more detailed description of our solution method appears in Appendix E.[29] In part because solving for the equilibrium of each auction is time-consuming, we do so for only one auction. We endow this auction with average values of x_i, w_{ij}, and z_{ij}. The value for x_i is the average of observed x_i across all auctions with a

[28]For example, competition may affect agencies' incentive to invest in more accurate models or other dynamic considerations.

[29]As described in the appendix, we simplify the computation by assuming there is only one unobservable, which we approximate as a linear function of the z_{ij} and u_{ij}. Note that the two elements of z_{ij} are highly correlated, and u_{ij} are correlated across within j, so the assumption appears reasonable.

single winner. For the average values of w_{ij} and z_{ij}, we take the average across all auctions for which j is the only winner. Because these values systematically differ across j, it does not make sense to average them across bidders.

We calculate $E\lambda$, the expected distortion in a representative auction under a given scenario. Unlike in the estimation, here the identity of the pivotal bidder is simulated and thus known. Table 8 summarizes the distortion under the counterfactual and baseline scenarios. For the baseline scenario, Table 8 shows that $E\lambda = 1.026$. Removing one competitor reduces the average distortion to $E\lambda = 0.762$. These numbers are in normalized bid space, and so do not have a natural interpretation. We compute standard errors via the bootstrap, and as the table shows, this change of 0.264 is large relative to the standard error of 0.079.

If we dropped a bidder but did not allow the remaining bidders to adjust bids, equilibrium distortion would be 0.981 (the second panel of Table 8). Thus, the ratings shopping effect accounts for the change from 1.026 to 0.981. Note that although dropping a competitor without adjusting bids must lead the issuer to take a weakly lower bid, that does not necessarily reduce the distortion, because in any given auction, a high bidder may have a low distortion. Indeed, we find that the reduction from 1.026 to 0.981 is not statistically significant, and some of our bootstrap samples show an increase in distortion. This result contrasts with the full counterfactual, in which distortions are less than in the baseline in all of our bootstrap samples.[30]

Thus, we find that dropping one bidder holding bids constant reduces the average distortion from 1.026 to 0.981, and allowing bids to adjust further reduces the expected distortion to 0.762. Therefore, about 17% of the change in distortion from eliminating a competitor is due to a reduction in ratings shopping, and the remaining 83% is due to a reduction in ratings catering.

In our model, reducing the number of competitors to one leads the remaining agency to choose $b_{ij} = \tilde{b}_{ij}$, or equivalently, $\lambda = 0$. Thus, the change in distortion in going from three to two (1.026 to 0.762) is small relative to the change in going from two to one (0.762 to

[30]The tables shows that the standard error for the change from the baseline to the full counterfactual for the 75th percentile is large relative to the change, but this results from a skewed distribution. The confidence interval from our bootstrap samples shows this change is negative.

0). We ascribe this in part to the high correlation across bidders in unobserved terms (u_{ij} and z_{ij}) that we find, which means that competitors are relatively homogenous. Analogous to a Bertrand pricing game with homogenous products, adding the first competitor has a larger effect than adding further competitors.

A further effect of reducing the number of agencies would be to increase the average level of incumbency in future auctions. Our calculations hold incumbency fixed at observed levels for ease of comparison, but this would be inaccurate as remaining firms divided up issuers and eventually appeared as incumbents in future auctions. As we find that higher incumbency leads to increased margins, we can expect market concentration to lead to increased distortion through this channel. While we do not calculate that effect here, it is worth considering in evaluating policy changes.

10.2 Disclosure of private information

Our second counterfactual considers the effect of forcing ratings agencies to share research with each other. We operationalize this idea by reclassifying some of the privately observed variables (z_{ij}, but not u_{ij}) to be common knowledge. Consistent with our private values paradigm, we assume that knowledge of rivals' private information does not affect how an agency values a security (that is, $\tilde{b}_{ij}$ is unaffected in the counterfactual). However, it does affect an agency's bid, because the agency knows more about the values of its rivals.

The effects of this counterfactual scenario are theoretically ambiguous. If the bidders gain knowledge of each other's private information, the direct effect on bidder j's behavior if j has a high draw relative to its competitors is to distort its bid less, whereas if it has a relatively low draw, the effect is to increase its distortion. There would also be additional, indirect effects due to the competitive response to this direct effect.

To implement the second counterfactual, we reclassify the bidders' common information for each auction i to include z_i in addition to x_i and w_i. Recall that z_i denotes the collection of z_{ij} for all bidders that produce a pre-sale report.[31] The remaining information, comprising the idiosyncratic component u_{ij}, stays private to each bidder. Just as in the first

[31] We abstract from the possibility that requiring disclosure of the pre-sale report variables before bidding may introduce incentives for the rating agencies to choose their pre-sale report variables strategically.

counterfactual, we assume that bidders' beliefs about their competitors' types are consistent with the empirical joint distribution of the bidders' private information.

For this counterfactual, we consider the case in which there are three bidders and two winners. As before, we limit computational time by calculating the counterfactual for a single representative auction. For the representative auction, we set x_{ij} to the average of the observed x_{ij} for auctions with *two* winners where both winners produce pre-sale reports. The variables w_{ij} and z_{ij} are set to the average from these auctions for each bidder j, conditional on j being a winner in i.

For each auction, we compute the equilibrium bid function of each bidder, now conditioning on z_i rather than z_{ij}, which can be denoted by $b(x_i, w_i, z_i, u_{ij})$. As above, we solve for equilibrium in bids for each simulation draw (which requires integrating over simulated draws for rivals), and average over the pivotal distortion for each simulation. The baseline is different from the previous counterfactual because we use a different representative auction.

As shown in the lower panel of Table 8, we find that making z_i public increases the expected distortion from 0.946 to 0.983. The change in expected distortion suggests that, under the particular conditions of the market studied, greater transparency leads to higher distortions overall. However, it is important to note that the overall direction of change masks considerable heterogeneity across auctions. For individual auctions (not shown), the distribution of the distortion for the pivotal bidder may either shift up or down. Indeed, the change in expected distortion does not appear large, and it is not statistically significant.

Overall, Counterfactual 2 finds that if anything, disclosure increases distortion, rather than decreasing it, which is presumably the opposite of the policy goal. An interpretation of this result is that by giving all bidders the same information, it makes them more homogenous which leads to more intense bidding competition. However, if the effect of disclosure is also to cause exit, Counterfactual 1 shows that the effect can be to reduce competition, and by a more significant amount than disclosure. Note that in our model, because competition causes ratings catering and ratings shopping, competition is generally bad for the market. Naturally, we are focusing on these issues because they are central in the policy debate, but we are ignoring price competition by ratings agencies as well as any investment in quality, on which competition presumably also has an effect.

11 Conclusion

Biases in ratings are potentially important in markets in which product reviewers behave strategically. If the reviewers must compete to perform the review, ratings may be distorted toward being too favorable. Quantifying the magnitude of the distortion is empirically challenging because the reviewers' true assessments of quality are unobserved, precluding a direct approach to quantification.

We instead use an indirect approach that treats competition among rating agencies to rate a deal as an auction. Our approach permits comparisons of the amount of rating distortion in different auctions. We relate the implied degree of distortion associated with each deal to the ex post performance of the deal, and thereby confirm the economic significance of the distortion. CMBS is an important market in which to study these issues, both because of the size of the industry and because policymakers are debating various proposals for reforming the rating agency industry, which we address using counterfactual simulations.

We find that the distortions are economically significant. A one standard-deviation increase in the amount of distortion is associated with a 34 basis-point increase in realized losses on a deal—almost one fourth of the sample mean of the ex post losses. Our estimates also indicate that rating distortions decreased after the recent financial crisis. Empirically, although we find some changes in observable characteristics of the market after the crisis, such as an increase in the dispersion of private information, these changes cannot fully explain the reduction in ratings post-crisis. Thus, our model captures this change with a reduction in the distortion, suggesting that there was an increase in investor scrutiny of ratings after the crisis.

Two caveats point to future directions for research. First, we do not explicitly capture the dynamics of rating agencies' reputational incentives. In particular, one might conjecture that if an agency rates a deal favorably but that the deal suffers large losses ex post, the agency's credibility would be diminished, making it difficult for the agency to charge issuers as high a fee to rate future deals. Quantifying this effect would require additional data, both to identify investor demand for ratings and to identify the time series properties of mortgage losses. Furthermore, some investors may be sophisticated enough to anticipate ratings distortion and respond. Second, because we do not model investor behavior, our findings are agnostic on total-welfare implications. As mentioned in the introduction, some observers have ascribed a causal role to ratings distortion in the creation of the financial crisis of 2008, which suggests welfare implications are quite important.

References

Athey, S. (2001). Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, *69*(4), 861–889.

Athey, S. & Haile, P. A. (2002). Identification of standard auction models. *Econometrica*, *70*(6), 2107–2140.

Athey, S. & Haile, P. A. (2007). Nonparametric approaches to auctions. *Handbook of econometrics*, *6*, 3847–3965.

Baghai, R. & Becker, B. (2018). Reputations and credit ratings: Evidence from commercial mortgage-backed securities. CEPR Discussion Paper No. DP12648.

Bajari, P. (2001). Comparing competition and collusion: A numerical approach. *Economic Theory*, *18*, 187–205.

Bajari, P., Benkard, L., & Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, *75*, 1331–1370.

Bajari, P., Hong, H., Krainer, J., & Nekipelov, D. (2010). Estimating static models of strategic interactions. *Journal of Business and Economic Statistics*, *28*.

Basel Committee (2012). Consultative document: Revisions to the Basel securitisation framework.

Becker, B. & Milbourn, T. (2011). How did increased competition affect credit ratings? *Journal of Financial Economics*, *101*(3), 493–514.

Black, L. K., Chu, C. S., Cohen, A., & Nichols, J. B. (2012). Differences across originators in CMBS loan underwriting. *Journal of Financial Services Research*, *42*(1-2), 115–134.

Bolton, P., Freixas, X., & Shapiro, J. (2012). The credit ratings game. *The Journal of Finance*, *67*(1), 85–111.

Camanho, N., Deb, P., & Liu, Z. (2012). Credit rating and competition. In *22nd Australasian Finance and Banking Conference*.

Cohen, A. & Manuszak, M. D. (2013). Ratings competition in the CMBS market. *Journal of Money, Credit and Banking*, *45*(s1), 93–119.

