Collateral Value Uncertainty and Mortgage Credit Provision

Erica Xuewei Jiang and Anthony Lee Zhang*

Abstract

Houses with higher value uncertainty receive less mortgage credit: mortgages backed by these houses are more likely to be rejected, have higher interest rates, and have lower loan-to-price ratios. The relationship between house value uncertainty and credit availability is driven partly by a classic channel in which uncertainty lowers debt recovery rates, and partly by a novel channel where more uncertain appraisals make regulatory constraints on loan size more likely to bind. We build a structural model to quantify the effects of each channel, and show how a shift toward computerized asset appraisals could influence credit access.

*Jiang is from the University of Southern California, Marshall School of Business. erica.jiang@marshall.usc.edu. Zhang is from the University of Chicago Booth School of Business. anthony.zhang@chicagobooth.edu. Pranav Garg, Willy Yu-Shiou Lin, Yijing Ren, Michael Yip, and Zichen Zhao provided outstanding research assistance. We are grateful to Kenneth Ahern, Jason Allen, Eric Budish, Sylvain Catherine, Anthony DeFusco, Jason Donaldson, Kevin Duncan, Marco Giacoletti, Adam Guren, Larry Harris, Mark Jansen, Ben Keys, Sam Kruger, David Low, Gregor Matvos, Tim McQuade, Norm Miller, Andrii Parkhomenko, Giorgia Piacentino, Tomasz Piskorski, Michael Reher, Dan Ringo, Amir Sufi, Selale Tuzel, Joe Vavra, Daniel Greenwald, Edward Glaeser, Jinyuan Zhang, Arthur Korteweg, Rodney Ramcharan, participants at Columbia Workshop in New Empirical Finance, the Guthrie Center Housing and Macroeconomics Conference, USC brownbag, UIUC, the Bank of Canada, NUS, UChicago Booth, the Craig W. Holden Memorial Conference, the FRB Philadelphia Mortgage Market Research Conference, UCLA, the FHFA EconSummit, SFS Cavalcade, PUC Chile, NBER SI Real Estate, Georgetown University, McGill University, FRB Chicago, and FRB Boston for comments. We are grateful for financial support from the Fama-Miller Center for Research in Finance at the University of Chicago Booth School of Business. TransUnion (who provided some of the data for this project) has the right to review the research before dissemination to ensure it accurately describes TransUnion data, does not disclose confidential information, and does not contain material it deems to be misleading or false regarding TransUnion, TransUnion's partners, affiliates or customer base, or the consumer lending industry. *First version: August, 2021*

1 Introduction

Households and firms often finance investments in durable assets using collateralized debt. A sizable literature has shown that some assets are better collateral than others, are able to sustain more debt, and thus are easier for households and firms to purchase (Titman and Wessels, 1988; Shleifer and Vishny, 1992; Rampini and Viswanathan, 2010, 2013, 2020). This paper analyzes how collateral values influence credit available in the US residential mortgage market. Mortgages are essential for homeownership: first-time homebuyers in the US borrow over 80% of house prices in mortgage credit on average.[1] In this paper, we ask: What kinds of houses are better collateral for debt? What kinds of households live in these houses? How do housing collateral values influence credit availability: are the channels of effect basically efficient, or do they reflect the effects of regulatory distortions?

Using rich residential property transaction data from 2000 to 2020 in the US, we show that older and less standardized houses have higher price dispersion; that is, these houses' sale prices are harder to predict based on house characteristics. Mortgages backed by these houses are more likely to be rejected, receive worse interest rates, and have lower loan-to-price ratios (LTP). We propose that this relationship is driven by two main channels. One is a classical *collateral recovery* channel: less standardized houses have lower expected recovery rates, so lenders will offer less credit against these houses even in frictionless markets. The other is a novel *regulatory* channel. Regulators impose loan-to-value (LTV) restrictions on most US mortgages, limiting the borrowing amount to a fraction of the house's estimated value. The collateral value used in LTV calculations is set to the lesser of the house transaction price and the house *appraisal value*. High-value-uncertainty houses have noisier appraisals, making regulatory constraints on loan sizes more likely to bind for these houses. While its effects qualitatively resemble the classical collateral recovery effects, the appraisal channel is a potentially *inefficient* channel through which value uncertainty influences credit availability. We build a structural model, illustrating quantitatively how each channel influences different measures of credit availability, how value uncertainty affects consumer willingness-to-pay for houses, and how policy interventions could influence mortgage credit availability and consumer welfare.

[1] Source: CFPB report, Market Snapshot: First-Time Homebuyers.

We begin by estimating value uncertainty for all individual houses transacted in the U.S. from 2000 to 2020 by measuring what kinds of houses tend to have larger errors when priced using a hedonic model (Case and Shiller, 1989; Giacoletti, 2021; Sagi, 2021; Hartman-Glaser and Mann, 2017; Sklarz and Miller, 2016; Buchak et al., 2020). We document a number of stylized facts, some of which have been demonstrated in prior literature. Price dispersion is highly persistent over time at the zipcode level, suggesting that the cross-zipcode variation is driven by persistent differences in the characteristics of local housing stocks, rather than time-varying market conditions. Price dispersion is high for older and less standardized houses and for houses that trade in thinner markets. High-dispersion houses also have noisier appraisal values: they are more likely to under-appraise, and under-appraise by larger amounts on average.

We find that value uncertainty is associated with mortgage credit provision on multiple margins. The largest effects are on mortgage rejection rates: a 1SD increase in house price dispersion is associated with a 25% increase in the prevalence of collateral-related mortgage rejections, and a 10% increase in total rejection rates. We also find statistically significant, but economically smaller, effects on mortgage interest rates and loan-to-price ratios: a 1SD dispersion increase is associated with around 0.9bps higher interest rates, and around 20bps lower loan-to-price ratios. The relationship between value uncertainty and credit provision has important distributional consequences: value uncertainty tends to be high precisely in areas with low-income and minority households, implying that value uncertainty may limit mortgage credit availability to some of the households that are most dependent on credit for homeownership.

Lenders could in principle lend less against high-dispersion houses, not because the *houses* are worse, but because *buyers* of high-dispersion houses systematically have higher credit risks. In the cross-section, we find that high-dispersion areas indeed have buyers with lower incomes and FICO scores. To address this concern, we construct instruments for price dispersion, based on the heterogeneity of houses relative to their local housing stock. Intuitively, when a zipcode has very nonstandardized houses, the market for any given house will tend to be thin, and price dispersion will tend to be high. Our instrument is correlated with price dispersion but is not correlated with ex-ante buyer creditworthiness. Moreover, buyers of houses with high instrumented price dispersion are not ex-post more likely to default on mortgages. These results thus suggest that our findings are driven by houses being worse collateral, rather than buyers of these houses having higher credit risks.

Why is house price dispersion associated with mortgage credit provision? We hypothesize that there are two economic channels driving this relationship. The first is a classic "collateral recovery" channel. The value of debt is concave in the sale price of the collateral asset, since lenders cannot keep the upside if collateral sells for more than the value of outstanding debt, but suffer a loss if collateral sells for less. Thus, lenders rationally offer less credit against high-value-uncertainty houses. The second channel, which we believe is novel to our paper, is a regulatory effect based on appraisals. Loan-to-value ratios are an important factor that regulators use to determine the risk of lending. For loans held on banks' balance sheets, the amount of capital required by the regulators typically depends on the LTV ratios.[2] For securitized loans, there are LTV constraints on whether they can receive government guarantees, and also the guarantee fees also depend on LTV ratios.[3] In both cases, houses backing mortgages undergo appraisals, and the lesser of the appraised value and the transaction price is used as the collateral value in LTV calculations. Higher-dispersion houses tend to have noisier appraisals, so these regulatory appraisal constraints are more likely to bind, limiting mortgage credit availability for such houses.

These two channels both directionally imply higher value uncertainty should be associated with less mortgage credit; however, they are driven by different economic forces. The first channel is a result of the fair market pricing of collateralized debt and should in principle function even in frictionless markets. The second channel, on the other hand, is a product of regulation; it is not necessarily efficient, and in theory its incidence could vary depending on the specifics of how appraisal regulation is implemented. It is thus important to disentangle the extent to which each channel contributes to the outcomes we observe.

We thus build a structural model to quantify how each of the two channels influences mortgage market outcomes. In the model, competitive lenders offer menus of interest rate-LTP pairs to a borrower, such that lenders break even, given the exogenous risk of default and expected recovery rates from the house upon foreclosure. The borrower chooses a target loan size from the menu. The house then undergoes an appraisal. We model appraisals

[2] See, for example, the BIS document on bank capital requirements.

[3] See, for example, Fannie Mae loan eligibility requirements; and loan-level price adjustment matrix.

as noisy, upward-biased signals of house prices, consistent with the distribution of house appraisals in practice. If the house over-appraises, the borrower proceeds with the mortgage as planned. If the appraisal is sufficiently low that the mortgage would violate regulatory LTV constraints, the buyer must choose to either make a costly increase in her down payment or pay a fixed cost to renege on the transaction and find a new house, which we interpret as a mortgage rejection. The tradeoff homebuyers face is that larger mortgages improve consumption smoothing, but increase the risk of under-appraisals. When price dispersion is higher, lenders offer worse menus, and under-appraisal risk is larger, leading to more mortgage rejections, higher interest rates, and lower LTPs.

We then calibrate the model to data, matching moments on how rate menus, appraisal distributions, mortgage rejection rates, and loan size depend on house price dispersion. The calibrated model suggests that the collateral recovery channel mainly drives changes in loan interest rates, whereas appraisal risk is the main driver of increased mortgage failures and lowered loan-to-price ratios. Our calibrated model also implies that price dispersion and underappraisal risk present nontrivially large costs to consumers. To offset higher rates and underappraisal risks, we find that consumers in the highest decile of counties by price dispersion would have to face 1.679% lower prices for identical houses, to attain the same expected utility as consumers in the lowest decile of counties.

The appraisal channel is a product of regulation, implying that shifts in the regulatory environment around appraisals could change the effects of value uncertainty on credit provision. One such change is the ongoing transition toward automated appraisals.[4] We estimate the impact of automated appraisals in our model under two different sets of assumptions. First, suppose appraisal software simply removed human appraisers' tendency to upward bias appraisals, so that appraisal were symmetrically distributed around transaction prices. In our model, this would actually dramatically increase underappraisal risk: we find that naively removing human appraises' biases would lead mortgage failures to increase by roughly 10.540% to 13.561%, and that prices would have to decrease substantially, by 5.347% to 6.262%, for consumers attain the same expected utility as they would under human appraisers. On the other hand, suppose automated appraisals were calibrated to maintain the status quo rate of underappraisals, but were half as noisy than human appraisals. We find that underap-

[4]In 2021, the FHFA announced that banks and mortgage lenders could use automated appraisal software in the place of human appraisals.

praisal pressure would then be somewhat alleviated: mortgage failure rates would decrease modestly, by up to 1.600%, and consumers would be willing to pay up to 0.694% more for identical houses, relative to what they would be willing to pay under human appraisers. Our results thus imply that automated appraisals can thus have both positive and negative effects on mortgage credit availability, depending on how they are implemented.

A central theme in the study of financial institutions is that regulatory interventions designed to address specific issue tend to produce other issues as side effects. Regulators commonly impose LTV restrictions on collateralized debt contracts. In the US, for example, the Fed's Regulation T limits the leverage brokers and dealers can offer on margin loans; regulators also impose LTV requirements on bank loans backed by land and commercial real estate, as well as housing.[5] LTV restrictions are relatively straightforwards to enforce for debt backed by liquid financial assets with "live" prices, such as stocks and bonds. They are more difficult to impose for *real* assets: no two houses are identical, and houses do not have live prices, so LTV restrictions can only be imposed based on estimates of the market value of these assets. Many approaches for conducting asset appraisals exist, based on either comparable sales as in the US housing market, or other methods such as cash flow or replacement-cost approaches.[6] But asset appraisal can never be perfect: appraisals are always noisy signals of asset values, so appraisal-based regulatory constraints will tend to be differentially binding, and may limit credit provision more, for assets which are harder to appraise. This effect qualitatively resembles the classic, and basically efficient, "collateral recovery" effect, where higher-uncertainty assets are worse collateral for debt even in frictionless markets. However, the "appraisal channel" is a different force driven by imperfect regulation, which has distinct effects on credit market outcomes and has no particular reason to be efficient.

Our paper fits into two main strands of literature. First, we relate to a number of papers analyzing how collateral values affect the properties of debt contracts collateralized by these assets (Titman and Wessels, 1988; Shleifer and Vishny, 1992; Kermani and Ma, 2020). Conceptually, the literature has argued that illiquid assets are inferior collateral. Several papers have tested this idea, and how collateral values affect financing and investment in

[5] For residential and commercial real estate, the Interagency Guidelines for Real Estate Lending Policies recommends specific LTV limits for bank loans backed by various types of real estate collateral.

[6] The USPAP discusses appraisal standards for real estate in the US. For international valuation standards, see the RICS Valuation Global Standards and the International Valuation Standards.

general, in a variety of empirical settings (Benmelech and Bergman, 2008; Benmelech, 2009; Benmelech and Bergman, 2009; Bian, 2021; Pan et al., 2024; Collier et al., 2021). Our contribution to this literature is to propose that the collateral-liquidity link reflects a novel "appraisal channel", in addition to the classic "collateral recovery" channel, and to show how these channels can be distinguished. To the best of our knowledge, this is the first paper in any setting to argue that appraisal-based regulatory constraints play a role in the collateral-credit relationship.

Second, we contribute to a literature on frictions that affect mortgage credit (Lustig and Van Nieuwerburgh, 2005; Mian and Sufi, 2011; Greenwald, 2016; Agarwal et al., 2017; Piskorski and Seru, 2018; Beraja et al., 2019; DeFusco et al., 2020; Adelino et al., 2020; Buchak et al., 2018; Jiang, 2023) and the corresponding real effects of such frictions (Glaeser and Shapiro, 2003; Di Maggio and Kermani, 2017; Agarwal et al., 2022; Di Maggio et al., 2017; DeFusco, 2018; Dokko et al., 2019; Gupta et al., 2024; DeFusco and Mondragon, 2020; Kermani and Wong, 2021). Relative to this literature, we are the first to study the effects of house *value uncertainty* on mortgage credit provision. We show that collateral value has modest effects on outcomes such as LTV and interest rates but sizable effects on mortgage rejection rates; our quantitative results imply that the effects of collateral values on consumers' expected utility from home purchasing can be nontrivial. The value uncertainty channel also has interesting distributional implications: low-income and minority households tend to live in areas with low collateral quality, implying that value uncertainty limits mortgage credit availability to some of the households who are most dependent on credit for homeownership.