Coval, J., Jurek, J., & Stafford, E. (2009). The economics of structured finance. *The Journal of Economic Perspectives*, *23*(1), 3–25.

Dewatripont, M., Rochet, J.-C., & Tirole, J. (2010). *Balancing the Banks: Global Lessons from the Financial Crisis*. Princeton University Press.

Efing, M. & Hau, H. (2015). Structured debt ratings: Evidence on conflicts of interest. *Journal of Financial Economics*, *116*(1), 46–60.

Faure-Grimaud, A., Peyrache, E., & Quesada, L. (2009). The ownership of ratings. *The RAND Journal of economics*, *40*(2), 234–257.

Financial Crisis Inquiry Commission (2011). The financial crisis inquiry report. U.S. Government Printing Office.

Fitch (2008). Rating criteria for Fitch's U.S. CMBS multiborrower rating model.

Flynn, S. & Ghent, A. (2014). When low standards are a winning strategy: How credit rating agencies compete. Unpublished manuscript, Arizona State University.

Furfine, C. (2014). Complexity and loan performance: Evidence from the securitization of commercial mortgages. *Review of Corporate Financial Studies*, *2*, 154–187.

Griffin, J. M., Nickerson, J., & Tang, D. Y. (2013). Rating shopping or catering? An examination of the response to competitive pressure for CDO credit ratings. *The Review of Financial Studies*, *26*(9), 2270–2310.

Griffin, J. M. & Tang, D. Y. (2012). Did subjectivity play a role in CDO credit ratings? *The Journal of Finance*, *67*(4), 1293–1328.

Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, *68*, 525–574.

He, J., Qian, J., & Strahan, P. E. (2011). Credit ratings and the evolution of the mortgage-backed securities market. *The American Economic Review*, *101*(3), 131–135.

Hortacsu, A. & Kastl, J. (2012). Valuing dealer's informational advantage: A study of Canadian treasury auctions. *Econometrica*, *80*.

Hubbard, T. P. & Paarsch, H. J. (2009). Investigating bid preferences at low-price, sealed-bid auctions with endogenous participation. *International Journal of Industrial Organization*, *27*, 1–14.

Judd, K. L. & Su, C.-L. (2012). Constrained optimization approaches to estimation of structural models. *Econometrica*, *80*, 2213–2230.

Kastl, J. (2016). Recent advances in empirical analysis of financial markets: Industrial organization meets finance. Invited Lecture at the 11th World Congress of the Econometric Society, Available at Princeton University.

Krasnokutskaya, E. (2011). Identification and estimation of auction models with unobserved heterogeneity. *The Review of Economic Studies*, *78*(1), 293–327.

Krugman, P. (2010). Berating the raters. *New York Times*, *April 26*, A23.

Lehmann, E. L. & Romano, J. P. (2005). *Testing Statistical Hypotheses* (3 ed.). Springer-Verlag New York.

Lowenstein, R. (2011). *The End of Wall Street*. Penguin Random House.

Luca, M. & Zervas, G. (2016). Fake it till you make it: Reputation, competition and Yelp review fraud. *Management Science*, *62*, 3412–3427.

Mathis, J., McAndrews, J., & Rochet, J.-C. (2009). Rating the raters: Are reputation concerns powerful enough to discipline rating agencies? *Journal of monetary economics*, *56*(5), 657–674.

Moody's (2000). CMBS: Moody's approach to rating U.S. conduit transactions.

Newey, W. K. & Powell, J. L. (2003). Instrumental variable estimation of nonparametric models. *Econometrica*, *71*.

Popper, N. (Aug 1, 2013). Banks find S&P more favorable in bond ratings. *New York Times*, A1.

Prager, R. (2016). Industrial Organization research by Federal Reserve Board economists. *Review of Industrial Organization*, *48*, 149–180.

Rivlin, A. M. & Soroushian, J. B. (2017). Credit rating agency reform is incomplete. Brookings Institution.

Sangiorgi, F., Sokobin, J., & Spatt, C. (2009). Credit-rating shopping, selection and the equilibrium structure of ratings. Working Paper, Stockholm School of Economics and Carnegie Mellon University.

Sangiorgi, F. & Spatt, C. (2017). Opacity, credit rating shopping, and bias. *Management Science*, *63*, 4016 – 4036.

Securities and Exchange Commission (2008). Summary report of issues identified in the commission staff's examinations of select credit rating agencies.

Securities and Exchange Commission (SEC) (2014). Release no. 34-72936 (final rule release).

Skreta, V. & Veldkamp, L. (2009). Ratings shopping and asset complexity: A theory of ratings inflation. *Journal of Monetary Economics*, *56*(5), 678–695.

S&P (2009). U.S. CMBS rating methodology and assumptions for conduit/fusion pools.

Stiglitz, J. E. (2010). The financial crisis of 2007-2008 and its macroeconomic consequences. In S. Griffith-Jones, J. A. Ocampo, & J. Stiglitz (Eds.), *Time for a Visible Hand: Lessons from the 2008 World Financial Crisis*. Oxford University Press.

Sun, Q., Li, S., & Tan, W. (2013). An empirical analysis of credit ratings game in subprime mortgage-backed security market. Manuscript, Shanghai University of Finance and Economics.

Taibbi, M. (July 18, 2013). The last mystery of the financial crisis. *Rolling Stone*.

U.S. Senate Committee on Banking, Housing, and Urban Affairs (2008). Turmoil in the U.S. credit markets: The role of the credit rating agencies (April. 22, 2008), testimony by John Coffee. Columbia University Law School.

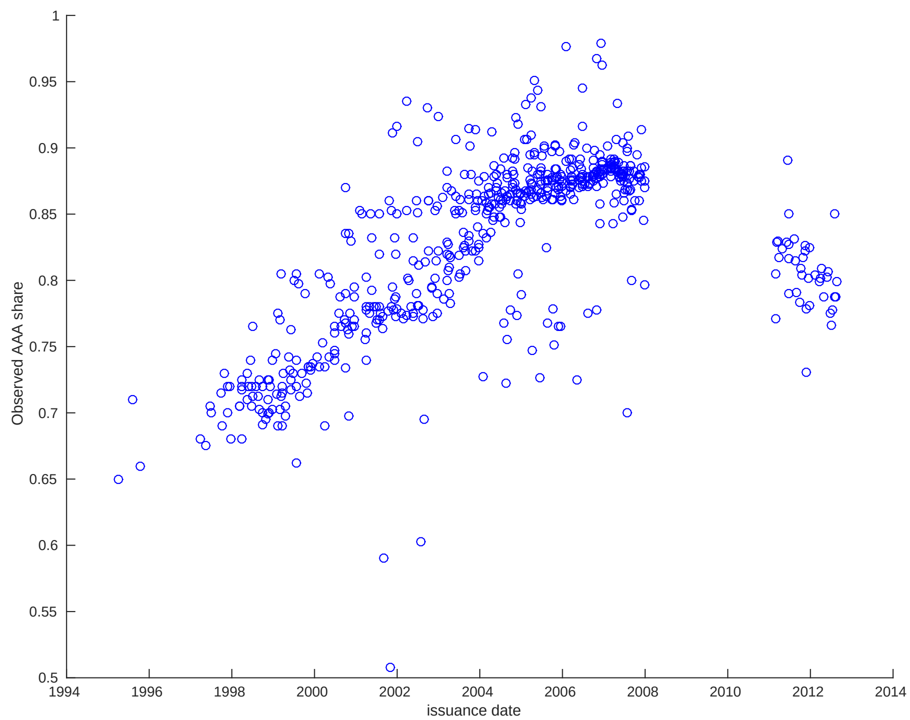

Figure 1: Share of security rated AAA, by time.

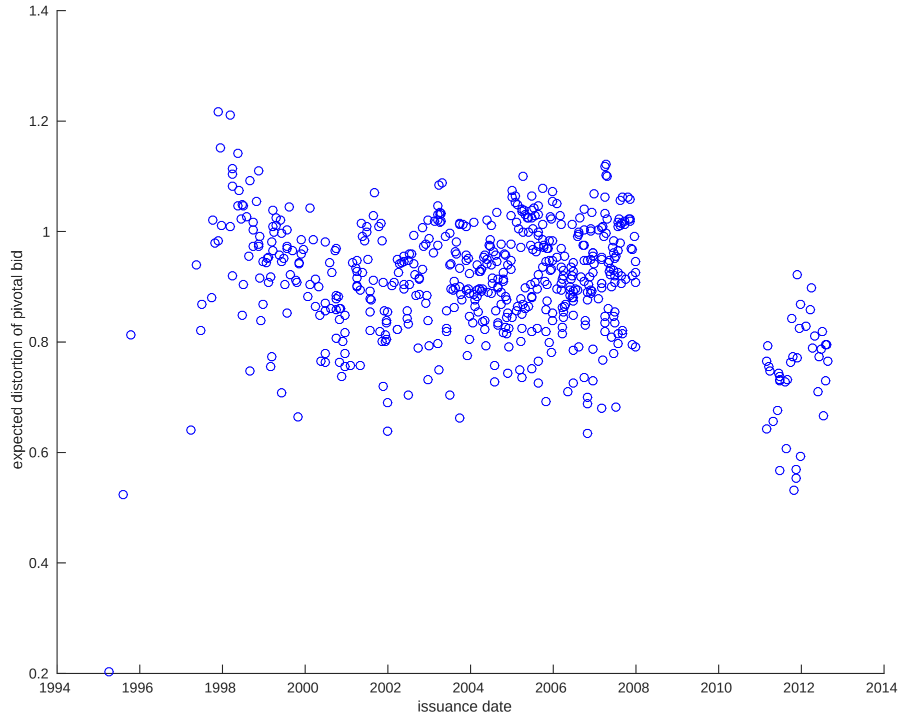

Figure 2: Distortion (λ_i) by date of issuance for deals.

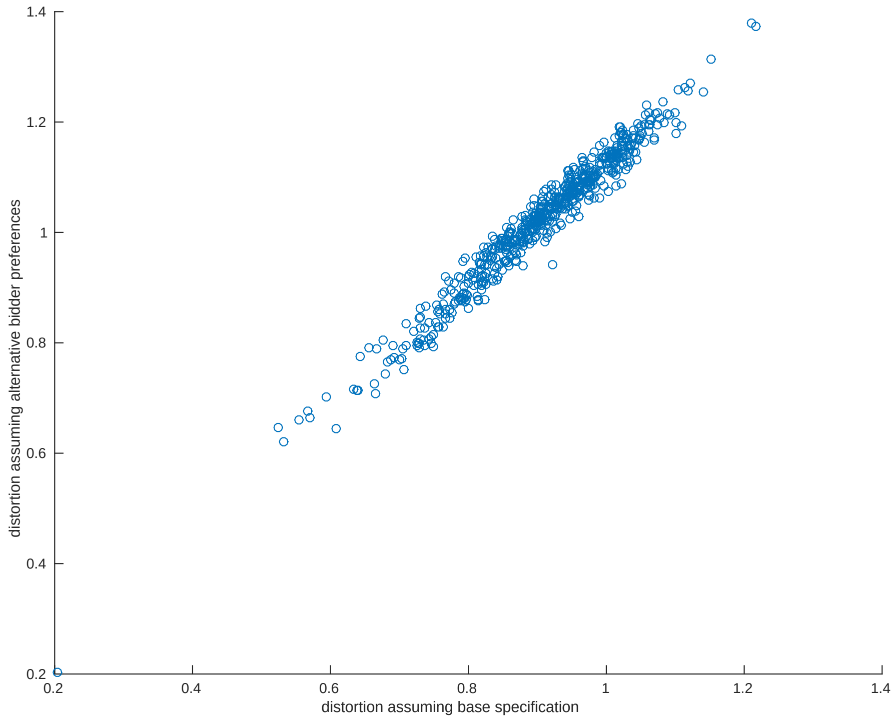

Figure 3: Distortion (λ_i) under baseline specification versus robustness check with alternative bidder preferences.

Figure plots the distortion λ_i for CMBS deals for the baseline specification (x-axis), in which the bidder's loss function depends on the discrepancy between the bidder's type and the bidder's *own* bid, versus the robustness check (y-axis), in which the bidder's loss function depends on the discrepancy between the bidder's type and the *pivotal bidder's* bid (i.e., the equilibrium AAA share). See Subsection 8.1 for details.

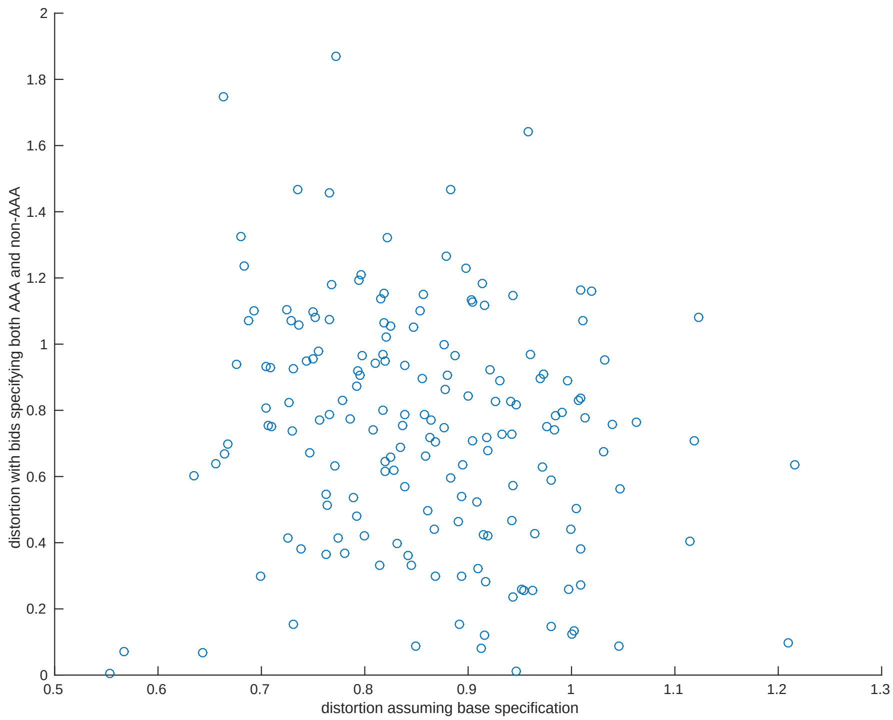

Figure 4: Distortion (λ_i) under baseline specification versus robustness check in which bids specify structure of both AAA and non-AAA securities.

Figure plots the distortion for CMBS deals for the baseline specification versus the robustness check in which the rating agencies specifies the structure of both AAA and non-AAA securities. See Subsection 8.2 for details.

Table 1: Summary Statistics for CMBS Deals

Rating outcomes	**Mean**	**Std Dev**
Share of AAA security	0.828	0.067
S&P rated deal	0.701	
Moody's rated deal	0.705	
Fitch rated deal	0.581	
Deal covariates	**Mean**	**Std Dev**
Weighted-average loan characteristics at origination		
Balloon payment	0.951	0.138
Cross-collateralization*	0.162	0.204
Original DSCR	1.507	0.277
Original LTV	0.677	0.067
Loan balance ($M)	1591.5	1131.6
Loan originated by deal issuer	0.467	0.349
Concentration indices (HHI) for loan characteristics		
By originator	0.550	0.249
By property type	0.324	0.156
By region	0.239	0.238
By MSA	0.146	0.264
By loan's share of pool balance	0.031	0.029
Pool total principal ($B)	1.591	1.132
Ex post performance of pool	**Mean**	**Std Dev**
Principal writedown and interest shortfall, as proportion of original principal, as of Sep. 2012	0.0243	0.0265

S&P bidder-deal-specific covariates	**Mean**	**Std Dev**
S&P produced pre-sale report, conditional on rating deal	0.928	
reunderwritten DSCR (weighted average for deal)	1.457	0.207
reunderwritten LTV (weighted average for deal)	0.835	0.158
S&P's market share: among last 10 deals (by issuance date)	0.708	0.228
" " : among last 3 deals issued by same bank as current issuer	0.698	0.315
" " : among last 10 deals issued by a different bank	0.704	0.233
Moody's bidder-deal-specific covariates	**Mean**	**Std Dev**
Moody's produced pre-sale report, conditional on rating deal	0.722	
reunderwritten DSCR (weighted average for deal)	1.350	0.288
reunderwritten LTV (weighted average for deal)	0.792	0.160
Moody's market share: among last 10 deals (by issuance date)	0.700	0.172
" " : among last 3 deals issued by same bank as current issuer	0.705	0.273
" " : among last 10 deals issued by a different bank	0.699	0.176
Fitch bidder-deal-specific covariates	**Mean**	**Std Dev**
Fitch produced pre-sale report, conditional on rating deal	0.955	
reunderwritten DSCR (weighted average for deal)	1.362	0.255
reunderwritten LTV (weighted average for deal)	0.778	0.142
Fitch's market share: among last 10 deals (by issuance date)	0.581	0.200
" " : among last 3 deals issued by same bank as current issuer	0.566	0.301
" " : among last 10 deals issued by a different bank	0.572	0.197

N = 613

* Most loans are collateralized by a single property. However, some mortgages are collateralized by multiple properties, which market participants consider to be a form of risk-diversification. We define the "cross-collateralization" of a loan as the logarithm of the number of properties that serve as collateral for the loan. Bidder-specific summary statistics are for the 591 observations for which we have the full set of data.

Table 2: Hazard regressions for Loan Default Time

	Estimate	Std error
Loan characteristics at deal cutoff date		
Loan seasoning	0.014	0.001
Original DSCR	-0.295	0.059
Occupancy	-2.036	0.116
No occupancy data	-2.015	0.112
Original LTV	3.499	0.162
Contractual interest rate*	0.416	0.021
Original loan amount	0.820	0.056
Interest-only loan	0.192	0.037
Fixed-rate mortgage	-0.727	0.178
Insurance co. loan	-0.314	0.050
I-bank loan	0.125	0.027
Domestic conduit loan	0.306	0.055
Finance co. loan	-0.004	0.047
Foreign conduit loan	0.220	0.027
Characteristics of deal containing the loan		
AAA share (linear effect)	-0.265	0.332
No pre-sale reports available	-0.425	0.100
Bidder-deal-specific variables, averaged over winning bidders		
(orig. DSCR - presale DSCR)/(orig. DSCR)**	0.300	0.151
(presale LTV - orig. LTV)/(orig. LTV)**	0.342	0.120
Fixed effects		
Loan origination year	Included	
Loan origination month	Included	
Region and property-type (interacted)	Included	
N	59433	

Table shows a Cox proportional hazard regression for individual loans' default times. Dependent variable is the time to default, defined as the time between loan origination and the point at which the loan is 60 days delinquent or in special servicing.

* We express the contractual interest rate as a spread relative to the yield on U.S. Treasuries with the same maturity as the loan, as of the month of loan origination.

** Letting $DSCR_{il}^{orig}$ and LTV_{il}^{orig} denote the original DSCR and LTV for loan l in deal i, and letting $DSCR_{ijl}$ and LTV_{ijl} denote the corresponding reunderwritten values by agency j, the first two bidder-deal-specific covariates above are defined as follows:

$$\frac{1}{\#(winning\ bidders\ j\ for\ deal\ i)} \sum_{winning\ bidders\ j\ for\ deal\ i} (DSCR_{il}^{orig} - DSCR_{ijl})/DSCR_{il}^{orig}$$

$$\frac{1}{\#(winning\ bidders\ j\ for\ deal\ i)} \sum_{winning\ bidders\ j\ for\ deal\ i} (LTV_{ijl} - LTV_{il}^{orig})/LTV_{il}^{orig}$$

When one or more of the winning bidders does not produce a pre-sale report for deal i, we instead average over the remaining set of winning bidders. For the rare event in which none of the winning bidders produces a pre-sale report for deal i (indicated by the dummy "No pre-sale reports available"), we set these variables to zero.