Two older papers which discuss related effects of appraisal noise on mortgage credit are Lang and Nakamura (1993) and Blackburn and Vermilyea (2007). Lang and Nakamura (1993) propose a theoretical model in which, when sales volume is low, lenders are more uncertain about house values, so mortgage payoffs are lower and lenders demand higher down payments as a result. Though the paper frames this as an effect of appraisal noise, this force is in fact more similar to the collateral recovery channel in our model. The appraisal channel in our model is a distinct *regulatory* effect, which is not present in Lang and Nakamura (1993). Blackburn and Vermilyea (2007) empirically tests the relationship between market thickness and credit availability proposed in Lang and Nakamura (1993), by regressing mortgage approval rates on "market size" measures, such as sales volume. Relative to Blackburn and Vermilyea (2007), we use a more direct measure of collateral quality – value uncertainty – as well as being able to directly measure the effects of value uncertainty on the distribution of appraisals, allowing us to construct our quantitative model and to disentangle the two channels of effect.

The paper proceeds as follows. Section 2 describes our data, measurement strategy, and stylized facts on our price dispersion measure. Section 3 studies the effect of price dispersion on mortgage provision. Section 4 describes our model, and Section 5 calibrates the model to the data. We estimate model counterfactuals and discuss implications of our results in Section 6, and conclude in Section 7.

2 Measurement, Data, and Stylized Facts

2.1 Measuring Value Uncertainty

A number of recent papers have shown that house prices display nontrivial dispersion, likely driven by the fact that houses trade in thin markets (Case and Shiller, 1989; Giacoletti, 2021; Sagi, 2021; Hartman-Glaser and Mann, 2017; Sklarz and Miller, 2016; Buchak et al., 2020). As in Buchak et al. (2020), we estimate house price dispersion at the level of individual house sales, by measuring what kinds of houses have smaller errors when priced with a hedonic regression. We first regress transaction prices on house characteristics:

$$p_{it} = \eta_{kt} + f_k\left(x_i, t\right) + \epsilon_{it}, \tag{1}$$

We then regress the squared residuals, $\hat{\epsilon}_{it}^2$, from (1) on a flexible function of characteristics and time to predict which house characteristics make them difficult to price:

$$\hat{\epsilon}_{it}^2 = g_k\left(x_i, t\right) + \xi_{it} \tag{2}$$

In (1) and (2), i indexes properties, k indexes counties, and t indexes months. p_{it} is the log transaction price of house i at time t. $f_k\left(x_i, t\right)$ and $g_k\left(x_i, t\right)$ are generalized additive models in observable house characteristics x_i and time t, which we describe in Appendix A.1. $f_k(x_i, t)$ allows house characteristics to affect prices in a manner that varies over time. $g_k(x_i, t)$ allows the variance of price dispersion to vary with characteristics and over time. η_{kt} is a county-month fixed effect. Intuitively, specification (1) estimates a hedonic specification for house prices, and specification (2) projects the squared residuals ϵ_{it}^2 from the hedonic regression on house features and time, to predict which kinds of houses are difficult to value. We use the square roots of the predicted values from specification (2) as our house-level measure of idiosyncratic price dispersion:[7]

$$\hat{\sigma}_{it} \equiv \sqrt{\hat{g}_k(x_i, t)} \tag{3}$$

Our main specification uses a hedonic model of house prices; one concern is that there are house-level features which affect prices, which are observed by market participants but are not in our dataset. To alleviate this concern, in Appendix A.2, we repeat the analysis using a repeat-sales specification to predict prices in (1). This specification absorbs all time-invariant components of house quality, whether or not they correspond to observable characteristics in our data, into house fixed effects. The resultant price dispersion estimates are highly correlated with our baseline specification, and our empirical results continue to hold.

When considering whether to lend against a house, lenders should care about the *total* volatility of a house. Our measurement strategy focuses on the idiosyncratic component, which Piazzesi and Schneider (2016) estimate to be approximately half of total house price volatility. Most of our empirical analysis compares houses within region-years; these houses should have similar local index exposures, so most of the differences in total volatility which are relevant for our results should be driven by differences in idiosyncratic volatility.

[7]Note that it is important to use the predicted values of $\hat{\sigma}_{it}^2$ in stage 2 rather than the residuals $\hat{\epsilon}_{it}^2$ in stage 1 directly. This is because the expected value of idiosyncratic dispersion, σ_{it}^2, is the analog of σ in our model, which is relevant for the LTV. Each realization of $\hat{\epsilon}_{it}^2$ is a noisy measure of σ_{it}^2. If we regressed outcomes such as house-level LTP on the regression residuals $\hat{\epsilon}_{it}^2$ directly, the coefficients would be biased towards 0, relative to the first-best of regressing LTPs on σ_{it}, due to measurement error bias.

2.2 Data

Corelogic Deed & Tax Data: We obtain house transaction records in the entire US from 2000 to 2020 from the Corelogic Deed dataset. The dataset reports each house transaction attached to a specific property and provides information on the sale amount, mortgage amount, transaction date, and property location. We merge the transaction records with the Corelogic Tax records, which contain property characteristics such as year built and square footage. We estimate price dispersion for each house in this merged dataset. Appendix B provides detailed description about data cleaning steps. To link our price dispersion estimates with mortgage rates, loan rejections, and appraisals, we aggregate the price dispersion estimates to the 5-digit zipcode level because the data usage agreement prohibits us from linking the individual records across these datasets.

Corelogic Loan-Level Market Analytics (LLMA) Data: We obtain mortgage information from 2000 to 2020 from the Corelogic LLMA data, which provides detailed information on mortgage and borrower characteristics at origination – interest rates, down payments, sale prices, credit score, and debt-to-income ratio – and monthly loan performance after origination, including delinquency status and investor type. Importantly for our analysis, the LLMA provides both transaction price and the house's appraisal value. We use this dataset to estimate the menu of LTP-interest pairs in any given market, to examine loan performance, and to analyze appraisal values relative to prices.

Home Mortgage Disclosure Act (HMDA): The HMDA covers the near universe of U.S. mortgage applications from 2000 to 2017, including both originated and rejected applications. For all loan applications, we observe the application outcomes (whether a loan is approved or rejected) and the borrowers' locations. For rejected loans, we observe the rejection reasons, from which we determine whether a loan is rejected due to collateral-related reasons. We use the HMDA for extensive margin analysis on mortgage application rejections. We keep all completed home purchase mortgage applications with non-missing key variables, such as location, loan amount, income, and loan types.

Other Sources: We use the Booth TransUnion Consumer Credit Panel to calculate the average credit score by county to measure the creditworthiness of the entire borrower population. We obtain zipcode level demographic data from the American Community Survey

(ACS) 1-year and 5-year samples.

Table 1 provides summary statistics.

2.3 Stylized Facts

Price Dispersion is Persistent Over Time Figure 1 Panel (a) plots zipcode idiosyncratic price dispersion in 2020 against zipcode dispersion in 2010. While there is large cross-sectional variation in price dispersion, dispersion is very persistent over time. This suggests that differences in price dispersion are driven by persistent characteristics of the local housing stock, rather than time-varying factors such as local housing market conditions.

Price Dispersion and House Characteristics Table 2 presents the association between estimated value uncertainty and house characteristics. Panel A analyzes house features. Throughout, we control for linear and squared terms in log house prices, comparing houses with similar prices and different characteristics. Older houses have higher price dispersion (column 1). Controlling for building age, houses which were renovated within 5 years of the transaction date (column 2) have lower price dispersion.[8] Columns 3-4 present the association between property size, measured by square-footage and number of bedrooms, and price dispersion. There is a U-shaped relationship: price dispersion is low for moderately large houses and higher for houses which are very large or very small. In terms of local housing market conditions, Panel B of Table 2 shows that houses in zipcodes with larger income inequality, less population density, and more vacancies tend to have higher price dispersion. Together, Table 2 suggests that house price dispersion is essentially driven by house standardization and market thickness. This finding is consistent with evidence from other papers.[9] In Appendix C, we discuss a number of factors and theoretical forces that may drive dispersion, such as information asymmetry (Kurlat and Stroebel, 2015; Stroebel, 2016), search frictions, and so on.

[8] We can partially measure house renovations, as the Corelogic tax data contains an "effective year built" variable, which tracks the last date at which a property was renovated.

[9] See, for example, Kotova and Zhang (2021) and Andersen et al. (2022).

Price Dispersion and Appraisal Noise Figure 2 illustrates that high-dispersion houses tend to have noisier appraisal values. Panel (a) shows the percentage absolute difference between appraisals and transaction prices, $\frac{|a_i - p_i|}{p_i}$. This difference is around 1.5 percentage points larger in high-dispersion areas.

Appraisals only constrain borrowing when appraisals are below the sale price. In Panel (b), we define the "appraisal deviation" for each loan as $\frac{|a_i - p_i|}{p_i}\mathbf{1}(a_i < p_i)$, the product of the under-appraisal percentage $\frac{|a_i - p_i|}{p_i}$ and an indicator for under-appraisal, $\mathbf{1}(a_i < p_i)$. This is a summary measure of the downwards pressure on loan size induced by appraisals, combining the probability of under-appraisal and the size of under-appraisals. The figure shows that appraisal deviations are larger in high-dispersion areas.

The appraisal distribution is known to be very asymmetric: appraisals are often equal to sale prices, and under-appraisals are rare. This may reflect a combination of selection of successful sales, and human appraisers' incentives to bias appraisal prices upwards towards transaction prices. We will account for both effects in our model; however, to illustrate that our results are not driven only by appraiser bias, in Panel (c), we analyze the average size of the appraisal gap conditional on *over-appraisal*, $a_i > p_i$. Increasing the appraisal past the transaction price does not increase the amount borrowers can borrow, so appraisers should have no incentive to bias appraisals which are already above transaction prices.

Regional Variation: Zipcode Income and Race Price dispersion tends to be higher in low-income and black-dominant zipcodes. Panel (b) of Figure 1 shows the relationship between price dispersion and zipcode demographics. Comparing zipcodes with similar levels of median income, price dispersion in Black-dominant zipcodes tends to be 0.03 (1/4 SD) higher than in non-Black dominant zipcodes. Comparing zipcodes with similar racial composition, price dispersion in low-income zipcodes tends to be 0.06 (1/2 SD) higher than in high-income zipcodes.

3 Price Dispersion and Mortgage Credit Provision

We show a stylized depiction of how price dispersion can affect the home purchase and mortgage application process in Figure 3; we will formalize this framework in our model in Section 4. A homebuyer first decides to purchase a home. After the homebuyer's offer is accepted, she applies for a mortgage. The lender offers a menu of interest rate-LTV pairs to the borrower. Debt which is backed by higher-dispersion houses has lower expected recovery rates, because the value of debt is concave in the foreclosure price of collateral: if the collateral sells for more than the outstanding debt amount, lenders do not keep the upside, whereas lenders are responsible for some of the downside if the collateral sells for less. Thus, when houses have higher price dispersion, lenders should offer worse rate menus: higher interest rates for any given LTV, and vice versa. The buyer then chooses an option from the menu: in high-dispersion areas, buyers are thus forced to choose either lower LTVs, higher interest rates, or both. We call this effect of price dispersion the "collateral recovery" channel; this is a classic effect documented in a number of other collateralized debt markets (Benmelech and Bergman, 2008; Benmelech, 2009; Benmelech and Bergman, 2009; Bian, 2021; Pan et al., 2024; Collier et al., 2021).

After the buyer chooses her targeted mortgage LTV, the house is appraised, and the value of the house for LTV calculation is set to the lesser of the transaction price and the appraised value. If the house over-appraises – that is, the appraisal price is at least the transaction price – the transaction proceeds as planned. If the house under-appraises, the homebuyer may need to decrease her mortgage size to meet regulatory LTV constraints by making higher down payments.[10] If the buyer is unable to make the higher down payments, she may have to renege on the transaction. We showed in Figure 2 that appraisals are noisier when house price dispersion is high, resulting in higher under-appraisal in high-dispersion areas. Thus, mortgages in high-dispersion areas are more likely to be rejected, and buyers may decrease their targeted LTVs to lower the impact of under-appraisal risk. We refer to this effect as the "appraisal risk" channel.

In summary, Figure 3 illustrates that, when house price dispersion is high, mortgage

[10] For loans held on banks' balance sheets, the amount of capital required by the regulators typically depends on the LTV ratios (See, for example, the BIS document on bank capital requirements). For securitized loans, there are LTV constraints on whether they can receive government guarantees, and also the guarantee fees usually depend on the LTV ratios (See, for example, Fannie Mae loan eligibility requirements; and loan-level price adjustment matrix).

applications should have a higher mortgage rejection likelihood, receive higher interest rates, and loan-to-price ratios should be lower. We now show that these relationships hold in a variety of empirical specifications.

3.1 Effects on Mortgage Credit

3.1.1 Mortgage Rejection Rates

Figure 4(a) plots county-level mortgage rejection rate against price dispersion, in which rejection rates are calculated as the number of rejected mortgage applications divided by total completed mortgage applications in a county-year recorded in the HMDA data. The figure shows that mortgage applications are more likely to be rejected in counties with higher price dispersion. The HMDA data also records lender-reported reason for rejecting an application. In Figure 4(b), we use this information to construct county-level rates of rejection due to collateral-related reasons, calculated as the number of rejected mortgage applications that are reported as due to collateral reasons, divided by total applications. The figure shows that more applications are rejected for collateral-related reasons in high-dispersion counties.

We then exploit within county-year variation by estimating the following loan application-level specification:

$$Reject_{ikt} = \beta ZipDispersion_{ikt} + X_{ikt}\Gamma + \mu_{kt} + \nu_{lt} + \epsilon_{ikt} \tag{4}$$

$Reject_{ikt}$ is an indicator that equals 100 if the mortgage collateralized by property i in county k in year t is rejected and 0 otherwise. $ZipDispersion_{ikt}$ is the average price dispersion of houses in property i's zipcode that are transacted in year t.[11] X_{ikt} is a set of controls, including zipcode house transaction price, credit score and its squared term, individual income, loan-to-income ratio and its squared term, and mortgage type.[12] μ_{kt} and ν_{lt} are county-year and lender-year fixed effects, respectively.

Panel A of Table 3 reports the results. We first confirm the effect of local house price

[11] We aggregate property-level price dispersion measures estimated using Corelogic Deeds to zipcode level and assign it to every loan application in HMDA based on borrowers' location.

[12] We aggregate mortgage-level interest rates from Corelogic LLMA to zipcode level and assign it to every loan application in HMDA based on borrowers' location.

dispersion on mortgage rejection rates using the full sample (column 1). Zipcode house price dispersion is positively and significantly associated with mortgage rejections. This result holds for both securitized loans (column 2) and portfolio loans (column 3). The rejection rate increases by about 1.4 percentage-points for every standard deviation increase in house price dispersion. The effect is economically significant: given the sample average rejection rate of about 16 percentage points, the estimate amounts to about a 10% increase in rejection likelihood.