Table 3: Structural Estimates

	Base Specification		Alternative Spec.	
Bidder-deal-specific covariates, commonly observed (β_1)				
	Estimate	Std. Err.	Estimate	Std. Err.
Own share of last 10 deals	0.345	0.304		
Avg. competitors' share of last 10 deals	0.834	0.502		
Own share of last 3 deals by same bank	0.095	0.135		
Avg. competitors' share of last 3 deals by same bank	-0.141	0.154		
Own share of last 10 deals by a different bank			0.263	0.214
Avg. competitors' share of last 10 deals by a different bank			0.434	0.319
Deal covariates (β_2)				
Constant	0.330	0.820	0.496	0.552
Balloon payment (wtd avg)	-0.481	0.374	-0.356	0.384
Cross-collateralization (wtd avg)	-0.069	0.060	-0.048	0.051
Deal size (total principal)	0.023	0.010	0.022	0.010
Originator HHI	-0.105	0.044	-0.076	0.045
Property type HHI	-0.408	0.300	-0.318	0.276
Region HHI	0.069	0.162	0.084	0.160
2001 vintage	0.250	0.049	0.223	0.037
2002 vintage	0.355	0.073	0.297	0.062
2003 vintage	0.430	0.052	0.395	0.050
2004 vintage	0.532	0.050	0.489	0.043
2005 vintage	0.569	0.054	0.538	0.046
2006 vintage	0.648	0.053	0.590	0.048
2007 vintage	0.597	0.068	0.554	0.057
2010-2012 vintages	0.528	0.144	0.397	0.083
Bidder-deal-specific covariates, private information (β_3)				
Reunderwritten DSCR	-0.057	0.146	0.026	0.134
Reunderwritten LTV	-0.419	0.213	-0.287	0.184
Bidder produced no pre-sale report	-0.073	0.179	-0.072	0.181
Number of Observations:	591		591	

Table shows estimates of the structural parameters that determine the relationship between observed covariates and bidders' types. The distribution of the residual u_{ij} is computed by simulating the distribution of types for the full set of bidders and netting out the effects of the covariates. First-step estimates for the "Alternative Specification" are in Appendix Tables D.1 and D.2.

Table 4: Covariance of residual u_{ij} in structural estimation

	Base Specification			**Alternative Spec.**		
	S&P	Moody's	Fitch	S&P	Moody's	Fitch
S&P	0.111			0.133		
Moody's	0.090	0.137		0.102	0.151	
Fitch	0.106	0.094	0.112	0.108	0.097	0.139

The distribution of the residual u_{ij} is computed by simulating the distribution of types for the full set of bidders and netting out the effects of the covariates. The residuals are drawn from the specifications in Table 3.

Table 5: Structural Estimates, robustness checks

	Uniform-"price" specification		**Bidding AAA and non-AAA**	
Bidder-deal-specific covariates, commonly observed (β_1)				
	Estimate	Std. Err.	Estimate	Std. Err.
Own share of last 10 deals	0.435	0.313	1.208	0.444
Avg. competitors' share of last 10 deals	0.859	0.452	-1.229	0.943
Own share of last 3 deals by same bank	0.105	0.144	-1.135	0.347
Avg. competitors' share of last 3 deals by same bank	-0.154	0.165	1.390	0.412
Deal covariates (β_2)				
Constant	0.142	0.746	1.456	1.183
Balloon payment (wtd avg)	-0.447	0.355	-0.415	0.427
Cross-collateralization (wtd avg)	-0.069	0.059	-0.053	0.102
Deal size (total principal)	0.030	0.011	-0.001	0.020
Originator HHI	-0.105	0.046	-0.195	0.078
Property type HHI	-0.342	0.278	-0.383	0.363
Region HHI	0.118	0.153	0.085	0.147
2001 vintage	0.248	0.047	0.192	0.080
2002 vintage	0.347	0.071	0.272	0.074
2003 vintage	0.430	0.052	0.332	0.074
2004 vintage	0.525	0.047	0.421	0.095
2005 vintage	0.556	0.052	0.514	0.095
2006 vintage	0.638	0.051	0.502	0.095
2007 vintage	0.574	0.067	0.633	0.126
2010-2012 vintages	0.528	0.131	0.339	0.230
Bidder-deal-specific covariates, private information (β_3)				
Reunderwritten DSCR	-0.049	0.145	0.188	0.109
Reunderwritten LTV	-0.397	0.218	-0.343	0.413
Bidder produced no pre-sale report	-0.078	0.163	-0.206	0.196
Number of Observations:	591		591	

Table shows structural estimates for specifications described in the robustness checks. The results on the left modify bidders' preferences such that they depend on the discrepancy between a bidder's type and the *pivotal bidder's* bid (i.e., equilibrium AAA share). The results on the right modify the interpretation of rating agencies' shadow rating such that a bid specifies the structure of non-AAA securities as well as the share of AAA. First-step estimates for the former specification are the same as for the baseline model. First-step estimates for the latter specification are in Appendix Tables D.3 and D.4.

Table 6: Covariance of residual u_{ij} in robustness checks

	S&P	Moody's	Fitch	S&P	Moody's	Fitch
S&P	0.099			0.238		
Moody's	0.087	0.120		0.052	0.298	
Fitch	0.096	0.090	0.102	0.032	-0.016	0.254

See the note in Table 4. The residuals are drawn from the specifucations in Table 5.

Table 7: Tobit regressions of ex post deal outcomes (principal losses and interest shortfall) on distortion (λ_i) and control variables

	Full sample of deals		**Deals rated by Moody's and S&P**	
	Estimate	Std. Error	Estimate	Std. Error
Distortion (λ_i)	0.031	0.013	0.045	0.022
Splines for AAA share	Included		Included	
Deal rated by 3 agencies	-0.001	0.003		
Deal rated by 1 agency	-0.024	0.004		
Constant	-0.008	0.014	-0.009	0.021
Square root of error variance	0.026	0.003	0.020	0.001
Observations	579		203	

	Deals rated by S&P and Fitch		**Deals rated by Moody's and Fitch**	
Distortion (λ_i)	-0.005	0.018	0.072	0.028
Splines for AAA share	Included		Included	
Constant	0.022	0.033	-0.022	0.041
Square root of error variance	0.031	0.006	0.026	0.004
Observations	147		131	

Dependent variable is sum of principal loss and interest payment shortfalls on the deal's loan pool, as of the censoring date (September 2012), expressed as a share of the original pool principal. Letting ψ denote linear coefficients and ε a normal error, the assumed model is:

$$\begin{aligned} \text{dependent variable} &= \psi'(\text{covariates}) + \varepsilon \text{ if } \psi'(\text{covariates}) + \varepsilon > 0, \\ &= 0 \text{ otherwise.} \end{aligned}$$

Standard errors are bootstrapped and take into account estimation error in both the first step and in structural estimation.

Table 8: Summary of simulated bid distortions under counterfactuals

Counterfactual 1: Eliminating Fitch from bidder set for set of auctions 1 observed winner

Variable and scenario	*Statistic*		*Std. err.*	*Change from baseline*	*Std. err.*
Distortion for pivotal bidder (λ), baseline scenario	Expectation	1.026	0.431		
	25th percentile	0.927	0.432		
	75th percentile	1.120	0.429		
Distortion after dropping one bid, chosen randomly, baseline scenario prices	Expectation	0.981	0.440	-0.045	0.034
	25th percentile	0.881	0.449	-0.046	0.055
	75th percentile	1.071	0.433	-0.049	0.027
Distortion after dropping one bid, chosen randomly allowing prices to adjust	Expectation	0.762	0.414	-0.264	0.079
	25th percentile	0.607	0.427	-0.320	0.083
	75th percentile	0.912	0.413	-0.208	0.134

Counterfactual 2: Pre-sale reports commonly observed for set of auctions with 2 observed winners

Variable and scenario	*Statistic*		*Std. err.*	*Change from baseline*	*Std. err.*
Distortion for pivotal bidder (λ), baseline scenario (reports not commonly observed)	Expectation	0.946	0.458		
	25th percentile	0.868	0.457		
	75th percentile	1.003	0.458		
Distortion for pivotal bidder, counterfactual scenario (reports commonly observed)	Expectation	0.983	0.449	0.037	0.049
	25th percentile	0.904	0.448	0.036	0.037
	75th percentile	1.056	0.450	0.053	0.062

The upper panel compares baseline in a representative auction to a counterfactual in which one bidder is dropped at random according to market share. The panel provides the outcome with and without allowing prices to adjust to the new equilibrium. The lower panel compares baseline to a counterfactual in which bidders commonly observe all pre-sale report information observed in our data (z_i^*).

Appendices

Appendix A Reduced-form relationship between equilibrium AAA share and covariates

Table A.1 shows the reduced-form relationship between the AAA share on deals and deal-level covariates.

The regression is imprecisely estimated because the dependent variable is determined by the pivotal bid, whereas the covariates that are bidder-deal-specific are averages over the set of winning bidders, which gives rise to measurement error. We take the average because we cannot perfectly observe the identity of the pivotal bidder. The variables relating to DSCR and LTV are explained in the table footnote. We consider three alternative definitions of incumbency (corresponding to columns (1), (2), and (3) of the table), and take the average of this measure over the winning bidders.

As expected, a higher reunderwritten LTV predicts a lower AAA share, but the reunderwritten DSCR is not significant at the 5 percent level. The incumbency measures have negative coefficients, suggesting that bidders submit higher bids (are more generous with their ratings) when they have a lower recent market share.

Analysis of serial correlation

When estimating the structural model, we shall be concerned about potential serial correlation over auctions i in the unexplained residual of bidder j's type (u_{ij} in Equation 2), which would make the incumbency measures endogenous. As an informal test to see whether serial correlation is a concern, we tested for autocorrelation in the AAA share, which would be an implication if bidders' types are autocorrelated. Specifically, we examined a regression similar to the one reported in Table A.1 (excluding the incumbency measures from the set of regressors), estimated separately for each subsample of auctions as defined by the identities of the set of winning bidder(s). For each subsample, we perform Durbin's alternative test for autocorrelation, sequencing the deals by their cutoff dates. Autocorrelation cannot be meaningfully assessed in the case of deals with one winner due to small sample size. However, for each of the remaining subsamples, we cannot reject the null hypothesis of no autocorrelation at the 5-percent significance level. This lack of autocorrelation is unsurprising, considering that we already control for year effects.

Table A.1: OLS regressions with AAA share (as proportion of deal principal) as dependent variable

	(1)		(2)		(3)	
Pool characteristics at cutoff—weighted averages over loans						
	Estimate	Std error	Estimate	Std error	Estimate	Std error
Balloon payment	-0.078	0.019	-0.078	0.019	-0.077	0.019
Cross-collateralization	-0.014	0.007	-0.015	0.007	-0.014	0.007
Original DSCR	-0.022	0.009	-0.023	0.009	-0.022	0.009
Original LTV	-0.368	0.046	-0.369	0.046	-0.368	0.046
Pool characteristics at cutoff—concentration indices (HHI)						
By originator	-0.008	0.006	-0.008	0.006	-0.008	0.006
By property type	-0.112	0.016	-0.112	0.016	-0.112	0.016
By region	-0.058	0.041	-0.060	0.041	-0.057	0.041
By MSA	0.057	0.038	0.059	0.038	0.056	0.038
By loan's share of pool balance	-0.530	0.108	-0.536	0.107	-0.531	0.107
Other deal characteristics						
Pool total principal	0.006	0.002	0.007	0.002	0.007	0.002
No pre-sale reports available	-0.007	0.009	-0.007	0.009	-0.007	0.009
Deal rated by 3 agencies	-0.021	0.005	-0.021	0.005	-0.021	0.005
Deal rated by 1 agency	0.010	0.008	0.009	0.008	0.010	0.008
Vintage dummies	Included		Included		Included	
Bidder-deal-specific variables, averaged over winning bidders						
Average over loans in deal of*:						
$\frac{\text{orig. DSCR - presale DSCR}}{\text{orig. DSCR}}$	0.008	0.019	0.009	0.019	0.008	0.019
$\frac{\text{presale LTV - orig. LTV}}{\text{orig. LTV}}$	-0.059	0.015	-0.059	0.015	-0.059	0.015
Winning bidders' market share:						
among last 10 deals	-0.018	0.014				
among last 3 deals issued by same bank			-0.011	0.008		
among last 10 deals issued by different bank					-.015	0.013
N	578		579		579	
R-squared	0.7685		0.7717		0.7714	

* Let $DSCR_{il}^{orig}$ and LTV_{il}^{orig} denote the original DSCR and LTV for loan l in deal i, let $DSCR_{ijl}$ and LTV_{ijl} denote the corresponding reunderwritten values by agency j, and let w_l denote loan l's proportion of the total pool principal. The two covariates under this heading are defined as follows:

$$\frac{1}{\#(winning\ bidders\ j\ for\ deal\ i)} \sum_{winning\ bidders\ j\ for\ deal\ i} \sum_l w_l \cdot (DSCR_{il}^{orig} - DSCR_{ijl})/DSCR_{il}^{orig}$$

$$\frac{1}{\#(winning\ bidders\ j\ for\ deal\ i)} \sum_{winning\ bidders\ j\ for\ deal\ i} \sum_l w_l \cdot (LTV_{ijl} - LTV_{il}^{orig})/LTV_{il}^{orig}$$

Note that these variables are each defined such that a higher value indicates a riskier reunderwritten assessment, relative to the original assessment.

When one or more of the winning bidders does not produce a pre-sale report, we instead average over the remaining set of winning bidders. For the rare event in which none of the winning bidders produces a pre-sale report (indicated by the dummy "No pre-sale reports available"), we set these variables to zero.

Appendix B First-step estimates, within-sample fit, and out-of-sample fit

Table B.1: First-step estimates (base specification): distribution of presale-report variables ($G_z(\cdot)$)

Means (μ_z) relative to sample means

	S&P		Moody's		Fitch	
	DSCR	LTV	DSCR	LTV	DSCR	LTV
Estimate	1.422	0.928	1.095	0.884	1.269	0.863
Std Error	0.037	0.007	0.028	0.010	0.027	0.009

Covariances (Ω_z)

		S&P		Moody's		Fitch	
		DSCR	LTV	DSCR	LTV	DSCR	LTV
S&P	DSCR	0.127					
	LTV	-0.006	0.011				
Moody's	DSCR	0.039	-0.005	0.090			
	LTV	0.005	0.011	-0.008	0.018		
Fitch	DSCR	0.070	-0.006	0.044	-0.002	0.064	
	LTV	0.004	0.011	-0.009	0.016	-0.004	0.017

Standard errors of covariance

		S&P		Moody's		Fitch	
		DSCR	LTV	DSCR	LTV	DSCR	LTV
S&P	DSCR	0.007					
	LTV	0.003	0.001				
Moody's	DSCR	0.005	0.004	0.013			
	LTV	0.005	0.003	0.004	0.002		
Fitch	DSCR	0.006	0.004	0.006	0.003	0.004	
	LTV	0.005	0.002	0.004	0.002	0.003	0.002

Tables B.1 and B.2 report maximum likelihood estimates for the first-step parameters of the "Base Specification" (whose structural parameter estimates are in Table 3). All of the first-step parameters are jointly estimated but are reported in two separate tables due to space considerations. Table B.1 shows the estimated joint distribution of the weighted-average reunderwritten DSCR and LTV for each agency (at the deal-level), assuming joint normality.

Table B.2: First-step estimates (base specification): bid functions ($G_B(\cdot)$)

Sieve parameters ($\{\gamma_j\}_{j=1,2,3}$)

Bidder-specific-deal covariates

	S&P		Moody's		Fitch	
	Estimate	Std Error	Estimate	Std Error	Estimate	Std Error
Bidder fixed effect	1.132	0.107	1.116	0.131	1.201	0.106
Reunderwritten DSCR	-0.035	0.036	0.504	0.088	-0.025	0.046
Reunderwritten LTV	-0.533	0.108	0.180	0.125	-0.435	0.095
Bidder produced no pre-sale report	0.144	0.049	0.262	0.077	0.171	0.080
Own share of last 10 deals	0.087	0.059	0.024	0.076	0.0003	0.049
Avg. competitors' share of last 10 deals	-0.057	0.092	-0.008	0.121	-0.071	0.098
Own share of last 3 deals by same bank	-0.075	0.033	-0.029	0.046	-0.030	0.029
Avg. competitors' share of last 3 deals by same bank	-0.044	0.060	-0.049	0.068	-0.104	0.060

Deal covariates

	Estimate	Std Error
Balloon payment (wtd avg)	-0.309	0.076
Cross-collateralization (wtd avg)	-0.076	0.037
Deal size (total principal)	0.013	0.012
Originator HHI	-0.077	0.033
Property type HHI	-0.484	0.073
Region HHI	-0.149	0.053
2001 vintage	0.222	0.036
2002 vintage	0.302	0.031
2003 vintage	0.420	0.037
2004 vintage	0.514	0.037
2005 vintage	0.618	0.037
2006 vintage	0.657	0.034
2007 vintage	0.650	0.041
2010-2012 vintages	0.267	0.044

Covariance of residual (Ω), point estimates

	S&P	Moody's	Fitch
S&P	0.022		
Moody's	0.020	0.020	
Fitch	0.021	0.020	0.023

Standard errors of covariance of residual

	S&P	Moody's	Fitch
S&P	0.037		
Moody's	0.021	0.022	
Fitch	0.021	0.014	0.021

See the note in the previous table. Table B.2 shows the estimated joint distribution of equilibrium bids. The sieve parameters capture the effect of covariates on bidding behavior for individual agencies. The covariance parameters capture the joint distribution of the component of agencies' bids that is not explained by covariates.

Table B.3: Within-sample fit of first-step estimates

Frequency of each bidder being among the auction winners

	S&P	Moody's	Fitch
Empirical mean	0.721	0.719	0.591
Predicted mean	0.676	0.688	0.667

Mean and standard deviation of pivotal bid (b_i^*) over sample of deals, by winner set

Winner Set	Empirical mean	Predicted mean	Empirical std dev	Predicted std dev
Full sample	0.828	0.853	0.067	0.061
S wins in data	0.858	0.895	0.056	0.060
M wins in data	0.849	0.924	0.050	0.049
F wins in data*	0.754	0.832	0.091	0.020
S, M win in data	0.838	0.867	0.062	0.046
S, F win in data	0.834	0.826	0.070	0.062
M, F win in data	0.812	0.858	0.064	0.047
S, M, F win in data	0.812	0.858	0.064	0.047
2000-2003 vintages	0.778	0.816	0.065	0.064
2004-2007 vintages	0.871	0.879	0.037	0.041
2010-2012 vintages	0.806	0.874	0.027	0.040

Table shows fit statistics for the baseline model specification. A unit of observation is an auction. Bottom panel summarizes predicted and empirical pivotal bids for the full sample and for various subsamples defined by the set of winning bidders in the data.

Table B.4: Out-of-sample fit of first-step estimates for validation sample

Frequency of each bidder being among the auction winners

	S&P	Moody's	Fitch
Empirical mean	0.722	0.725	0.589
Predicted mean	0.687	0.610	0.739

Mean and standard deviation of pivotal bid (b_i^*) over sample of deals, by winner set

Winner Set	Empirical mean	Predicted mean	Empirical std dev	Predicted std dev
Full sample	0.828	0.852	0.065	0.062
S wins in data	0.884	0.906	0.010	0.057
M wins in data	0.844	0.923	0.038	0.036
F wins in data*	0.680	0.782	0.042	0.082
S, M win in data	0.831	0.874	0.066	0.055
S, F win in data	0.843	0.830	0.055	0.059
M, F win in data	0.809	0.840	0.066	0.061
S, M, F win in data	0.809	0.840	0.066	0.061
2000-2003 vintages	0.780	0.813	0.063	0.060
2004-2007 vintages	0.868	0.883	0.038	0.045
2010-2012 vintages	0.808	0.861	0.033	0.056

Table shows out-of-sample fit statistics for the validation sample after estimating the baseline model on a random 50-percent sample of the data. Bottom panel summarizes predicted and empirical pivotal bids for the full sample and for various subsamples defined by the set of winning bidders in the data.

Appendix C Alternative specification with shadow ratings specifying subordinate securities

The baseline model in the main text assumes that bidders' "shadow rating" specifies the subordination amount for the AAA security, based on the salience of this figure as an institutional feature of CMBS. In this appendix, we consider an alternative assumption under which the shadow rating stipulates subordination amounts for all of the securities, both AAA and non-AAA.

For tractability, we make an assumption that allows us to continue regarding bids as one-dimensional (as in the baseline case). Namely, we assume that the loss distribution on the pool principal is drawn from a stochastically ordered family of distributions (assumed to be common knowledge). A "bid" can be thought of as the bidder's assertion about the specific distribution of losses, with the stochastic-ordering assumption allowing bids to be ranked along a single dimension.