In Panel B of Table 3, we focus on collateral-related rejections. A mortgage application is about 50bps more likely to be rejected due to collateral reasons in a zipcode with one standard deviation higher house price dispersion, which is about a 25% increase in the likelihood of collateral-related rejections. Again, the result holds in the full sample (column 1) as well as sub-samples of securitized loans (column 2) and portfolio loans (column 3).

Mortgage Rejection Reasons As a robustness check, we examine the relationship between house price dispersion and different rejection reasons among rejected loans. We restrict the sample to only rejected loans and estimate Specification 4 using various rejection reason indicators as the outcome variables. Intuitively, this specification estimates, conditional on a mortgage being rejected, whether rejections are more likely to be attributed to collateral-related reasons in high-dispersion areas.

Table A1 reports the results. As the sample means indicate, the most common rejection reasons in the entire sample are creditworthiness-related reasons (i.e., credit score and debt-to-income ratios). However, as house price dispersion increases, the results show that the mortgage rejections are significantly more likely due to collateral reasons, and less likely to be due to creditworthiness reasons, thereby supporting our baseline findings.

3.1.2 Interest Rates

In high price dispersion areas, lenders offer worse rate menus, so mortgage interest rates are higher for any given LTP ratio. To visually demonstrate this, we estimate the entire menu of LTP-interest rate pairs available to borrowers in high- and low-dispersion areas. We first residualize interest rates using borrowers' credit scores, loan type, and time fixed effects; we then plot the residuals against LTP separately for zipcodes with above-median and below-median dispersion in Figure 5. The figure shows that the entire menu of interest rate-LTP pairs shifts upwards in high-dispersion zipcodes: for any given LTP, borrowers in high-dispersion zipcodes can expect to pay higher prices. The difference is about 3bps for loans with LTP below 80, and enlarges to 7bps for loans with LTP above 80.

We then estimate the following loan-level specification:

$$Rate_{ikt} = \beta_1 ZipDispersion_{ikt} + \beta_2 LTP_{ikt} + X_{ikt}\Gamma + \mu_{kt} + \nu_m + \epsilon_{ikt} \tag{5}$$

$Rate_{ikt}$ is the interest rate on mortgage i collateralized by house in county k in year t. $ZipDispersion_{ikt}$ is the average price dispersion of houses in property i's zipcode that are transacted in year t.[13] X_{ikt} is a set of controls, including house transaction price, credit score and its squared term, LTV and its squared term, debt-to-income ratio and its squared term, and loan type. μ_{kt} and ν_m are county-year fixed effect and loan origination month fixed effect, respectively.

Table 4 presents the results. Column 1 uses the full sample. Higher loan-to-price ratios are associated with higher interest rates: a one percentage point increase in LTP is associated with an 69bps increase in interest rate. Controlling for LTP, houses in zipcodes with higher house price dispersion have higher interest rates. The mortgage rate increases by 0.89bps in zipcodes with one standard deviation higher average house price dispersion. Columns 2 to 3 show the results for securitized loans and portfolio loans; the results hold in both samples. For every 1SD increase in zipcode average house price dispersion, the mortgage rate of securitized loans increases by 1.31bps, and the rates on portfolio loans increases by 1.26bps.

3.1.3 Loan-to-Price Ratio

Lastly, we show that price dispersion is associated with smaller loan sizes, as measured by loan-to-price ratios (LTP). Figure 4(c) illustrates the relationship by plotting county average

[13] We use zipcode dispersion instead of property-level dispersion because our price dispersion measure is estimated using Corelogic Deeds, and the data vendor prohibited us from merging loan-level records in LLMA with property-level records in Corelogic Deeds and Tax. We therefore aggregate property-level price dispersion measures to the most granular geographic region in LLMA.

LTP against average house price dispersion.[14] Counties with higher price dispersion have lower average LTPs. The pattern holds for all types of loans: GSE loans, FHA loans, and jumbo loans (Figure A2).

We then exploit within county-year variation by comparing two properties that are bought in the same county-year at the same price and by buyers with similar credit profiles and incomes. To implement this strategy, we estimate the following property-level specification:

$$LTP_{ikt} = \beta Dispersion_{ikt} + X_{ikt}\Gamma + \mu_{kt} + \nu_d + \epsilon_{ikt} \tag{6}$$

LTP_{ikt} is the loan-to-price ratio of a mortgage, collateralized by property i in county k in year t. $Dispersion_{ikt}$ is the estimated price dispersion of the underlying property. X_{ikt} is a set of controls, including property transaction price, mortgage type, mortgage term, and resale indicator.[15] μ_{kt} and ν_d are county-year and transaction date fixed effects, respectively.

Table 5 presents the results. Column 2 corresponds to Specification 6. Column 1 includes only transaction date fixed effects, and column 3 adds lender-year fixed effects. For two houses in the same county that are transacted on the same date at the same price, the one with higher estimated price dispersion tends to receive a smaller sized loan. In the most saturated specification, the loan-to-price ratio is more than 20bps lower for houses with one standard deviation higher estimated price dispersion across these specifications.

3.2 Identification

Mortgage outcomes could be worse for high-dispersion houses, not because the *houses* are worse, but because *borrowers* purchasing these houses are systematically less creditworthy. Lenders would then rationally lend less against these houses; however, the channel of effect would be due to borrower quality rather than house quality. One indication that this does not explain our results is that, in Subsection 3.1.1, we showed that lenders in high-dispersion areas systematically indicate that they are rejecting more mortgages because the collateral has low quality. To further alleviate this concern, we develop an instrument for house price

[14] To make this plot, we first remove the average LTP differences across levels of individual house prices and then plot their county average value against county average price dispersion.

[15] Our results are not sensitive to the inclusion of transaction price, though we believe including price is the right specification. We discuss this in detail in Appendix D.

dispersion, show that the instrument is not correlated with ex-ante or ex-post buyer credit characteristics, and show that all our baseline results hold using the instrument.

3.2.1 Instrumental Variables

Our estimates of house price dispersion are functions of house characteristics. House characteristics could be confounded with buyer characteristics because, for example, less creditworthy buyers could tend to buy smaller houses in a region. However, it is plausible that the relationship between buyer creditworthiness and house characteristics is mostly monotone: it would intuitively be less likely that both buyers of very small and very large houses in a region are less creditworthy than buyers of average-sized houses in that region. On the other hand, the literature on house price dispersion has shown that the relationship between house characteristics and price dispersion has a strong component that is quadratic around mean house characteristics: houses that are either very large or very small tend to have higher dispersion than average-sized houses. Thus, if we construct a measure of how "unusual" a house's characteristics are relative to other nearby houses, the measure may be associated with price dispersion, but less related to buyer creditworthiness.[16]

Formally, we construct a set of house-level instruments for the price dispersion of each individual house i in county c by measuring its heterogeneity relative to the local housing stock. For all home purchases transacted in each county c in a given year, we first calculate the average value of each key house features ($\overline{X}_c^m$), where

$$m \in \{\text{building age}(age), \text{size}(sqft), \text{bedrooms}(bed), \text{bathrooms}(bath), \text{geo-coordinates}(geo)\}.$$

For each house, we calculate the "distance" between its feature m and the average value of m in county c in the same transaction year, denoted by Z_i^m:

$$Z_i^m = (X_i^m - \overline{X}_c^m)^2, \;\; \forall m \in \{age, sqft, bed, bath, geo\}, \tag{7}$$

House features are not symmetrically distributed around their mean in the data, meaning

[16] The approach of using measures of house nonstandardization as instruments is not new to the literature: similar ideas are used in Andersen et al. (2022), and the approach can be micro-founded in a search and matching framework as done in Guren (2018).

that the distances constructed in (7) may still be increasing or decreasing in house features. To further alleviate this effect, we purge Z_i^m of the variation that is linear in house features by taking the residuals of the following linear regression:

$$Z_i^m = \sum_k \alpha_k X_i^k + \mu_d + \epsilon_i^m, \tag{8}$$

where X_i^k are the five house features and μ_d is house transaction date fixed effects. Our instruments are the residuals in the above equation:

$$\tilde{Z}_i^m = \hat{\epsilon_i^m} \tag{9}$$

The instruments measure how locally thin the market is for a given house i, by benchmarking it to other houses within the same county. Small houses, for example, will have large Z^{sqft} in a county with mostly large houses, but will have small Z^{sqft} in a county with mostly small houses. Since markets for small houses are thinner in the former than in the latter, there are likely fewer buyers at any given point in time in the former.

Using the five instruments, we estimate the following 2SLS specification:

$$\begin{aligned} \text{Stage 1: } Dispersion_{it} &= \beta_1 \tilde{Z}_{it}^{age} + \beta_2 \tilde{Z}_{it}^{sqft} + \beta_3 \tilde{Z}_{it}^{bed} + \beta_4 \tilde{Z}_{it}^{bath} + \beta_5 \tilde{Z}_{it}^{geo} \\ &\quad + Controls_{ikt}\Gamma_1 + \mu_{kt} + \nu_d + \epsilon_{ikt} \\ \text{Stage 2: } Y_{ikt} &= \gamma \widehat{Dispersion}_{ikt} + Controls_{ikt}\Gamma_2 + \eta_{kt} + \zeta_d + \xi_{ikt}, \end{aligned} \tag{10}$$

where $\widehat{Dispersion}_{ikt}$ is the predicted value from stage 1, X_{ikt} are control variables, and μ_{kt}, ν_d (η_{kt} and ζ_d) are county-year and transaction date fixed effects, respectively.

For our analyses that use zipcodes, we take geographical averages of $\tilde{Z}_i^m$. The aggregated instrument is essentially a measure of the heterogeneity of the zipcode's housing stock along characteristic m. Zipcodes with more heterogeneous housing stocks tend to have higher price dispersion, since they have thinner local markets for any individual house.

In order for the instruments to be valid, we must argue that they are relevant in the sense that they are associated with house price dispersion and excluded in the sense that they are not associated with other factors which may affect mortgage market outcomes, in particular, buyers' observable and unobservable creditworthiness. For relevance, consistent with the literature, Table A2 shows that our instruments are correlated with the raw price dispersion measure in a statistically significant manner.

For exclusion, we show that the instrument is uncorrelated with house prices and ex-ante measures of buyer creditworthiness and also that borrowers in areas with higher values of the instrument are not ex-post more likely to default. We first conduct a balance test on ex-ante characteristics in Table A3 by correlating raw or instrumented zipcode price dispersion with house price, borrower FICO, and borrower income in that zipcode. Our results suggest that the raw price dispersion measure is negatively correlated with house prices, FICO, and income, but the instrumented price dispersion does not appear to be statistically significantly correlated with these characteristics. If anything, borrowers in zipcodes with higher instrumented dispersion tend to have higher, instead of lower, income.

While these results suggest that the instrument is not associated with ex-ante *observable* measures of borrowers' creditworthiness, a further concern is that borrowers in heterogeneous zipcodes may be *unobservably* worse credit risks. To address this concern, we can further do an *ex post* test, measuring whether borrowers in more heterogeneous zipcodes are more likely to default on mortgages. Table 6 Panel A estimates specifications 5 (columns 1-3) and the 2SLS version of it (columns 4-6), but sets the outcome variable equal to 100 for loans that become 60 or more day-delinquent within 2 years after origination and zero otherwise. Columns 1 and 4 include the full sample. Columns 2 and 5 restrict the sample to securitized loans. Columns 3 and 6 restrict the sample to portfolio loans. All regressions include the full set of borrower and loan characteristics as in our main regression specifications.

In these specifications, interest rates, FICO scores, DTI, and LTV are generally associated with ex-post default rates as expected. Consistent with the results from Table A3, in the OLS results, price dispersion is positively associated with default rates: buyers in high-dispersion zipcodes are more likely to default, even after controlling for observable mortgage and buyer features. However, instrumented price dispersion is not associated with default rates: buyers in high heterogeneity zipcodes are not more likely to default on their mortgages. This lends support for our exclusion restriction, that zipcode heterogeneity shifts house price dispersion without shifting buyer creditworthiness.

3.2.2 IV Results

We estimate specification 10 for every credit outcome in our baseline analyses in Section 3.1. We confirm our baseline results qualitatively and get reasonably stronger estimated effects.

Approvals Table 3 reports the loan approval likelihood results. Zipcode house price dispersion is positively and significantly associated with mortgage rejections: the rejection rate increases by more than 2 percentage points as house price dispersion increases by one standard deviation (Panel A columns 4-6). As shown in Panel B, a mortgage application is about 80bps more likely to be rejected due to collateral reasons in a zipcode with one standard deviation higher house price dispersion (Panel B column 4-6). Both results — overall rejection or rejection due to collateral reasons — hold in the full sample as well as sub-samples of securitized loans and portfolio loans.

Interest Rates Table 4 columns 4-6 present the interest rate results. For every one standard deviation increase in zipcode average house price dispersion, the mortgage rate increases by 1.4bps in the full sample (column 4), increases by 1.7bps in the sample of securtized loans (column 5), and increases by 3bps for portfolio loans (column 6).

LTP Lastly, Table 5 columns 4-6 present the LTP results, where column 5 corresponds to Specification 10, and column 4 and 6 are less and more saturated specifications, respectively. LTP decreases by 1.3 percentage points for every one standard deviation increase in the estimated price dispersion in the most saturated IV specification.

The IV coefficient estimates are mostly larger than the OLS estimates. This is potentially driven by the fact that the independent variable, price dispersion, is measured imperfectly by our first-stage regression, causing the coefficients in the OLS specifications to be biased toward 0. When instrumental variables alleviate measurement error in the independent variable, they tend to lead to larger coefficient estimates; the pattern that IV estimates tend to be larger than OLS estimates is common in empirical studies across many areas (Pancost and Schaller, 2021).

Economic Magnitudes The largest effects of price dispersion are on mortgage rejection rates: according to our OLS estimates, a 1SD increase in house price dispersion is associated with a 25% increase in the prevalence of collateral-related mortgage rejections, and a 10% increase in total rejection rates.[17] The effects on loan sizes and interest rates are statistically significant but economically smaller. However, as we illustrate in Appendix E, a seemingly small percentage change in the down payment requirement can in principle have moderate-size effects on house affordability, considering the low level of annual savings by the marginal home buyers in the US.

3.3 Distributional Consequences

The relationship between price dispersion and mortgage credit has important distributional consequences: price dispersion tends to be higher in low-income zipcodes, so price dispersion limits mortgage credit disproportionately to some of the households most reliant on credit for homeownership.

To illustrate this point quantitatively, for each transacted house in our dataset, we calculate how much rejection rate, down payment, and mortgage rate would decrease according to our reduced-form estimates, if the house's price dispersion decreased to be equal to the bottom decile of price dispersion in our dataset. Specifically, we calculate the difference between a house's price dispersion and the bottom decile of price dispersion, and then multiply it respectively by the estimated effect of price dispersion on mortgage rejection (column 4 of Table 3), by the estimated effect on LTP (column 6 of Table 5), and by the estimated effect on mortgage rate (column 4 of Table 4). We then analyze how these effects vary across zipcodes with different income levels in Figure 6 by plotting the average changes in mortgage rejection rate, LTP, and mortgage rates in each zipcode income quintile.

We find that, if the price dispersion of local housing stocks decreased to be equal to the bottom decile of price dispersion in our dataset, mortgage failure rates would decline by about 6 percentage points in the bottom quintile of zipcodes by income; this would be about 1.5 times the magnitude of the impact in the top zipcode income quintile. While the impact on loan sizes and interest rates would be economically smaller, there would still be

[17] The IV estimates are slightly larger: a 1SD increase in house price dispersion is associated with a 40% increase in the prevalence of collateral-related mortgage rejections, and a 14.5% increase in total rejection rates.

meaningful heterogeneity across zipcodes with different income levels: the increase in loan size in the bottom zipcode income quintile would be 1.5 times the magnitude of the impact in the top zipcode income quintile, while the decline in mortgage rates in the bottom zipcode income quintile would be 60% higher than the magnitude of the impact in the top zipcode income quintile.

3.4 Robustness Checks

We conduct two additional robustness checks. In Appendix F.1, we show that our results hold in a subsample of transactions with sale prices below conforming loan limits, suggesting that our results are not driven by homebuyers' incentives to keep prices below conforming loan limits. In Appendix F.2, we address the role of lender market power in three steps. First, we show that our main findings survive controls for lender-county-year fixed effects fixed effects when data permits, suggesting that the results are not driven by market power at the lender-county level. However, a remaining concern is that only a small number of lenders may be willing to lend against houses with high value uncertainty, giving these lenders market power to extract rents from home buyers of such houses, but not from other home buyers purchasing houses with lower value uncertainty.

To address this concern, we investigate whether the number of lenders serving high-value-uncertainty market segments is indeed smaller than the number of lenders serving low-value uncertainty market segments *within* the same county-year. In Table A5, we divide each county-year into four market segments based on the price dispersion of transacted houses recorded in the Corelogic Deed data, and examine the correlation between price dispersion and lender concentration within a county-year. The results indicate a lack of significant positive correlation between lender concentration and house price dispersion within a county-year. In fact, in some specifications, lender concentration appears to be negatively correlated with house price dispersion: we observe a higher number of lenders lending against houses with higher price dispersion, yielding lower lender concentration in the higher price dispersion segments within a county-year. Therefore, although suggestive, the evidence is not consistent with the narrative that a lender can earn monopolist rents on buyers of high-dispersion houses but not on buyers of low-dispersion houses in the same county-year.

Finally, to further isolate the effect of collateral value uncertainty from the effect of lender market power, we conduct subsample analyses for houses located in zipcodes with low lender concentration level versus houses located in zipcodes with high lender concentration level. The results suggest that the effects of price dispersion on loan rejection, loan sizes, and interest rates remain significant and have similar magnitudes after controlling for mortgage supply concentration and the effects are quantitatively similar across zipcodes with different lender concentration levels. The results suggest that the main findings of this paper — the effects of price dispersion on mortgage provision — are not driven by the correlation between price dispersion and lender market power within a county-year.

4 Model

We build a structural model showing how price dispersion affects application failures, interest rates, and mortgage loan-to-value ratios (LTVs) through the collateral recovery and appraisal risk channels. The model follows the structure in Figure 3. A prospective homebuyer chooses a targeted mortgage size to finance a house at an exogenous transaction price. By choosing a larger mortgage, the buyer smoothens consumption more effectively, but also faces higher interest rates and a greater risk of under-appraisal and mortgage rejection. When idiosyncratic price dispersion is higher, lenders offer borrowers worse interest rate menus, and under-appraisals are more likely; both forces push buyers towards choosing smaller mortgages.

4.1 Setup

4.1.1 The Buyer's Problem

A homebuyer attempts to finance a house that is sold at price P, by choosing a target loan size L. The choice of L determines the buyer's consumption in two time periods: the first period is when the buyer purchases the house, and the second is when the mortgage loan is paid back. The buyer has CRRA utility, discounting consumption at rate β^T between periods:

$$U(c_1, c_2) = \frac{c_1^{1-\eta} - 1}{1-\eta} + \beta^T u_2' c_2 \tag{11}$$

where u_2' is an exogenous constant. Hence, the buyer solves a consumption smoothing problem, where utility is concave in the first period, and linear in the second. The buyer receives exogenous labor income W_1 in period 1, and W_2 in period 2. Expression (11) can be thought of as a reduced-form of a richer model in which a consumer smoothens consumption between a single period, in which the house purchase is made, and a large number of future periods in which the mortgage is paid down. The term $\beta^T u_2' c_2$ can be thought of as a linear approximation to the consumer's value function over wealth in future periods after the house purchase.[18]

The mortgage application process has two stages:

1. Lenders offer an interest rate menu $r(L, \sigma)$, determining the mortgage interest rate if the buyer targets loan size L and idiosyncratic price dispersion is σ. The buyer chooses a target loan size L, receiving interest rate $r(L, \sigma)$. We introduce how rate menu is determined in the next section.

2. The house appraisal value A is determined. The collateral value used to calculate the LTV of the mortgage takes the smaller of the appraisal value A and the transaction price P:[19]

$$L_{final} \leq \phi \min(P, A). \tag{12}$$

   If $A < P$, the final loan amount L_{final} will be below the target size L, so the buyer will need to make an additional down payment. Conditional on A, the buyer can choose to continue the transaction, or to renege, pay a fixed penalty cost, and search for a new house, returning to period 1.

[18] A similar linear approximation to utility in future periods is used in Jansen et al. (2022). The concavity of utility over consumption in the single house purchase period is high relative to the concavity over the value function of wealth in future periods, since there are many future periods to smooth consumption over, so assuming post-purchase utility is linear in consumption is likely a reasonable approximation. In our setting, this modeling simplification is needed in order to make the appraisal problem recursive, allowing us to use tools from the search literature to model the buyer's response to under-appraisals.

[19] This constraint is imposed by both bank regulators and mortgage securitizers in reality.

In the following, we normalize final loan size, target loan size, and appraisal values:

$$l_{final} = \frac{L_{final}}{P}, l \equiv \frac{L}{P}, a \equiv \frac{A}{P} \tag{13}$$

Hence, the target LTV is l, the final LTV is l_{final}, and the ratio of appraisals to transaction prices is a. We will write $r(l)$ to mean the interest rate if the target LTV is l. We proceed to describe the buyer's payoffs if she chooses to continue with a transaction, then if she decides to renege.

Continuation From (12), if $a < 1$, the final loan size is capped at:

$$\phi \min(P, A) = \phi P \min(1, a) = \phi a P \tag{14}$$

Since we have restricted the target loan size to $l < \phi P$, the buyer's final loan size is $P \min(l, \phi a)$. If the buyer originally planned to borrow l, making down payment $P(1-l)$, the appraisal constrains loan size further whenever $a < \frac{l}{\phi}$. With appraisal a, the required down payment is $P(1-\phi a)$, which is $P \max[0, l - \phi a]$ larger than the targeted down payment. We assume that, if the buyer faces such a down payment gap, this decreases her period-1 consumption c_1 by $\psi P \max[0, l - \phi a]$, where $\psi > 1$. That is, for every dollar in additional downpayments she must make, the buyer's period-1 consumption decreases by $\psi > 1$ dollars. This is a reduced-form modelling device, capturing the idea that an unanticipated increase in down payments, induced by an under-appraisals, is more costly than an anticipated increase, because it is harder to smooth consumption in response to unanticipated shocks; we demonstrate this point quantitatively in Appendix G.1.

Given an appraisal a, the buyer's consumption in period 1 is:

$$c_1 = \underbrace{W_1}_{\text{labor income}} - \underbrace{P(1-l)}_{\text{target down payment}} - \underbrace{\psi P \max[0, l - \phi a]}_{\text{penalty term from under-appraisal}} \tag{15}$$

That is labor income less the target down payment for the house, less the penalty term from under-appraisal. Consumption in period 2 is:

$$c_2 = \underbrace{W_2}_{\text{labor income}} - \underbrace{\left(1 + r\left(l\right)\right)^T P\left(l - \max\left[0, l - \phi a\right]\right)}_{\text{mortgage principal and interest}} \tag{16}$$

This is labor income, minus the principal and interest on the mortgage, which we assume is paid in a single lump sum in period 2. Since utility in period 2 is linear, the term W_2 simply increases the level of utility and does not affect any outcomes, so for notational simplicity we will set $W_2 = 0$ going forwards.

Reneging If the appraisal is too low, the buyer can renege on the transaction, paying a cost ζ (as a fraction of house price), and then searching for a new house. For tractability, to make the problem recursive, we think of ζ as being paid in period 2 dollars. We think of this as capturing, for example, foregone deposits if there is no appraisal contingency in the sales contract or hassle costs of searching for another house. They then revert to stage 1, to purchase another house, and have continuation value:

$$-\beta^T u_2' \zeta P + E_a\left[V\left(a, l\right)\right] \tag{17}$$

where $V\left(a, l\right)$ is the value of choosing loan size l, when the appraisal is a.

4.1.2 Interest Rate Menus

We assume that the interest rate lenders offer depends on price dispersion and the size of the mortgage. Mortgages which are larger, and which are in higher-dispersion areas, are riskier, and lenders will thus charge higher interest rates as a result. In the main text, we assume a simple reduced-form model of the rate menu:

$$r\left(l, \sigma\right) = \bar{r} + \theta_l l + \theta_\sigma \sigma \tag{18}$$

where θ_l and θ_σ capture the dependence of the interest rate on loan size and price dispersion respectively.[20]

[20] We adopt this reduced-form model of the rate menu in the model for simplicity; however, in Appendix G.2, we construct a more detailed microfoundation of the interest rate menu. We assume competitive profit-maximizing

4.1.3 The Distribution of Appraisal Values

It is known in the literature that house appraisals are systematically biased upwards, and there is substantial bunching at house transaction prices. In our data, we find that large over-appraisals are also rare, suggesting that appraisers largely only bias appraisals upwards to the point where they are equal to sale prices. We model appraisals in a way that matches these stylized facts. We assume that there is an unbiased appraisal value which is normally distributed around the house transaction price, $A_{raw} \sim N(P, \sigma)$. The appraisal value A given to the borrower is then determined by:

$$A = \begin{cases} A_{raw} + Pb & A_{raw} < P(1-b) \\ P & P(1-b) \leq A_{raw} < P \\ A_{raw} & P \leq A_{raw} \end{cases} \tag{19}$$

Expression (19) states that, when A_{raw} is above P, appraisers simply report the raw appraisal price $A = A_{raw}$. When A_{raw} is below P but above $P(1-b)$, the appraisers biases A just enough so that it is equal to P, generating bunching at P. When A_{raw} is below $P(1-b)$, appraisers still attempt to bias A upwards, but are only able to push the appraisal to $A_{raw} + Pb$. This is still useful to the buyer, since any upwards bias in appraisals still allows the buyer to borrow more. We will estimate b based on the distribution of appraisal-to-sale ratios in our data, as we describe in Subsection 5.1 below.[21]

lenders make loans, setting prices and LTPs such that they at least break even, given imperfect collateral recovery rates. When price dispersion is higher, lenders must offer a worse rate menu to break even. In a simple calibration of the model, the observed dependence of the interest rate menu on price dispersion in the data can be matched fairly well, under reasonable assumptions for average foreclosure discounts. Since the dependence of the rate menu on price dispersion is quantitatively consistent with this microfoundation in the data, we proceed with the reduced-form model (18), as a simpler linear approximation to the microfounded model.

[21] In Appendix G.3, we show that (19) can be microfounded in a simple model based on Calem et al. (2021). In the model, appraisers have a convex cost of biasing appraisals upwards, and receive some linear side benefit – for example, from increased future business – to the extent that they are able to increase the amount that buyers can borrow on the loan. In this model, appraisals bunch at sale prices, because appraisers face positive costs, but no benefit, of biasing appraisals upwards past the transaction price, since the transaction price then binds in (12), and further increases in A do not affect the amount that can be borrowed.

4.2 Model Outcomes

Optimal behavior in the model is described by buyers' optimal target loan size choice l and buyers' optimal decision about whether to continue or renege on the transaction for each possible value of a. The following theorem characterizes optimal buyer behavior.

Theorem 1. *For any parameter settings, and for any target loan size l, there is an optimal appraisal cutoff $\bar{a}(l)$, which is the unique value that satisfies:*

$$\omega(\bar{a}, l) = -\beta^T u_2' \zeta P + \int_0^\infty \max\big(\omega(a, l), \omega(\bar{a}, l)\big)\, dF_a(a) \tag{20}$$

where $\omega(a, l)$ is defined as:

$$\omega(a, l) \equiv u_1\big(W_1 - P(1-l) - \psi P \max(0, l - \phi a)\big) + u_2' \beta^T \big(1 + r(l)\big)^T P \max[0, l - \phi a] \tag{21}$$

The buyer optimally continues with the purchase for any $a > \bar{a}(l)$, and reneges on the transaction for any $a < \bar{a}(l)$. The buyer chooses target loan size l to solve:

$$l^* = \arg\max_l \left(-\beta^T \big(1 + r(l)\big)^T u_2' P l + \int_0^\infty \max\Big(\omega(a, l), \omega\big(\bar{a}(l), l\big)\Big)\, dF_a(a)\right) \tag{22}$$

The proof of Theorem 1, and further properties of the buyer's choice problem, are described in Appendix G.4. In words, Theorem 1 states the following. Conditional on any target loan size l, buyers will continue the transaction if the house appraises to at least $\bar{a}(l)$, and will renege otherwise. The cutoff $\bar{a}(l)$ is the value of the appraisal such that the consumer is just indifferent between continuing with the transaction and making a higher down payment, thus receiving the LHS of (20); and reneging, thus receiving the RHS of (20), which is negative the cost ζ multiplied by house prices and period-2 marginal utility, plus the expected value from buying a new house.

To find the optimal loan size target, (22) states that buyers simply maximize expected utility from the second-stage problem over l. In Appendix G.5, we derive a first-order condition for optimal loan choice. The buyer faces a tradeoff: larger loan sizes smooth consumption more effectively if the house over-appraises, but lead to higher interest rates, and also larger under-appraisals and thus larger consumption penalties in period 1 upon under-appraisal.

Buyers thus optimally choose a target loan size slightly smaller than they would if the house never under-appraised, limiting consumption smoothing in order to decrease interest rates and under-appraisal risk.

5 How Does Price Dispersion Affect Mortgage Outcomes?

5.1 Calibration

We calibrate several parameters externally based on existing literature. We set the intertemporal elasticity of substitution (η) to 2, as chosen in standard lifecycle models. We set period 1 wealth to \$60,000 and the house price to \$200,000. We set $\beta = 0.96$. We set $T = 7$, approximately equal to the duration of a 30-year mortgage.[22] The maximum LTV parameter ϕ is set to 0.8, which is the most common regulatory threshold. We evaluate the sensitivity of our estimates to varying some of these input parameters in Appendix H.1: while estimated parameter values are somewhat sensitive to input choices, estimated counterfactual quantities are relatively insensitive to using different inputs.