As in the baseline case, in order for the model to be identified, we must normalize the scaling of the bids to something observable. For this purpose, we define a measure of how favorably an entire deal is rated overall—the weighted-average "rating-implied yield spread" (WARIS). Let ζ_r be the proportion of securities rated r. For a set of ratings $1 \ldots R$:

$$WARIS \;=\; \textstyle\sum_r^R \zeta_r \cdot RIS_r,$$

The constant RIS_r proxies for the level of risk associated with rating r. Following Efing & Hau (2015), we define RIS_r to be the fixed effect associated with rating r in a regression of observed bond yields at issuance on various deal- and security characteristics.[32] We take as given that a lower WARIS is more favorable to the issuer—conceptually, a lower WARIS implies a lower cost of funding due to lower yields that must be paid to investors.

Consider the problem of minimizing the WARIS over alternative choices of ζ_r given some set of constraints on the permissible loss rate for each rating. The stochastic-ordering assumption implies that a bid can equivalently be thought of as stipulating the lowest feasible value of WARIS given such constraints.[33]

The final deal structure is determined in the following manner. Recall that in the

[32]The regression is (yield spread at issuance$_{ir}$) $= \alpha' W_i + \beta' Z_{ir} + RIS_t + \varepsilon_{ir}$, where W_i are observed deal characteristics, Z_{ir} are observed characteristics of securities rated r, and the dependent variable is the yield premium on securities from deal i that have rating r at the time of issuance. Data on security yields are from Commercial Mortgage Alert.

[33]For example, according to the agencies' published investor guidelines, "AAA", "A1" and "B1" have idealized loss rates over four years of 0.001 percent, 0.104 percent, and 7.6175 percent, respectively. See "Probability of Default Ratings and Loss Given Default Assessments for Non-Financial Speculative-Grade Corporate Obligors in the United States and Canada," Moody's Investors Service (August 2006), Appendix 1 (available online).

baseline case, the final AAA share is determined by the K'th lowest winning bid, for an auction with K winners. Analogously, here the issuer minimizes the WARIS with respect to $\{proportion\ of\ securities\ rated\ r\}_{1,...R}$, subject to the constraints stipulated by each of the K winning bidders.

We can specify bounds on the bidders' actual bids, based on the observed ratings of the winning bidders. Let $WARIS^*_{ij}$ denote the WARIS on deal i computed using bidder j's observed ratings.[34] For convenience, we transform $WARIS^*_{ij}$ by a known, monotone-decreasing, normalizing function $q(\cdot)$, in order to have a higher bid correspond to a more favorable outcome for the issuer (as in the baseline model).[35]

- Suppose, there is a single observed winner and, without loss of generality, let this be bidder 1. By the definition of a bid, $b_{i1} \geq q(WARIS^*_{i1})$. Moreover, by the optimality of the issuer's behavior, this inequality must hold strictly: $b_{i1} = q(WARIS^*_{i1})$. Otherwise, the issuer could structure the deal in a way that achieves a lower WARIS while still complying with bidder 1's requirements.

- Suppose there are two observed winners and, without loss of generality, let these be bidders 1 and 2. By the definition of a bid, $b_{i1} \geq q(WARIS^*_{i1})$ and $b_{i2} \geq q(WARIS^*_{i2})$. Moreover, the optimality of the issuer's behavior implies that one of these inequalities must hold strictly: otherwise the issuer could structure the deal in a way that achieves a lower WARIS while still complying with both bidders' requirements. Also, bidder 3 not being a winner implies $b_{i3} \leq \min\{b_{i1}, b_{i2}\}$.

- When all three agencies win, it must be the case that $b_{i1} \geq q(WARIS^*_{i1})$, $b_{i2} \geq q(WARIS^*_{i2})$, and $b_{i3} \geq q(WARIS^*_{i3})$. Almost surely, exactly one of these inequalities must also hold strictly.

We perform the first-step estimation by maximizing a likelihood function that is similar to expression (9), but conditioning on the observed values of $WARIS^*_{ij}$ and the above inequalities, instead of the share of AAA.

The structural estimation is similar to the baseline case. First, for each auction i and each winning bidder j, we use bidder j's optimality condition to solve for the belief that would rationalize bidder i's bidding b^*_{ij}, which we denote by $\xi_{ij}(b^*_{ij}, x_i, w_i, z_{ij}; \theta, \beta)$. As in the baseline specification, we then compute moment conditions based on the expectation of the pivotal bidder's belief, $\hat{u}^*_i$, similar to Equation (10).

[34] Note that, in general, $WARIS^*_{ij} \neq WARIS^*_{ij'}$ for two bidders j and j'. This is the case due to the existence of "split ratings"—where the observed rating differs across rating agencies—for securities other than the AAA one.

[35] We define $q(\cdot)$ it to be an affine transformation such that $q(WARIS^*_{ij})$ has the same mean and standard deviation across deals i and bidders j as the observed AAA share b^*_i.

Appendix D Estimates for Additional Specifications

Tables D.1 and D.2 report the first-step estimates corresponding to the "Alternative Specification" reported in Table 3).

Tables D.3 and D.4 report the first-step estimates corresponding to the robustness check involving bids that specify the structure of both AAA and non-AAA securities, reported in Table 5.

Tables D.5 and D.6 report the post-estimation regressions of ex post deal outcomes on the ordinal distortion measure implied by the robustness checks in Sections 8.1 and 8.2, respectively.

Tables D.7, D.8, and D.9 report the first-step and structural estimates for another specification in which we endogenize the number of winning bidders. To accomplish this, we modify the issuer's maximization (Equation 1) as follows[36]:

$$\max_{d,K}\left[\left(\min\left\{B(d): d\in\{0,1\}^J, \sum_{j=1}^{J} d_j = K\right\}\right)(1+\kappa_2 \mathbb{1}\{K>1\}+\kappa_3\mathbb{1}\{K=3\})\right]$$

The objective function from Equation 1 is in the first parenthesis. The term in the second parenthesis is new. This equation implies the issuer maximizes the pivotal bid times a function of how many ratings the issuer obtains. This modified specification captures the following tradeoff faced by the issuer. On the one hand, choosing fewer winners increases the AAA proportion, which equals the K'th highest bid when the issuer chooses K winners. On the other hand, investors may place a premium on deals with more ratings—either because they value corroborating opinions or because they are sophisticated and recognize issuers' incentive to ratings shop. We do not explicitly model investor demand, but rather, specify that the issuer's payoff depends in an exogenous way on the number of bids, with κ_2 representing the premium for having two ratings versus only one, and κ_3 representing the premium for having three ratings versus two.

In principle, the issuer's premia on having at least two ratings (κ_2) or three ratings (κ_3) are identified by the relative frequency of auctions for which there are one, two, or three winners, and can be estimated jointly with the remaining first-step parameters. However, because the z_{ij}'s for losing bidders are known only in distribution, we have only weak identification of κ_2 and κ_3 separately from the degree of correlation in the bids (determined by the first-step parameters Ω and $\{\gamma\}_{j=1,2,3}$).[37] Intuitively, both greater correlation in the

[36] This modification changes the relationship between an agency's bid and the probability of winning, which we account for in computing the expectation terms in the first-order condition (4).

[37] The extent of correlation in the bids is determined by the covariance matrix Ω and by the magnitude of the coefficients for the bidder-deal-specific covariates relative to the magnitude of those for deal-specific

bids (which compresses the various order statistics of the bid profile) and a greater issuer premium on having more ratings would tend to result in more winners being selected. To finesse this issue, we do not attempt to estimate κ_2 and κ_3. Rather, we fix their values at a level that is *higher* than seems reasonable (5 percent and 2.5 percent respectively)—which maximally alters the likelihood function for the remaining parameters, relative to the base specification in which the number of winners is exogenous. Therefore, if the number of winners were truly endogenous, the difference between these estimates and the base specification estimates would "bound" the impact of erroneously assuming an exogenous number of winners. In fact, we do not find any qualitative differences between the estimates that endogenize the number of winners and the base specification.

covariates among the elements of the parameter vector $\{\gamma\}_{j=1,2,3}$.

Table D.1: First-step estimates (alternative specification): distribution of presale-report variables

Means (μ_z) relative to sample means

	S&P		Moody's		Fitch	
	DSCR	LTV	DSCR	LTV	DSCR	LTV
Estimate	1.350	0.928	1.083	0.885	1.288	0.861
Std Error	0.043	0.007	0.025	0.010	0.027	0.008

Covariances (Ω_z)

		S&P		Moody's		Fitch	
		DSCR	LTV	DSCR	LTV	DSCR	LTV
S&P	DSCR	0.145					
	LTV	-0.008	0.011				
Moody's	DSCR	0.013	-0.003	0.091			
	LTV	0.005	0.011	-0.008	0.018		
Fitch	DSCR	0.069	-0.006	0.042	-0.002	0.062	
	LTV	0.004	0.011	-0.007	0.016	-0.004	0.017

Standard errors of covariance

		S&P		Moody's		Fitch	
		DSCR	LTV	DSCR	LTV	DSCR	LTV
S&P	DSCR	0.007					
	LTV	0.004	0.001				
Moody's	DSCR	0.006	0.006	0.013			
	LTV	0.005	0.002	0.004	0.002		
Fitch	DSCR	0.005	0.003	0.005	0.003	0.003	
	LTV	0.005	0.002	0.004	0.002	0.003	0.002

Tables D.1 and D.2 report maximum likelihood estimates for the first-step parameters of the "Alternative Specification" (whose structural parameter estimates are in Table 3). All of the first-step parameters are jointly estimated but are reported in two separate tables due to space considerations. Table D.1 shows the estimated joint distribution of the weighted-average reunderwritten DSCR and LTV for each agency (at the deal-level), assuming joint normality.