We then estimate the remaining parameters by matching model-implied moments to data moments in the data, through a multi-step procedure.

Appraisal Distribution In our model, raw appraisal values a_{raw} are distorted only when they are below the transaction price, $a_{raw} < 1$, since appraisers have no incentives to further bias appraisals that are above the transaction price. Thus, the distribution of realized appraisals, conditional on over-appraisal, should be identical to the distribution of a_{raw}. Since we also assume raw appraisals have mean equal to the house price, we can thus estimate σ_i^A as:

$$\hat{\sigma}_{a,i} = \sqrt{E\left[(a_i - 1)^2 \mid a_i > 1\right]} \tag{23}$$

[22] Mortgages amortize and are prepayable, so their average duration is much lower than 30 years; see for example Krishnamurthy and Vissing-Jorgensen (2011).

That is, $\hat{\sigma}_{a,i}$ is simply the square root of the conditional mean squared error of appraisals around 1. Using expression (23), we calculate $\hat{\sigma}_{a,i}$ for each decile of σ values.

Interest Rate Menu To calibrate the interest rate menu, $r(l, \sigma)$ from expression (18), we assume:

$$r(l, \sigma) = \bar{r} + \theta_l (l - 0.8) + \theta_\sigma (\sigma - \bar{\sigma}) \tag{24}$$

That is, the interest rate $r(l)$ is equal to a constant $\bar{r}$, plus θ_l times the target LTV, plus θ_σ times idiosyncratic price dispersion. We set $\bar{r}$, the interest rate for a mortgage with $l = 0.8$, and $\sigma = \bar{\sigma}$, to $\frac{1}{\beta} - 1$, which is approximately 4.17%. We set θ_l and θ_σ to their values in Column 6 of Table 4.

Moment Matching There are four important parameters in the model that govern the homebuyer's tradeoff: the under-appraisal penalty (ψ), the cost of transaction failure (ζ), the appraisal bias (b), and period-2 marginal utilty (u_2'). We estimate these parameters by matching three sets of data moments: transaction failure probabilities; "appraisal deviations", measuring the prevalence and magnitude of underappraisals which do not result in failed transactions; and the relationship between final loan-to-price ratios and price dispersion.

We compute mortgage failure probabilities within each σ-decile of counties as the rate of collateral-related mortgage failures.[23] In the model, we calculate mortgage failure rates as $F_a(\bar{a})$, the probability that the appraisal a falls below the boundary $\bar{a}$ below which the buyer reneges on the transaction.

We then compute average "appraisal deviations" within each σ-decile of counties:

$$ApprDev_i = p_i^{under} E\left[1 - a_i \mid a_i < 1\right] \tag{25}$$

$ApprDev_i$ is the quantity plotted in Panel (b) of Figure 2: it is the product of the under-appraisal probability and the expectation of the percentage deviation of appraisal prices to sale prices conditional on under-appraisal. Figure 2 shows $ApprDev_i$ is strongly related to

[23]To be precise, we calculate the mortgage failure rate as $\frac{collFailure_c}{mortgage_c}$, where $collFailure_c$ is the total number of collateral-related mortgage failures in county c, from the HMDA data, and $mortgage_c$ is the total number of mortgage applications in county c.

price dispersion, since underappraisals are more likely to occur, and tend to be larger, in high-σ areas. Since we only observe appraisals on successful transactions in the data, we obtain the corresponding moment in the model by calculating $ApprDev_i$ in expression (25) conditional on appraisal values that do not result in transaction failure.[24]

Finally, we use the implied relationship between LTP and price dispersion from Column 3 of Table 5 as a target moment. We compute this in the model by running a simple OLS regression of the model-predicted average final loan size l_{final} on price dispersion σ, where one σ-decile is one data point.

Methodology and Identification With 10 σ-deciles, we have 21 moments in total: 10 failure probabilities, 10 appraisal deviations, and a single l_{final}-to-σ regression coefficient. The GMM problem is thus overidentified; so we choose parameters to minimize a weighted sum of moment errors. We set relative weights on each of the failure probability moments equal to the inverse of the approximate variances of each moment. For example, the ratio of weights on the squared moment errors of the 2nd and 8th σ-deciles is equal to the ratio of the estimated variances of these moments. Analogously, we use inverse variance relative weights for the appraisal deviation moments. We show how these variances are calculated in Appendix H.2.

We then manually increase the weight on the l_{final}-to-σ regression coefficient target moment substantially, and increase the weights on the appraisal deviation moments, relative to what would be implied by pure inverse variance weighting. We increase the regression coefficient's weight because, while it is estimated less precisely than the fail probability and appraisal deviation moments, it is economically important in pinning down the level of the costs facing buyers. We increase the appraisal deviation moments' weights because, while they are estimated around 10 times less precisely than fail probabilities, we judged the model fit to be better economically when the relative errors on these two sets of moments are similar, though this choice has a small quantitative effect on counterfactuals. Thus, our GMM

[24] In principle, we could target either $ApprDev_i$, or the probability of under-appraisal, in each σ-decile. We cannot target both, as the model has difficulty simultaneously matching both moments. This is because, as we show in panel (c) of Figure 7, the distribution of appraisal values, conditional on under-appraisal, is fairly long-tailed in the data. However, in the model, the consumer tends to renege on the transaction when appraisal values are too low, so the conditional appraisal distribution in the model is truncated from below. Thus, if we match appraisal probabilities in the model and the data, $ApprDev_i$ would tend to be much higher in the data than in the model. We choose to target the conditional appraisal deviation, because this appears to be a better measure of the downward pressure that under-appraisals generate for sale prices compared to the simple under-appraisal probability.

procedure essentially requires the model to match the l_{final}-to-σ relationship in the data well, and conditional on this, finds parameters which minimize a weighted sum of squared errors on the fail probability and appraisal deviation moments.

The intuition behind how these moments pin down the levels of our parameters is as follows. The magnitude of $ApprDev_i$ depends on the shape of the appraisal distribution, which is controlled largely by the appraisal bias parameter b. The relative magnitude of the penalty parameter ψ, compared to the cost of transaction failure ζ, pins down the relative prevalence of transaction failures, versus successful transactions with under-appraisals and larger down payments. In particular, if the failure cost ζ is low relative to the under-appraisal consumption penalty ψ, consumers have higher incentives to let transactions fail because re-starting with a new house allows them to re-draw a new appraisal. The absolute levels of the costs ψ and ζ determine the level of the model-predicted l_{final}-to-σ relationship: when the costs of underappraisal are high, consumers have a stronger incentive to precautionarily lower loan size to alleviate underappraisal costs, so the l_{final}-to-σ relationship is stronger. Finally, the level of u_2' controls the relationship between σ and consumers' demand for credit: u_2' must be in a certain range to rationalize why consumers in high σ-deciles borrow less.

To calculate confidence intervals of estimated parameters and counterfactual quantities, we run a parametric bootstrap of the moment inputs. We repeatedly resample moment values from independent normal distributions centered at their baseline values, with analytical approximations to variances calculated in Appendix H.2. In each bootstrap sample, we also resample the $\hat{\sigma}_{a,i}$ values, also using its approximated variance. For each bootstrap sample, we then re-run the moment matching procedure to estimate bootstrapped moment values, and re-run all counterfactual analyses.

5.2 Parameter Estimates and Model Fit

Table 7 presents the estimates of model parameters. Our point estimate of u_2' is 0.00243, with relatively tight standard errors. This implies that the Euler equation ratio,

$$\frac{\partial U}{\partial c_1} \frac{1}{\beta^T (1+r)^T u_2'},$$

has a point estimate of 1.030. That is, if the consumer could borrow more at rate r, she would gain 2.99% more utility on the margin from borrowing a dollar in period 1 and repaying it in period 2. Thus, the consumer in our model is liquidity constrained: she wants to increase borrowing, but higher interest rates and underappraisal risk limit her ability to do so. Our estimate of appraisers' bias parameter (b) is 0.086, with fairly tight standard errors.

Our point estimate of the underappraisal consumption penalty (ψ) is 2.782, and the transaction failure cost (ζ) is estimated to be 0.583. These parameter estimates have larger standard errors than other parameter estimates because they are identified mainly based on the l_{final}-to-σ regression coefficient, which we estimate less precisely than other empirical moments. In terms of the magnitude, we construct a microfoundation for ψ in Appendix G.1 to show that a high value of ψ is defensible if we view underappraisal as a large, sudden consumption shock to the consumer that cannot be saved for in advance.

The ζ estimate seems implausibly large. This is partly a modelling artifact driven by our assumption that the cost of transaction failures is borne in the second period. Future payments are more costly than present payments in the model, due to both discounting and the curvature of current-period utility. Quantitatively, in our estimated model, a consumer suffers the same utility loss from paying a fraction ζ of house prices in the future, and paying a penalty of around \$13,054, or around 6.53% of house prices, in the homebuying period.[25] This number is still fairly large. The estimation infers high values of ψ and ζ because underappraisals in the model must be fairly costly – that is, ψ must be large – in order to rationalize the observed loan-size-to-dispersion relationship. But, since underappraisals are fairly rare, consumers can avoid underappraisals with high probability if they are willing to simply buy a different house when their current house underappraises. To rationalize the fact that we observe many successful transactions with underappraisals empirically, the estimated model thus infers that it must be quite costly for consumers to "redraw" appraisals by choosing purchasing a new house. In our counterfactuals, our estimates of ψ and ζ influence our "compensating variation" estimates of the dollar cost of price dispersion to consumers: if the true values of ψ and ζ are lower than our estimates here, the inferred costs of price dispersion to consumers will also be lower than our estimates.

[25]It is somewhat more intuitive to think of the failure penalty being charged in the present period in reality. We charge the failure penalty in linear second-period consumption terms purely for model tractability: we cannot allow the failure penalty to be incurred in the current model because this would cause the model to no longer be recursive; the model would thus much less tractable.

Figure 7 evaluates model fit. Panel (a) shows the estimated appraisal standard deviations $\hat{\sigma}_{a,i}$. In the data, the standard deviation of appraisals is monotonically higher for higher σ-deciles, and we feed this directly into the model. Panel (b) shows the CDF of appraisals in the model and the data for the fifth σ-deciles. The appraisal distribution in the model has slightly thinner left and right tails compared to the data, but matches the observed bunching of appraisals at 1, the relatively low probabilities of under-appraisal, and the relatively large probabilities of over-appraisal.

Panels (c) and (d) show, respectively, the values of the two sets of targeted moments, mortgage failure probabilities and appraisal deviations, in the model and the data. Empirically, both moments are monotone with respect to changes in σ: counties with higher idiosyncratic price dispersion have monotonically higher collateral-related mortgage failures and higher appraisal deviations. The fitted model matches the average level of both moments fairly well; the main difference is that the relationship between both outcomes and σ is slightly stronger in the model than in the data.

Panel (e) shows the relationship between σ and average final loan size in the model, alongside a line with slope equal to the l_{final}-to-σ regression coefficient, normalized to have the same mean as the model series. These two lines are very close to each other, due to the fact that we set a high moment error weight on the the l_{final}-to-σ coefficient in our moment matching procedure.

5.3 Decomposition of Channels

Using our model, we evaluate how each of the two channels contributes to driving variation in loan rejections, LTPs, and interest rates. Figure 8 presents our results. In short, we find that the appraisal risk channel has a larger effect on loan-to-price ratios and rejection rates, whereas the collateral recovery channel has a large effect on interest rates.

We evaluate the magnitude of the collateral recovery channel by allowing lenders' rate menus to vary according to σ but shutting down the appraisal risk channel by setting appraisal noise constant across σ-deciles – $A_{raw} \sim N(P, \bar{\sigma})$ for all deciles – and re-calculating mortgage market outcomes. Analogously, to evaluate the magnitude of the appraisal risk channel, we shut down the collateral recovery channel, making lenders' rate menus constant across σ-deciles, but assuming that the appraisal distribution varies across σ-deciles.

Panel (a) shows results on the interest rate-price dispersion relationship. We find that this relationship is mainly driven by the collateral recovery channel. Shutting down the appraisal risk channel leaves the σ-to-interest-rate relationship quantitatively unaffected. Shutting down the collateral recovery channel in fact causes the relationship to change sign: consumers in high-dispersion areas receive *lower* interest rates. This is because, facing higher appraisal noise, consumers lower target loan sizes, and as a result also receive lower interest rates. However, this effect is more than 10 times smaller than compared to the collateral recovery effect.

Panel (b) analyzes mortgage failures. We find that appraisal risk is the main driver, with an effect more than 100 times the magnitude, and of opposite sign, to the collateral recovery effect. Panel (c) analyzes loan-to-price ratios, which are also mostly driven by appraisal risk. We further divide the effect of appraisals on LTP into two separate effects: an *ex-ante* effect based on borrowers choosing lower-target LTPs, and another *ex-post* effect based on realized appraisals. The ex-post effect captures the fact that, when appraisals are noisier, the gap between l_{final} and l tends to be larger, putting downward pressure on l_{final}. We find that ex-ante appraisal risk is the main driver, with ex-post appraisal pressure playing a negligible role.

6 Policy Counterfactuals and Implications

6.1 Price Dispersion and Consumer Willingness-to-Pay: A Compensating Variation Approach

How much does price dispersion affect consumers' willingness-to-pay for houses? We can give a partial answer to this question using an approach based on the idea of *compensating variation*. In our model, a consumer in the ith σ-decile achieves expected utility equal to the optimized objective value in (22), given price dispersion in her σ-decile. This optimized value tends to be lower when σ increases, since the consumer then faces higher interest rates and greater under-appraisal risk. The optimized value will *increase* if house *prices* P are lower, since consumers pay a lower amount for the same house. Thus, a natural way to quantify the utility costs of price dispersion is to ask: if consumer A faces higher price dispersion than consumer B, how much lower would consumer A's house price have to be, relative to consumer B, for A to achieve the same utility as B?

We apply this methodology to evaluate the cost of price dispersion in each σ-decile, relative to the lowest-dispersion σ-decile. Formally, we first calculate homebuyers' expected utility in the 1st σ-decile. Then, in every other σ-decile, we calculate expected utility for different values of the house price P, and find how much P would have to decrease for these homebuyers to achieve the same utility as homebuyers in the first σ-decile. The results, displayed in Figure 9, show that the utility costs of price dispersion are significant. Households in the 5th σ-decile would need to face 0.658% lower house prices, and consumers in the 10th σ-decile would need 1.679% lower prices, to attain the same expected utility as consumers in the 1st σ-decile.

Liquidity discounts. Beyond consumer WTP, a natural question is how dispersion affects house *prices*: if dispersion is higher in county X than county Y, and houses are otherwise identical, how much lower would equilibrium prices for identical houses be in county X? We refer to this price difference as a *liquidity discount*. Liquidity discounts are much harder to estimate than consumer WTP differences. A decrease in price dispersion increases consumer WTP, essentially shifting the housing demand curve outwards. How much a demand shift impacts equilibrium prices depends on the elasticities of housing supply and demand. Aggregate housing supply elasticity estimates exist in the literature (Saiz, 2010), but this elasticity likely varies substantially across regions and over time. The elasticity of housing *demand* is similarly difficult to estimate, and we do not believe the literature has reached consensus even on an aggregate value of this quantity.[26]

Under some assumptions, however, WTP changes are an upper bound for liquidity discounts. If house supply is perfectly elastic, house prices always equal production costs, so liquidity discounts are 0. If house supply is perfectly inelastic, prices equal marginal consumers' WTP for houses, so WTP changes pass through one-to-one to prices. With im-

[26]Our model cannot produce a nontrivial housing demand elasticity because we assume a representative consumer. If we assumed any fixed "outside option" level of utility the consumer would receive from not purchasing houses, house demand would thus be perfectly elastic at the price which makes the consumer indifferent between house purchasing and the outside option.

perfectly elastic supply and demand, liquidity discounts should thus fall between 0 and the "compensating variation" amounts that we calculate in Figure 9.

6.2 Implications for Desktop Appraisals

Our findings have implications for the shift from human appraisals to automated appraisals. In 2021, the FHFA announced that banks and mortgage lenders could use automated appraisal software in place of human appraisals.[27] There are a few different ways this could affect mortgage market outcomes. Automated appraisals may not display human appraisers' tendency to upwards bias under-appraisals towards the transaction price: this would tend to increase the prevalence and size of under-appraisals. Automated appraisals may also be less noisy than human appraisals, for example if they can more easily use a large amount of historical sales data; this would tend to decrease underappraisal pressure. We evaluate both potential effects within our calibrated model.

In Panel A of Table 8, we show results assuming that automated appraisals completely remove human biases from appraisals, thus setting $b = 0$. This scenario has severe negative effects on mortgage credit availability. Mortgage failures dramatically increase, by 10.540pp in the first σ-decile, and 13.561pp in the tenth: this represents a 4-7x increase in the rate of mortgage failures, relative to the human-appraiser status quo. Borrowers respond to increased underappraisal risk by decreasing target loan size, by an amount ranging from 2.360pp to 2.773pp; as a result, interest rates in fact slightly decrease. In column 4, we apply the "compensating variation" approach of Subsection 6.1: we find that house prices would have to decrease by 5.347% in the first σ-decile to 6.262% in the tenth, in order for consumers to be indifferent between automated appraisals and the human-appraiser status quo. Thus, in this scenario, automated appraisals would have large utility costs for consumers.

Thus, naively implementing automated appraisals, without accounting for the human tendency to cluster appraisals at the transaction price, could dramatically decrease mortgage credit availability. This is in principle easy to account for; under the assumptions of our model, if automated appraisals were simply uniformly shifted upwards by a fraction b of

[27] https://www.americanbanker.com/news/fhfa-will-make-desktop-home-appraisals-a-permanent-option

transaction prices, then outcomes would be unchanged from human appraisals.[28]

On the other hand, if automated appraisals decreased appraisal *variance*, mortgage credit availability would improve. We demonstrate this in another counterfactual, assuming that automated appraisals have 50% lower variance than human appraisers, within each σ-decile. To focus on the variance reduction effect, we continue to remove human biases by setting $b = 0$, but also shift the mean of appraisals upwards in each σ-decile, to keep underappraisal probability unchanged from the baseline calibration. Thus, in this counterfactual, automated appraisers produce underappraisals with equal probability to human appraisers, but underappraisals are on average smaller when they occur.

The results are shown in Panel B of Table 8: this scenario modestly improves mortgage credit availability. Mortgage failures decrease slightly, by 0.713pp in the first σ-decile and 1.600pp in the tenth. Buyers increase target loan size in response to reduced underappraisal risk, by around 0.109pp to 0.166pp. Since buyers are better off with automated appraisals, they are willing to pay slightly *higher* prices for identical houses; column 4 shows that this "compensating variation" ranges from 0.364% to 0.694% of house prices.

Thus, automated appraisals may have positive or negative effects on mortgage credit availability, depending on how they are implemented. Removing human biases, *ceteris paribus*, would dramatically increase underappraisal pressure; reducing appraisal variance, while compensating for human biases, would modestly reduce underappraisal pressure and increase mortgage credit availability.

7 Conclusion

An important policy goal of housing regulators in the US is to increase housing affordability. In this paper, we have shown that house value uncertainty affects mortgage credit provision in the US residential real estate market. Houses differ substantially in their degree of idiosyncratic price dispersion, which affects their value as collateral and thus the avail-

[28] Note that if we simply shifted automated appraisals upwards uniformly, *overappraisals* would be more frequent and larger when they occur. In our model, this has no effect on outcomes, because loan size is constrained by the small of the transaction price and the appraisal price, so overappraisals do not affect outcomes. In practice, if systematic overappraisals are viewed as undesirable, another policy which is equivalent within our model would be to add b to appraisals only if they are originally below the transaction price, thus mimicking the behavior of human appraisers.

ability of mortgage credit. This effect is partially due to a classical channel involving the fair pricing of collateral recovery risk, and partly through a novel channel involving the effect of idiosyncratic price dispersion on appraisal noise and its interaction with regulatory constraints.

An interesting implication of our results is that urban policy, in shaping characteristics of the housing stock, may also influence mortgage credit availability. If urban policymakers encouraged rebuilding and renovation, and set zoning rules in a way that promotes the development of newer and more standardized housing, the aggregate value uncertainty of the housing stock would decrease. Lenders would lend more against these houses, potentially contributing to increasing homeownership rates for low-income households, even if these policies do not decrease house prices. Interestingly, this is a channel through which housing stock renewal disproportionately benefits low-income households and first-time homebuyers, since these households tend to be most reliant on credit for homeownership.

References

Adelino, Manuel, William B McCartney, and Antoinette Schoar, 2020, The role of government and private institutions in credit cycles in the us mortgage market, Technical report, National Bureau of Economic Research.

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru, 2017, Policy intervention in debt renegotiation: Evidence from the home affordable modification program, *Journal of Political Economy* 125, 654–712.

Agarwal, Sumit, Gene Amromin, Souphala Chomsisengphet, Tim Landvoigt, Tomasz Piskorski, Amit Seru, and Vincent Yao, 2022, Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinance Program, *The Review of Economic Studies* rdac039.

Albrecht, James, Pieter A. Gautier, and Susan Vroman, 2016, Directed search in the housing market, *Review of Economic Dynamics* 19, 218 – 231, Special Issue in Honor of Dale Mortensen.

Andersen, Steffen, Cristian Badarinza, Lu Liu, Julie Marx, and Tarun Ramadorai, 2022, Reference dependence in the housing market, *American Economic Review* 112, 3398–3440.

Anundsen, André K, Erling Røed Larsen, and Dag Einar Sommervoll, 2020, Strategic price-setting and incentives in the housing market, Technical report, Oslomet Working paper series 2020: 1.

Bayer, Patrick, Christopher Geissler, Kyle Mangum, and James W Roberts, 2020, Speculators and middlemen: The strategy and performance of investors in the housing market, *The Review of Financial Studies* 33, 5212–5247.

Benmelech, Efraim, 2009, Asset Salability and Debt Maturity: Evidence from Nineteenth-Century American Railroads, *The Review of Financial Studies* 22, 1545–1584.

Benmelech, Efraim, and Nittai K. Bergman, 2008, Liquidation Values and the Credibility of Financial Contract Renegotiation: Evidence from U.S. Airlines, *The Quarterly Journal of Economics* 123, 1635–1677, Publisher: Oxford University Press.

Benmelech, Efraim, and Nittai K Bergman, 2009, Collateral pricing, *Journal of Financial Economics* 91, 339–360.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra, 2019, Regional heterogeneity and the refinancing channel of monetary policy, *The Quarterly Journal of Economics* 134, 109–183.

Bian, Bo, 2021, The real effects of legal and trading frictions in collateralization, *Working Paper* .

Blackburn, McKinley, and Todd Vermilyea, 2007, The role of information externalities and scale economies in home mortgage lending decisions, *Journal of Urban Economics* 61, 71–85.

Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru, 2018, Beyond the balance sheet model of banking: Implications for bank regulation and monetary policy, Technical report, National Bureau of Economic Research.

Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru, 2020, ibuyers: Liquidity in real estate markets?, *Available at SSRN 3616555* .

Calem, Paul, Jeanna Kenney, Lauren Lambie-Hanson, and Leonard Nakamura, 2021, Appraising home purchase appraisals, *Real Estate Economics* 49, 134–168.

Case, Karl E, and Robert J Shiller, 1989, The efficiency of the market for single-family homes, *The American Economic Review* 125–137.

Collier, Benjamin L, Cameron Ellis, and Benjamin J Keys, 2021, The cost of consumer collateral: Evidence from bunching, Technical report, National Bureau of Economic Research.

Conklin, James, N. Edward Coulson, Moussa Diop, and Thao Le, 2020, Competition and Appraisal Inflation, *The Journal of Real Estate Finance and Economics* 61, 1–38.

DeFusco, Anthony A, 2018, Homeowner borrowing and housing collateral: New evidence from expiring price controls, *The Journal of Finance* 73, 523–573.

DeFusco, Anthony A, Stephanie Johnson, and John Mondragon, 2020, Regulating household leverage, *The Review of Economic Studies* 87, 914–958.

DeFusco, Anthony A, and John Mondragon, 2020, No job, no money, no refi: Frictions to refinancing in a recession, *The Journal of Finance* 75, 2327–2376.

Di Maggio, Marco, and Amir Kermani, 2017, Credit-induced boom and bust, *The Review of Financial Studies* 30, 3711–3758.

Di Maggio, Marco, Amir Kermani, Benjamin J Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao, 2017, Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging, *American Economic Review* 107, 3550–88.

Dokko, Jane, Benjamin J Keys, and Lindsay Relihan, 2019, Affordability, financial innovation, and the start of the housing boom .

Eriksen, Michael D., Hamilton B. Fout, Mark Palim, and Eric Rosenblatt, 2019, The influence of contract prices and relationships on appraisal bias, *Journal of Urban Economics* 111, 132–143.

Eriksen, Michael D, Chun Kuang, and Wenyu Zhu, 2024, Attribute misreporting and appraisal bias, *Review of Finance* 28, 1663–1686.

Fout, Hamilton, Nuno Mota, and Eric Rosenblatt, 2021, When Appraisers Go Low, Contracts Go Lower: The Impact of Expert Opinions on Transaction Prices, *The Journal of Real Estate Finance and Economics* .

Genesove, David, and Christopher Mayer, 2001, Loss aversion and seller behavior: Evidence from the housing market, *The Quarterly Journal of Economics* 116, 1233–1260.

Genesove, David, and Christopher J Mayer, 1997, Equity and time to sale in the real estate market, *The American Economic Review* 255–269.

Giacoletti, Marco, 2021, Idiosyncratic risk in housing markets, *The Review of Financial Studies* 34, 3695–3741.

Giacoletti, Marco, and Victor Westrupp, 2017, The Risk-Adjusted Performance of Asset Flippers, SSRN Scholarly Paper ID 2892439, Social Science Research Network, Rochester, NY.

Gilbukh, Sonia, and Paul Goldsmith-Pinkham, 2024, Heterogeneous real estate agents and the housing cycle, *The Review of Financial Studies* 37, 3431–3489.

Glaeser, Edward L, and Jesse M Shapiro, 2003, The benefits of the home mortgage interest deduction, *Tax policy and the economy* 17, 37–82.

Greenwald, Daniel L., 2016, The Mortgage Credit Channel of Macroeconomic Transmission, *SSRN Electronic Journal* .

Gupta, Arpit, Christopher Hansman, and Pierre Mabille, 2024, Financial constraints and the racial housing gap, *Available at SSRN 3969433* .

Guren, Adam M, 2018, House price momentum and strategic complementarity, *Journal of Political Economy* 126, 1172–1218.

Hartman-Glaser, Barney, and William Mann, 2017, Collateral constraints, wealth effects, and volatility: Evidence from real estate markets, *Working Paper* .

Jansen, Mark, Fabian Nagel, Constantine Yannelis, and Anthony Lee Zhang, 2022, Data and welfare in credit markets, *Available at SSRN 4015958* .

Jiang, Erica Xuewei, 2023, Financing competitors: Shadow banks' funding and mortgage market competition, *The Review of Financial Studies* 36, 3861–3905.

Kermani, Amir, and Yueran Ma, 2020, Two tales of debt, Technical report, National Bureau of Economic Research.

Kermani, Amir, and Francis Wong, 2021, Racial disparities in housing returns, Technical report, National Bureau of Economic Research.

Kotova, Nadia, and Anthony Lee Zhang, 2021, Liquidity in residential real estate markets, *Available at SSRN 3386353* .

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, 2011, The effects of quantitative easing on interest rates: channels and implications for policy, Technical report, National Bureau of Economic Research.

Kurlat, Pablo, and Johannes Stroebel, 2015, Testing for information asymmetries in real estate markets, *The Review of Financial Studies* 28, 2429–2461.

Lang, William W, and Leonard I Nakamura, 1993, A model of redlining, *Journal of Urban Economics* 33, 223–234.

Lustig, Hanno N, and Stijn G Van Nieuwerburgh, 2005, Housing collateral, consumption insurance, and risk premia: An empirical perspective, *The Journal of Finance* 60, 1167–1219.

Mian, Atif, and Amir Sufi, 2011, House prices, home equity-based borrowing, and the us household leverage crisis, *American Economic Review* 101, 2132–56.

Pan, Guangqian, Zheyao Pan, and Kairong Xiao, 2024, The shadow cost of collateral, *The Review of Financial Studies* hhae073.

Pancost, N Aaron, and Garrett Schaller, 2021, Measuring measurement error, *Available at SSRN* .

Pennington-Cross, Anthony, 2006, The value of foreclosed property, *Journal of Real Estate Research* 28, 193–214.

Piazzesi, Monika, and Martin Schneider, 2016, Housing and macroeconomics, *Handbook of macroeconomics* 2, 1547–1640.

Piskorski, Tomasz, and Amit Seru, 2018, Mortgage market design: Lessons from the great recession, *Brookings Papers on Economic Activity* 2018, 429–513.

Rampini, Adriano A, and S Viswanathan, 2010, Collateral, risk management, and the distribution of debt capacity, *The Journal of Finance* 65, 2293–2322.

Rampini, Adriano A, and S Viswanathan, 2020, Collateral and secured debt, *Unpublished working paper, Duke University* .

Rampini, Adriano A, and Sridhar Viswanathan, 2013, Collateral and capital structure, *Journal of Financial Economics* 109, 466–492.

Sagi, Jacob S, 2021, Asset-level risk and return in real estate investments, *The Review of Financial Studies* 34, 3647–3694.