Table D.2: First-step estimates (alternative specification): bid functions

Sieve parameters ($\{\gamma_j\}_{j=1,2,3}$)

Bidder-deal-specific covariates

	S&P		Moody's		Fitch	
	Estimate	Std Error	Estimate	Std Error	Estimate	Std Error
Bidder fixed effect	1.034	0.099	0.950	0.123	1.077	0.098
Reunderwritten DSCR	0.222	0.078	0.499	0.091	0.007	0.045
Reunderwritten LTV	-0.207	0.123	0.247	0.126	-0.267	0.102
Bidder produced no pre-sale report	0.165	0.066	0.229	0.089	0.189	0.083
Own share of last 10 deals by different bank	0.032	0.046	0.076	0.068	0.028	0.040
Avg. competitors' share of last 10 deals by different bank	0.030	0.077	0.101	0.102	-0.064	0.080

Deal covariates

	Estimate	Std Error
Balloon payment (wtd avg)	-0.303	0.079
Cross-collateralization (wtd avg)	-0.032	0.037
Deal size (total principal)	0.008	0.012
Originator HHI	-0.088	0.032
Property type HHI	-0.383	0.076
Region HHI	-0.134	0.055
2001 vintage	0.216	0.036
2002 vintage	0.284	0.032
2003 vintage	0.383	0.036
2004 vintage	0.484	0.038
2005 vintage	0.580	0.037
2006 vintage	0.623	0.036
2007 vintage	0.601	0.045
2010-2012 vintages	0.254	0.047

Covariance of residual (Ω), point estimates

	S&P	Moody's	Fitch
S&P	0.020		
Moody's	0.019	0.020	
Fitch	0.021	0.021	0.024

Standard errors of covariance of residual

	S&P	Moody's	Fitch
S&P	0.011		
Moody's	0.008	0.012	
Fitch	0.013	0.012	0.026

Tables D.1 and D.2 report maximum likelihood estimates for the first-step parameters of the "Alternative Specification" (whose structural parameter estimates are in Table 3). All of the first-step parameters are jointly estimated but are reported in two separate tables due to space considerations. Table D.2 shows the estimated equilibrium bidding behavior. The sieve parameters capture the effect of covariates on bidding behavior for individual agencies. The covariance parameters capture the joint distribution of the component of agencies' bids that is not explained by covariates.

Table D.3: First-step estimates for robustness check with bids specifying structure of both AAA and non-AAA securities: distribution of presale-report variables

Means (μ_z) relative to sample means

	S&P		Moody's		Fitch	
	DSCR	LTV	DSCR	LTV	DSCR	LTV
Estimate	1.441	0.919	1.207	0.885	1.248	0.860
Std Error	0.038	0.007	0.026	0.009	0.026	0.009

Covariances (Ω_z)

		S&P		Moody's		Fitch	
		DSCR	LTV	DSCR	LTV	DSCR	LTV
S&P	DSCR	0.127					
	LTV	-0.005	0.011				
Moody's	DSCR	0.051	-0.007	0.076			
	LTV	0.006	0.011	-0.008	0.018		
Fitch	DSCR	0.067	-0.005	0.045	-0.003	0.062	
	LTV	0.005	0.011	-0.008	0.016	-0.004	0.017

Standard errors of covariance

		S&P		Moody's		Fitch	
		DSCR	LTV	DSCR	LTV	DSCR	LTV
S&P	DSCR	0.005					
	LTV	0.003	0.001				
Moody's	DSCR	0.004	0.002	0.008			
	LTV	0.004	0.002	0.003	0.002		
Fitch	DSCR	0.005	0.003	0.005	0.003	0.003	
	LTV	0.004	0.002	0.003	0.002	0.003	0.002

Tables D.3 and D.4 report maximum likelihood estimates for the first-step parameters of the robustness check with bids specifying structure of both AAA and non-AAA securities: distribution of presale-report variables (whose structural parameter estimates are in Table 5). All of the first-step parameters are jointly estimated but are reported in two separate tables due to space considerations. Table D.3 shows the estimated joint distribution of the weighted-average reunderwritten DSCR and LTV for each agency (at the deal-level), assuming joint normality.

Table D.4: First-step estimates for robustness check with bids specifying structure of both AAA and non-AAA securities: bid functions

Sieve parameters ($\{\gamma_j\}_{j=1,2,3}$)

Bidder-deal-specific covariates

	S&P		Moody's		Fitch	
	Estimate	Std Error	Estimate	Std Error	Estimate	Std Error
Bidder fixed effect	1.349	0.083	1.377	0.077	1.376	0.086
Reunderwritten DSCR	0.002	0.033	0.006	0.039	0.009	0.047
Reunderwritten LTV	-0.045	0.069	-0.016	0.065	-0.034	0.063
Bidder produced no pre-sale report	0.038	0.019	0.030	0.009	0.023	0.019
Own share of last 10 deals	0.005	0.045	-0.035	0.043	-0.026	0.042
Competitors' share of last 10 deals	-0.079	0.081	-0.091	0.080	-0.073	0.083
Own share of last 3 deals by same bank	0.002	0.024	0.018	0.025	0.016	0.022
Avg. competitors' share of last 10 deals by different bank	0.041	0.046	0.033	0.044	0.014	0.042

Deal covariates

	Estimate	Std Error
Balloon payment (wtd avg)	0.006	0.060
Cross-collateralization (wtd avg)	-0.035	0.020
Deal size (total principal)	-0.015	0.007
Originator HHI	-0.068	0.024
Property type HHI	0.027	0.051
Region HHI	-0.032	0.047
2001 vintage	0.140	0.032
2002 vintage	0.189	0.039
2003 vintage	0.282	0.034
2004 vintage	0.388	0.027
2005 vintage	0.381	0.028
2006 vintage	0.376	0.027
2007 vintage	0.455	0.036
2010-2012 vintages	0.320	0.050

Covariance of residual (Ω), point estimates

	S&P	Moody's	Fitch
S&P	0.021		
Moody's	0.020	0.021	
Fitch	0.020	0.020	0.021

Standard errors of covariance of residual

	S&P	Moody's	Fitch
S&P	0.019		
Moody's	0.009	0.003	
Fitch	0.011	0.006	0.013

Tables D.3 and D.4 report maximum likelihood estimates for the first-step parameters of the robustness check with bids specifying structure of both AAA and non-AAA securities: distribution of presale-report variables (whose structural parameter estimates are in Table 5). All of the first-step parameters are jointly estimated but are reported in two separate tables due to space considerations. Table D.4 shows the estimated joint distribution of the weighted-average reunderwritten DSCR and LTV for each agency (at the deal-level), assuming joint normality.

Table D.5: Tobit regressions for ex post deal outcomes (principal losses and interest shortfall) on distortion (λ_i) and control variables, for robustness check involving alternative bidder preferences (i.e., the uniform-"price" specification)

	Full sample of deals		**Deals rated by Moody's and S&P**	
	Estimate	Std. Error	Estimate	Std. Error
Distortion	0.031	0.013	0.051	0.021
Splines for AAA share	Included		Included	
Deal rated by 3 agencies	0.0003	0.004		
Deal rated by 1 agency	-0.022	0.004		
Constant	-0.010	0.014	-0.017	0.022
Square root of error variance	0.026	0.003	0.020	0.001
Observations	579		203	

	Deals rated by S&P and Fitch		**Deals rated by Moody's and Fitch**	
Distortion	0.004	0.019	0.068	0.027
Splines for AAA share	Included		Included	
Constant	0.014	0.033	-0.024	0.042
Square root of error variance	0.031	0.006	0.026	0.004
Observations	147		131	

Dependent variable is sum of principal loss and interest payment shortfalls on the deal's loan pool, as of the censoring date (September 2012), expressed as a share of the original pool principal. Letting ψ denote linear coefficients and ε a normal error, the assumed model is:

$$\begin{aligned}\text{dependent variable} &= \psi'(\text{covariates}) + \varepsilon \text{ if } \psi'(\text{covariates}) + \varepsilon > 0,\\ &= 0 \text{ otherwise.}\end{aligned}$$

* Standard errors are bootstrapped and take into account estimation error in both the first step and in structural estimation.

Table D.6: Tobit regressions for ex post deal outcomes (principal losses and interest shortfall) on distortion (λ_i) and control variables, for robustness check involving bids that specify the structure of non-AAA securities

	Full sample of deals		**Deals rated by Moody's and S&P**	
	Estimate	Std. Error	Estimate	Std. Error
Distortion	0.001	0.002	-0.013	0.004
Splines for AAA share	Included		Included	
Deal rated by 3 agencies	-0.0055	0.002		
Deal rated by 1 agency	-0.025	0.003		
Constant	0.014	0.006	0.033	0.013
Square root of error variance	0.026	0.003	0.020	0.002
Observations	579		203	

	Deals rated by S&P and Fitch		**Deals rated by Moody's and Fitch**	
Distortion	0.013	0.007	0.023	0.009
Splines for AAA share	Included		Included	
Constant	0.018	0.034	0.013	0.011
Square root of error variance	0.031	0.007	0.026	0.003
Observations	147		131	

Dependent variable is sum of principal loss and interest payment shortfalls on the deal's loan pool, as of the censoring date (September 2012), expressed as a share of the original pool principal. Letting ψ denote linear coefficients and ε a normal error, the assumed model is:

$$\begin{aligned} \text{dependent variable} &= \psi'(\text{covariates}) + \varepsilon \text{ if } \psi'(\text{covariates}) + \varepsilon > 0, \\ &= 0 \text{ otherwise.} \end{aligned}$$

* Standard errors are bootstrapped and take into account estimation error in both the first step and in structural estimation.

Table D.7: First-step estimates (specification endogenizing number of winning bidders): distribution of presale-report variables

Sieve parameters ($\{\gamma_j\}_{j=1,2,3}$)

Bidder-specific covariates

	S&P		Moody's		Fitch	
	Estimate	Std Error	Estimate	Std Error	Estimate	Std Error
Bidder fixed effect	1.147	0.098	1.221	0.095	1.200	0.101
Reunderwritten DSCR	-0.004	0.033	0.027	0.038	0.044	0.043
Reunderwritten LTV	-0.275	0.083	-0.137	0.074	-0.186	0.072
Bidder produced no pre-sale report	0.069	0.016	0.066	0.010	0.048	0.028
Own share of last 10 deals	0.050	0.048	-0.021	0.046	0.010	0.048
Competitors' share of last 10 deals	-0.017	0.089	-0.099	0.091	-0.060	0.098
Own share of last 3 deals by same bank	-0.051	0.027	-0.027	0.029	-0.027	0.026
Avg. competitors' share of last 10 deals by different bank	-0.050	0.053	-0.047	0.054	-0.079	0.054

Deal covariates

	Estimate	Std Error
Balloon payment (wtd avg)	-0.330	0.068
Cross-collateralization (wtd avg)	-0.072	0.039
Deal size (total principal)	0.008	0.011
Originator HHI	-0.092	0.030
Property type HHI	-0.480	0.063
Region HHI	-0.015	0.044
2001 vintage	0.212	0.030
2002 vintage	0.291	0.027
2003 vintage	0.405	0.032
2004 vintage	0.487	0.035
2005 vintage	0.578	0.033
2006 vintage	0.613	0.030
2007 vintage	0.592	0.039
2010-2012 vintages	0.274	0.045

Covariance of residual (Ω), point estimates

	S&P	Moody's	Fitch
S&P	0.021		
Moody's	0.020	0.022	
Fitch	0.020	0.020	0.022

Standard errors of covariance of residual

	S&P	Moody's	Fitch
S&P	0.021		
Moody's	0.011	0.004	
Fitch	0.013	0.009	0.017

Tables D.7 and D.8 report maximum likelihood estimates for the first-step parameters of the specification endogenizing the number of winning bidders (discussed earlier in this Appendix), fixing $\kappa_2 = 0.05$ and $\kappa_3 = 0.025$. (Structural estimates are reported in Table D.9). All of the first-step parameters are jointly estimated but are reported in two separate tables due to space considerations. Table D.7 shows the estimated joint distribution of the weighted-average reunderwritten DSCR and LTV for each agency (at the deal-level), assuming joint normality.