Saiz, Albert, 2010, The geographic determinants of housing supply, *The Quarterly Journal of Economics* 125, 1253–1296.

Shleifer, Andrei, and Robert W Vishny, 1992, Liquidation values and debt capacity: A market equilibrium approach, *The Journal of Finance* 47, 1343–1366.

Sklarz, Michael, and Norman Miller, 2016, Adjusting loan to value (ltv) ratios to reflect value uncertainty, Technical report, Collateral Analytics Research.

Stroebel, Johannes, 2016, Asymmetric information about collateral values, *The Journal of Finance* 71, 1071–1112.

Titman, Sheridan, and Roberto Wessels, 1988, The determinants of capital structure choice, *The Journal of Finance* 43, 1–19.

Wheaton, William C., 1990, Vacancy, Search, and Prices in a Housing Market Matching Model, *Journal of Political Economy* 98, 1270–1292.

Wood, Simon N, 2017, *Generalized additive models: an introduction with R* (CRC press).

Zhou, Hanqing, Yuan Yuan, Christopher Lako, Michael Sklarz, and Charles McKinney, 2015, Foreclosure discount: definition and dynamic patterns, *Real Estate Economics* 43, 683–718.

Figures

Figure 1. Stylized Facts about Price Dispersion Estimates

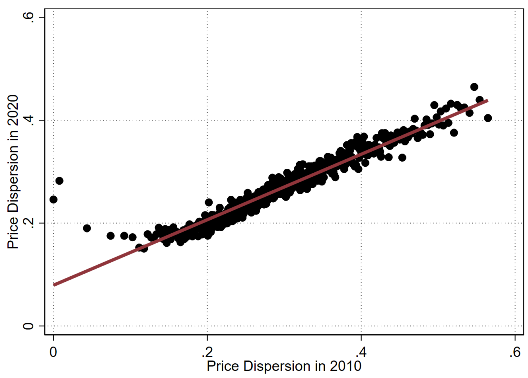

(a) Zipcode Dispersion: 2010 vs 2020

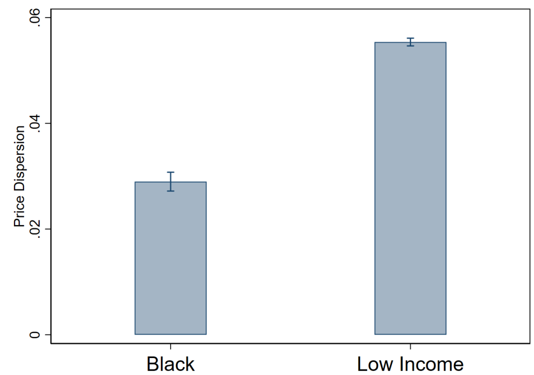

(b) Dispersion by Demographics

Notes: Panel (a) plots zipcode dispersion in 2020 against zipcode dispersion in 2010. Panel (b) shows the price dispersion difference between black-dominant zipcodes (black population share greater than 50%) and non-black dominant zipcodes (black population share less than 50%) conditional on income, as well as the price dispersion difference between high-income zipcodes and low-income zipcodes conditional on race. High and low income zipcodes are defined as above and below yearly median level, respectively. To obtain the values, we regress zipcode price dispersion on dummy variables for a zipcode being black-dominant, and whether the zipcode has below-median income; the figure shows the estimated coefficients and confidence intervals on these dummy variables. The sample includes annual zip level observations from 2000 to 2020. Zipcode demographic information is obtained from the 2008-2012 American Community Survey.

Figure 2. Price Dispersion and Appraisals

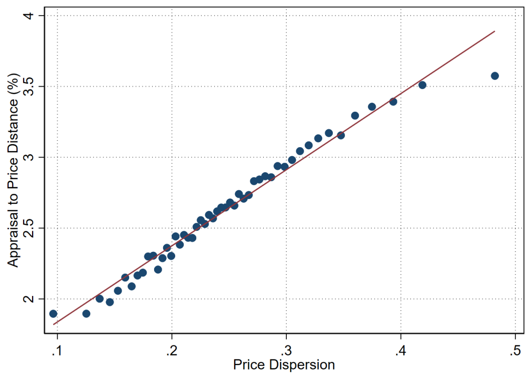

(a) Appraisal to Price Distance

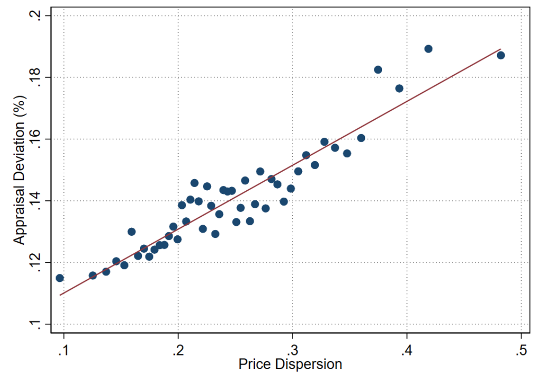

(b) Appraisal Deviation

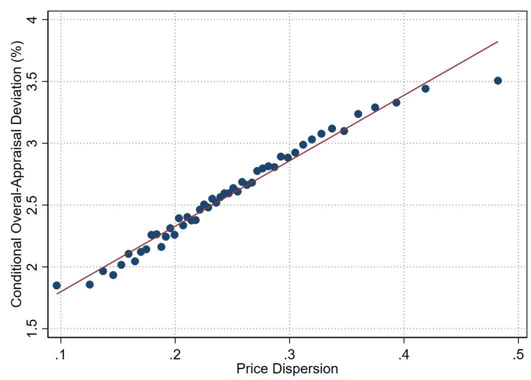

(c) Over-Appraisal Deviation

Notes: Panel (a) of this figure shows a binned scatter plot, where the y-variable is Appraisal-to-Price distance, defined as $\frac{|a_i-p_i|}{p_i}$. In panel (b), the y-variable is the appraisal deviation, defined as the product of the appraisal gap, and a dummy for under-appraisal, $\frac{|a_i-p_i|}{p_i}\mathbf{1}(a_i < p_i)$. In panel (c), the y-variable is the average over-appraisal percentage conditional on over-appraisal. In all panels, the x-variable is zipcode price dispersion. We divide all loans into 50 buckets based on zipcode house price dispersion. The sample includes loan level observations from 2000 to 2020. Appraisal values are obtained from the Corelogic LLMA data.

Figure 3. Home Purchase - Mortgage Origination Diagram

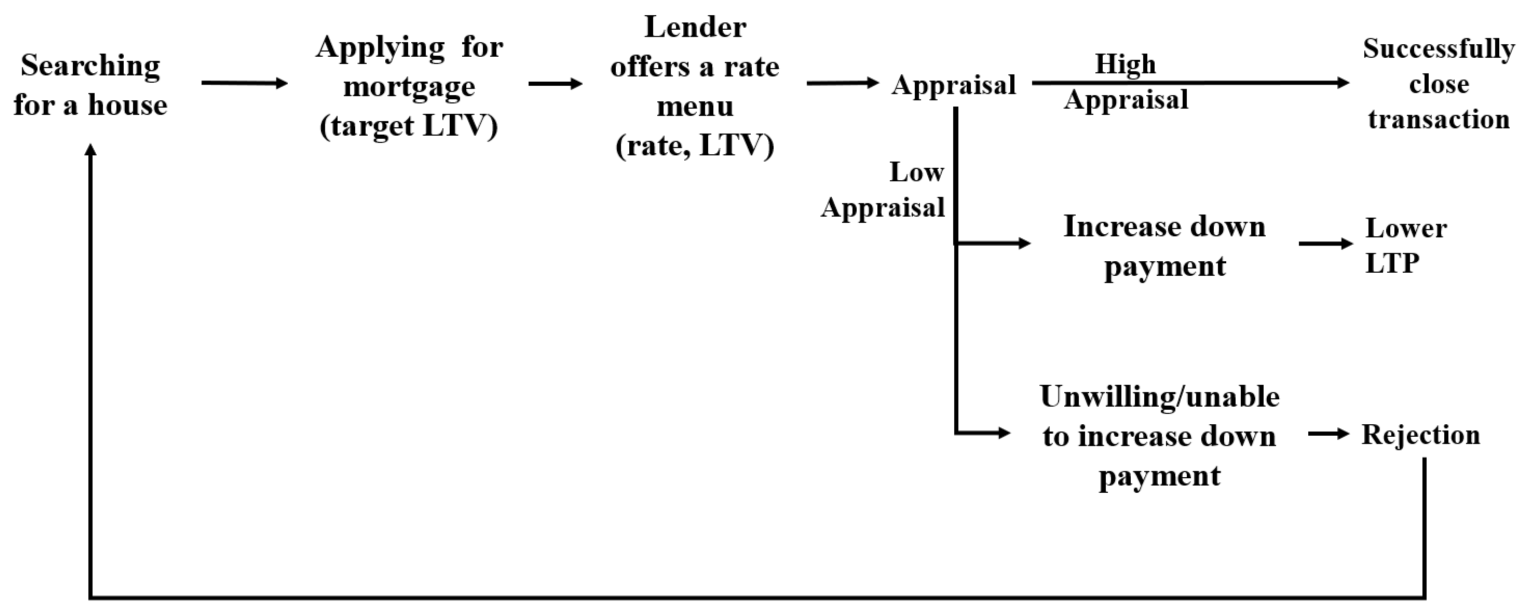

Figure 4. County Level House Price Dispersion and Credit Access

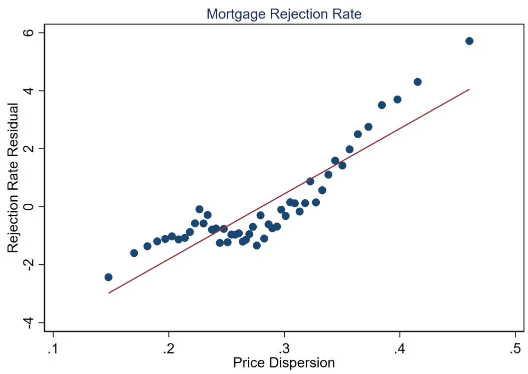

(a) Total Rejection

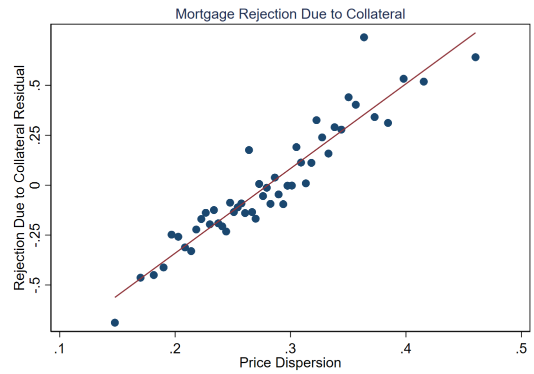

(b) Rejection due to Collateral

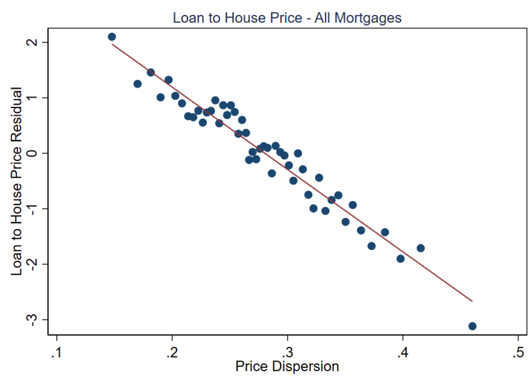

(c) Loan-to-Price

Notes: This figure shows the correlation between county level house price dispersion and various credit access outcomes. The y-variable in panels (a) and (b) are residualized county-level mortgage rejection rates and county-level rates of rejection due to collateral-related reasons, respectively, obtained from the HMDA data. We take the residuals of regressions of county-level rejection rates, or rates of rejection due to collateral-related reasons on county average log house price, credit score, and year fixed effects. In panel (c), the y-variable is the average county-level LTP residual from a regression of county average LTP on county-level house prices. The y-variable is obtained from the Corelogic LLMA data, and values are in percentage points. The underlying samples in all three panels include annual county observations from 2000 to 2017.

Figure 5. Property Level Mortgage Menu by House Price Dispersion

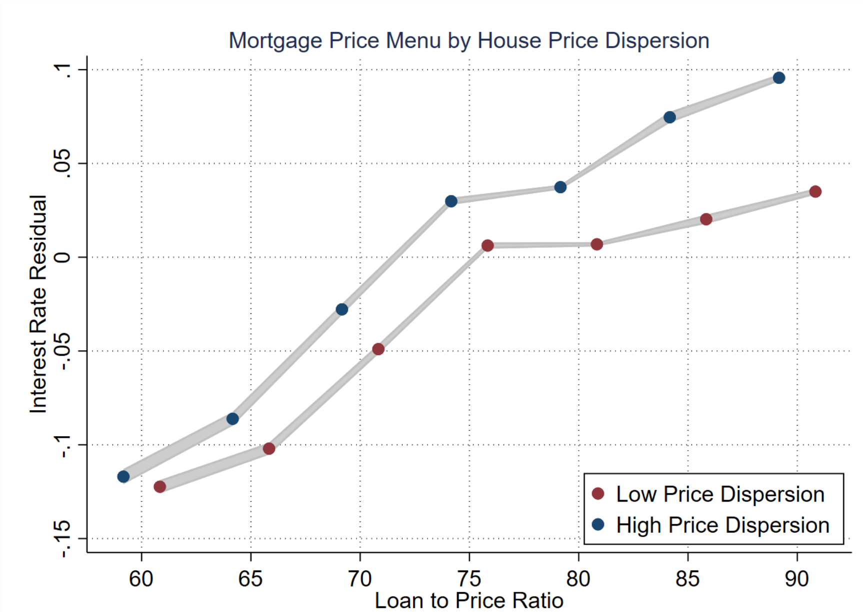

Notes: This figure shows the mortgage price menu (rate-LTP pair) by zip-level house price dispersion. The y-values are interest rate residuals from a regression of mortgage rates on borrower FICO, FICO-squared, DTI, DTI-squared conforming or jumbo indicator, and origination month fixed effects. The dots represent the average mortgage rate in each LTP bucket. The shaded area indicates a 95% confidence interval. The sample includes loan level observations of conventional loans in the Corelogic LLMA from 2000 to 2020.