Table D.8: First-step estimates (specification endogenizing number of winning bidders): bid functions

Sieve parameters ($\{\gamma\}_{j=1,2,3}$)

Bidder-specific covariates

	S&P		Moody's		Fitch	
	Estimate	Std error	Estimate	Std error	Estimate	Std error
Bidder fixed effect	0.8410	0.1157	0.9212	0.1136	0.8985	0.1139
Reunderwritten DSCR	-0.0336	0.0237	0.0167	0.0275	0.2781	0.0359
Reunderwritten LTV	-0.8494	0.0605	-0.6242	0.0499	-0.4438	0.0496
Bidder produced no pre-sale report	-0.0346	0.0252	0.0065	0.0141	-0.0340	0.0241
Own share of last 10 deals	0.1728	0.0411	0.0489	0.0405	0.0717	0.0426
Avg competitors' share of last 10 deals*	0.1426	0.0714	0.0899	0.0740	0.1014	0.0807
Own share of last 3 deals by same bank	-0.0330	0.0293	-0.0083	0.0301	-0.0167	0.0291
Avg. competitors' share of last 3 deals by same bank*	-0.0351	0.0574	-0.0064	0.0560	-0.0233	0.0580

Deal covariates

	Estimate	Std error
Balloon payment (wtd avg)	-0.3140	0.0887
Cross-collateralization (wtd avg)	-0.0927	0.0445
Deal size (total principal)	0.0301	0.0158
Originator HHI	-0.0748	0.0384
Property type HHI	-0.4701	0.0769
Region HHI	-0.1205	0.0505
2001 vintage	0.2347	0.0375
2002 vintage	0.3191	0.0350
2003 vintage	0.4037	0.0419
2004 vintage	0.4987	0.0436
2005 vintage	0.6018	0.0411
2006 vintage	0.6566	0.0393
2007 vintage	0.6696	0.0452
2010 vintage	0.3587	0.0550

Covariance of residual (Ω), point estimates

	S&P	Moody's	Fitch
S&P	0.0267		
Moody's	0.0256	0.0271	
Fitch	0.0251	0.0253	0.0262

Standard errors of covariance of residual

	S&P	Moody's	Fitch
S&P	0.0132		
Moody's	0.0073	0.0099	
Fitch	0.0076	0.0064	0.0089

Tables D.7 and D.8 report maximum likelihood estimates for the first-step parameters of the specification endogenizing the number of winning bidders (discussed earlier in this Appendix), fixing $\kappa_2 = 0.05$ and $\kappa_3 = 0.025$. (Structural estimates are reported in Table D.9). All of the first-step parameters are jointly estimated but are reported in two separate tables due to space considerations. Table D.8 shows the estimated equilibrium bidding behavior. The sieve parameters capture the effect of covariates on bidding behavior for individual agencies. The covariance parameters capture the joint distribution of the component of agencies' bids that is not explained by covariates.

Table D.9: Structural estimates (specification endogenizing number of winning bidders)

	Estimate	Std. Err.
Bidder-specific covariates, commonly observed (β_1)		
Own share of last 10 deals	0.317	0.480
Avg. competitors' share of last 10 deals	1.028	0.939
Own share of last 3 deals by same bank	0.201	0.429
Avg. competitors' share of last 3 deals by same bank	-0.310	0.493
Deal covariates (β_2)		
Constant	0.596	1.074
Balloon payment (wtd avg)	-0.699	0.446
Cross-collateralization (wtd avg)	-0.052	0.063
Deal size (total principal)	0.013	0.010
Originator HHI	-0.138	0.051
Property type HHI	-0.591	0.341
Region HHI	0.081	0.143
2001 vintage	0.253	0.055
2002 vintage	0.381	0.069
2003 vintage	0.461	0.054
2004 vintage	0.591	0.053
2005 vintage	0.602	0.060
2006 vintage	0.664	0.052
2007 vintage	0.572	0.092
2010-2012 vintages	0.540	0.180
Bidder-specific covariates, private information (β_3)		
Reunderwritten DSCR	0.034	0.160
Reunderwritten LTV	-0.070	0.288
Bidder produced no pre-sale report	0.003	0.172

Covariance of residual u_{ij}

	S&P	Moody's	Fitch
S&P	0.122		
Moody's	0.094	0.139	
Fitch	0.092	0.110	0.129

Number of observations: 578

Table shows structural estimates for specification endogenizing the number of winning bidders (discussed earlier in this Appendix), fixing $\kappa_2 = 0.05$ and $\kappa_3 = 0.025$. The distribution of the residual u_{ij} is computed by simulating the distribution of beliefs for the full set of bidders and netting out the effects of the covariates. First-step estimates corresponding to the specification reported in this table are in Tables D.7 and D.8.

Appendix E Solving for Counterfactual Equilibrium of Bidding Game

Existence of Pure Strategy Bayesian Nash Equilibrium

In this Appendix, we argue that a pure-strategy Bayesian Nash Equilibrium (PSNE) exists for the bidding game described in the model. If the possible set of actions were discrete (e.g., if bidders could bid only in increments of 0.01), the existence of a PSNE would be guaranteed so long as the game satisfies the Single-Crossing Condition (SCC) and certain other regularity conditions (see Definition 3 and Theorem 1 in Athey, 2001). The SCC can easily be shown to hold in our setup, and stipulates that, for each player $j = 1, \dots, J$, whenever every opponent $j' \neq j$ uses a strategy that is nondecreasing in its type, player j's objective function satisfies the *single crossing property of incremental returns* in (b_{ij}, t_{ij}). Because the objective function $\pi_{ij}(t_{ij}, b_i)$ is differentiable, it suffices to observe that $\frac{\partial \pi_{ij}(t_{ij}, b_i)}{\partial t_{ij} \partial b_{ij}} > 0$.

In the case of continuous actions, existence of a PSNE could be shown constructively by taking the limit of the finite-action equilibrium for successively finer action sets *if* the limit of this series were guaranteed to be an equilibrium of the continuous game. A complication arises in bidding games, such as in our setup, because the outcome (namely, the set of winners) is discontinuous in the actions. However, this problem goes away if, in the limit as the action becomes successively finer, "mass points" do not arise and the payoffs are continuous. The conditions for this to hold are discussed in Theorem 6 of Athey (2001), and are either standard or hold trivially in the current setting by virtue of the assumption of private values.

Solution method

The solution method falls under the general approach of mathematical programming with equilibrium constraints (MPEC). Let $\varphi_{ij}(b_{ij}) = \tilde{b}_{ij}$ denote the inverse bid function for bidder j in auction i (we assume a bidder's bid function is monotone and thus invertible). Following Hubbard & Paarsch (2009) and Bajari (2001), we approximate j's inverse bid function $\varphi_{ij}(b_{ij})$ by $\psi'_{ij} f(b_{ij})$, where $f(\cdot)$ is a family of basis functions (which we choose to be Chebyshev polynomials). Solving for the PSNE entails finding coefficients $\{\psi_{ij}\}$ for bidders $j = 1, \dots J$ that best fit a set of equilibrium conditions, evaluated on a set of grid points over the domain of possible bids. For all grid points b and b', the imposed equilibrium conditions are as follows:

- Monotonicity: $b' > b \;\Leftrightarrow\; \psi'_{ij} f(b') > \psi'_{ij} f(b)$.

- Optimality: setting $b_{ij} = b$ and $\tilde{b}_{ij} = \psi'_{ij} f(b)$ satisfies bidder j's first-order condition (4).

- Individual rationality: $V - (\Phi^{-1}(b) - w_{ij}\beta_1 - \psi'_{ij} f(b))^2 > 0$, where $\Phi^{-1}(\cdot)$ and $\beta'_3 w_{ij}$ are as defined in expression (3).

Because the support of bidders' type distributions is the real number line and thus infinite, in order to make the solution method feasible, we truncate the bidders' type spaces from above at an auction-specific truncation point $\bar{b}_i$ (chosen to be sufficiently high such that it is in the tails of the type distributions), and "fix" the upper bound of the domain of bids by imposing the boundary condition $\psi'_{ij} f(\bar{b}_i) = \bar{b}_i$. The lower bound of the domain of bids is determined in equilibrium, and, in principle could be estimated by imposing an additional boundary condition (see Hubbard & Paarsch, 2009). However, we found that doing so resulted in numerically unstable results. Instead, we attempt to estimate the inverse bid function only for bids greater than or equal to 0.40, a lower cutoff that we chose based on the observation that 0.40 is less than the minimum observed pivotal bid, 0.508.

Simulation is central to the computation. We simulate the joint distribution of bids conditional on x and w, treating z and ε in expression 8 as random variables distributed according to the first-stage estimates. We then compute the implied joint distribution of types based on the bidders' first-order conditions. At a given bid value, bidder j's first-order condition depends on the distribution of the competitors' bids. For each bidder j, we specify a discrete grid of T bid values $\{a(t)_{ij}, t = 1 \ldots T\}$ spanning a range that covers all but the extreme tails of the marginal distribution of j's bid distribution.

We use MPEC (Judd & Su, 2012) to minimize:

$$\sum_{j=1,2,3} \sum_{t=1,\ldots,T} [\omega(t,1)\varsigma_1 + \omega(t,2)\varsigma_2 + \omega(3)\varsigma_3]$$

where ς_i refers to the squared norm of the violation of monotonicity, optimality and individual rationality for $i = 1, 2, 3$. The terms $\omega(t,1)$, $\omega(t,2)$ and $\omega(3)$ are weights that we chose heuristically. In practice, we do this computation for only one auction, a representative auction as described in the text.