Figure 6. Cross-Sectional Heterogeneity

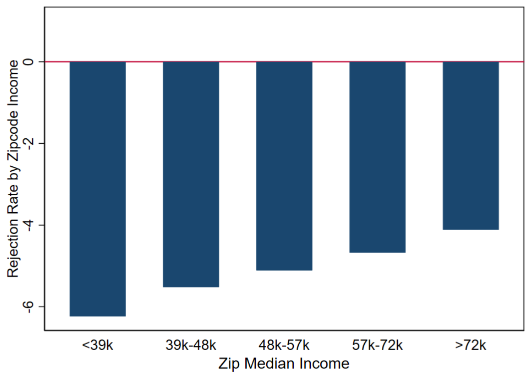

(a) Rejection

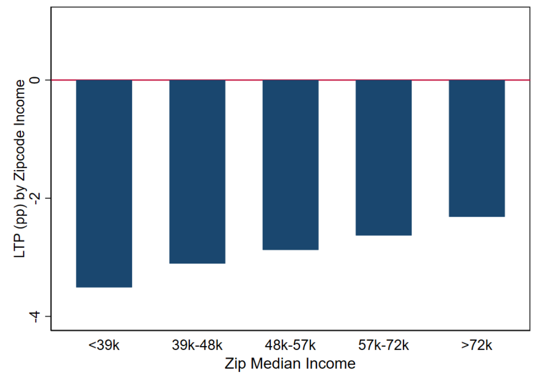

(b) LTP

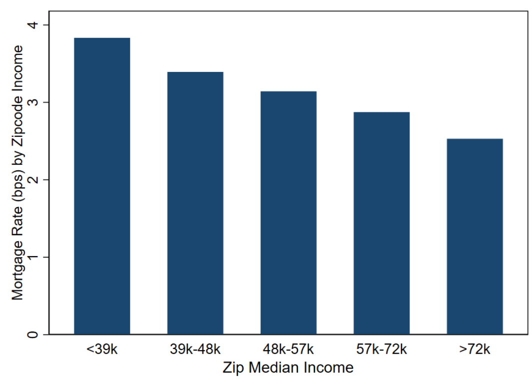

(c) Interest Rate

Notes: This figure presents how rejection rate, down payment, and mortgage rate would change if the house's price dispersion decreased to be equal to the bottom decile of price dispersion in our dataset. The x-axis are five equal-sized zipcode median income buckets, where zipcode median income is obtained from 2008-2012 American Community Survey. The y-variable is the impact on mortgage rejection rate in Panel (a), the impact on LTP in Panel (b), and the impact on interest rate in Panel (c). To obtain these values, we calculate the difference between a house's price dispersion and the bottom decile of price dispersion, and then multiply it respectively by the estimated effect of price dispersion on mortgage rejection (column 4 of Table 3), by the estimated effect on LTP (column 6 of Table 5), and by the estimated effect on mortgage rate (column 4 of Table 4). We find the average values in each zipcode income quintile.

Figure 7. Model Fit

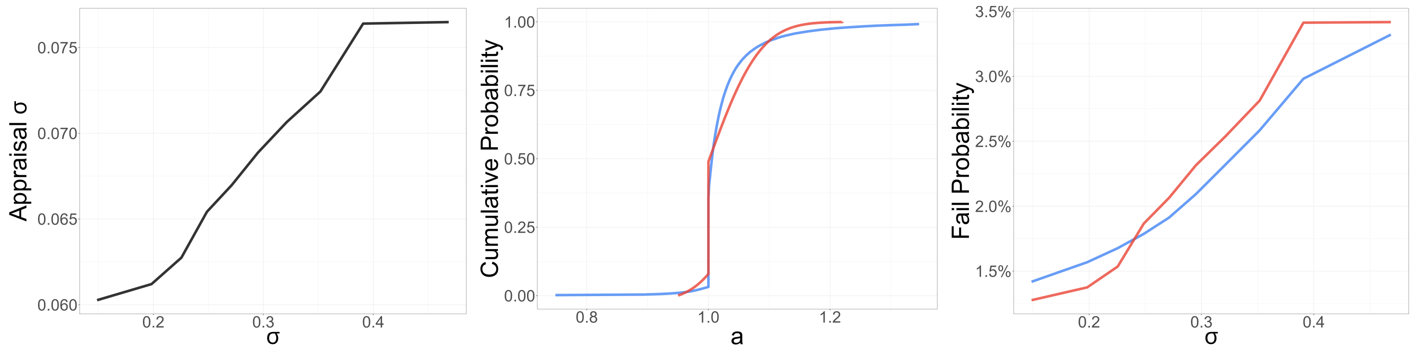

(a) Appraisal Noise **(b)** Appraisal-to-Price **(c)** Probabilities of Mortgage Failure

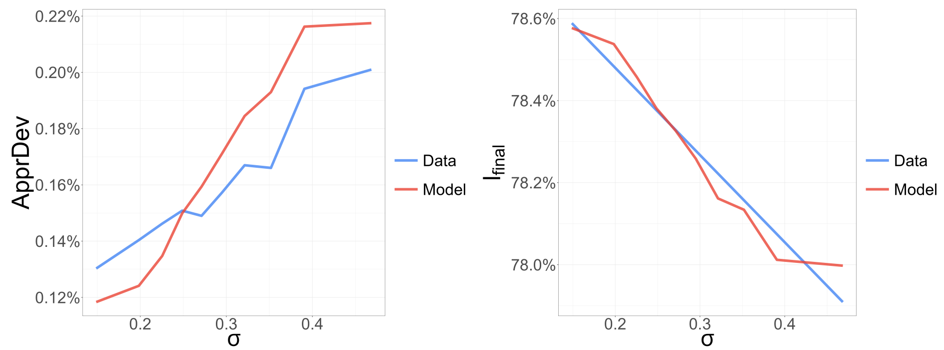

(d) Appraisal Deviation **(e)** Model Predicted LTP

Notes: Panel (a) shows estimated appraisal standard deviations σ_a on the y-axis, and estimated idiosyncratic price dispersion σ on the x-axis. Panel (b) shows the distribution of appraisal-over-price ratios a, in the data and the fitted model, for the 5th σ-decile (that is, counties with values of σ between the 40th and 50th percentiles). Panel (c) shows transaction failure probabilities in the data and in the fitted model. Panel (d) shows $ApprDev_i$, which is defined as $p_i^{under} E\left[\frac{a}{p} - 1 \mid under\right]$, in the data and in the fitted model. Panel (e) shows average values of l_{final} from the model, along with a line with slope equal to the l_{final}-to-σ regression coefficient we use as a target moment, normalized to have mean equal to the model-predicted loan-to-price ratio.

Figure 8. Decomposition of Channels

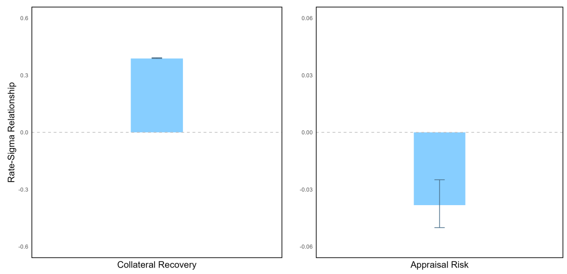

(a) Interest Rate (pp)

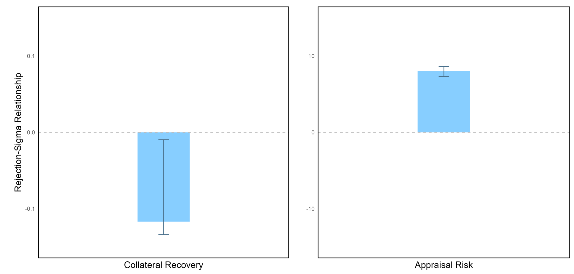

(b) Mortgage Failure (pp)

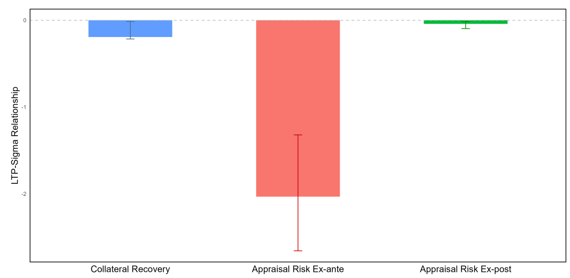

(c) Loan-to-Price (pp)

Notes: Panel (a), (b), and (c) respectively decompose the effect of price dispersion on interest rates, mortgage rejection rates, and loan-to-price ratios, into components attributable to the collateral recovery channel, and the appraisal channel. In all panels, the y-axis is the change in the outcome variable when sigma increases by 1. Note that the units differ from those in our reduced-form results, which standardize sigma: to convert the results in this figure to be comparable to the reduced-form results, all numbers should be multiplied by the standard deviation of sigma, which is 0.11. To calculate the effect of the collateral recovery channel, we shut off the effect of the appraisal channel by setting the appraisal standard deviation parameter to a constant across σ-deciles, but allowing the interest rate menu to vary across σ-deciles. Analogously, to calculate the effect of the appraisal risk channel, we shut off the collateral recovery channel, by assuming the rate menu does not vary across σ-deciles, allowing only appraisal standard deviations to vary. In Panel (c), we further decompose the appraisal risk channel into an ex-ante effect, which measures how target loan size l varies with interest rate buckets, due to buyers' precautionary decisions to decrease loan size; and an ex-post effect, which measures how the gap $l_{final} - l$ changes with σ, which measures how realized under-appraisals limit loan size. Note that, in both panels (a) and (b), the y-axes are not the same in the collateral recovery and appraisal risk panels.

Figure 9. Price Dispersion and Compensating Variation in Prices

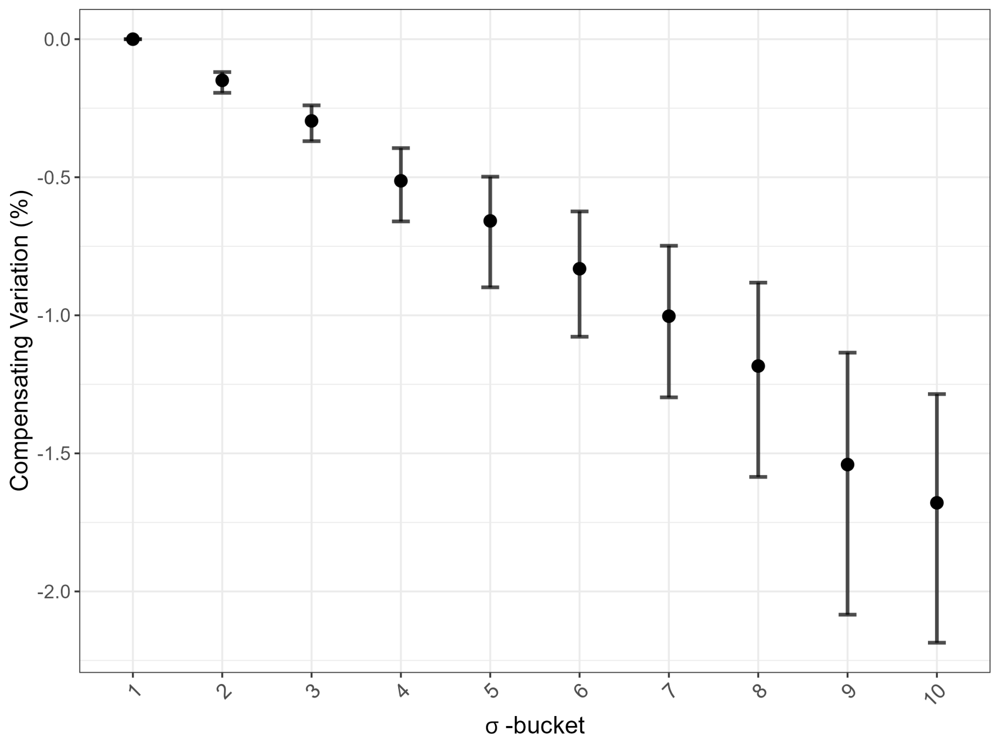

Notes: This figure presents the "compensating variation" in prices needed to offset the costs of price dispersion; that is, in each σ-decile, we show how much prices would have to decrease, in order for consumers to attain the same expected utility as in the first σ-decile. Error bars show 95% confidence intervals, calculated from 200 bootstrap moment samples, as we describe in Appendix H.2.1.

Table

Table 1: Summary Statistics

This table reports summary statistics for the three main datasets: the property sample from the Corelogic Deed and Tax datasets, the loan sample from the Corelogic LLMA dataset, and the mortgage application sample from the HMDA. The Corelogic samples span the time period 2000 to 2020. The HMDA sample spans 2000 to 2017.

	N	Mean	Stdev	P25	Median	P75
Property Level Sample						
Loan to Price	29M	85.42	15.65	80.00	89.68	98.19
Price Dispersion	29M	0.24	0.11	0.17	0.23	0.30
Sale Price (Thousand)	29M	273.02	224.93	140.30	215.00	332.50
Mortgage Amount (Thousand)	29M	222.40	163.53	121.80	182.16	275.79
Building Age	29M	27.12	25.95	6.00	20.00	42.00
Square Footage	29M	1,961.57	2,982.11	1,363.00	1,774.00	2,365.00
Loan Level Sample						
Loan to Price	4.8M	85.48	14.98	80.00	90.00	98.19
Zip Price Dispersion	4.8M	0.25	0.08	0.19	0.24	0.29
Sale Price (Thousand)	4.8M	280.83	242.84	143.50	218.00	340.00
Appraised to Price Ratio	4.8M	1.03	0.19	1.00	1.00	1.02
Mortgage Amount (Thousand)	4.8M	227.66	170.25	124.00	185.18	283.00
FICO	4.8M	725.35	61.39	681.00	735.00	778.00
Debt-to-Income	4.8M	37.23	11.28	29.85	38.00	44.69
Mortgage Application Sample						
Rejection Rate	49M	15.86	36.53	0.00	0.00	0.00
Rejection due to Collateral Reasons	49M	1.95	13.83	0.00	0.00	0.00
Zip Price Dispersion	49M	0.26	0.08	0.20	0.25	0.31
Applicant Income (Thousand)	49M	102.35	193.47	47.00	72.00	114.00
Loan-to-Income	49M	242.18	6,896.19	135.83	227.78	316.51
County Credit Score	49M	667.19	22.16	650.30	666.18	684.14

Table 2: Determinants of House Price Dispersion

This table presents the association between house price dispersion and house features (Panel A) and zipcode market condition (Panel B). All continuous variables are normalized by their standard deviations. We define a house as recently renovated if it has been renovated within 5 years before the transaction year. The sample includes house transactions from 2000 to 2020. Standard errors are clustered at county level. ***, **, * represent 1%, 5%, and 10% significance, respectively.

Panel A: House Features

	Estimated Price Dispersion				
	(1)	(2)	(3)	(4)	(5)
Building Age	0.04***	0.04***	0.04***		0.04***
	(0.002)	(0.002)	(0.002)		(0.002)
Recent Renovation		-0.01***			-0.01***
		(0.003)			(0.003)
Benchmark: Square-Footage < 1281					
[1282,1601]			-0.03***	-0.03***	-0.02***
			(0.002)	(0.002)	(0.002)
[1602,1970]			-0.03***	-0.03***	-0.02***
			(0.003)	(0.003)	(0.003)
[1971,2544]			-0.02***	-0.02***	-0.00
			(0.004)	(0.004)	(0.003)
> 2544			0.01**	0.01**	0.03***
			(0.005)	(0.004)	(0.004)
Benchmark: Bedrooms < 4					
=4				-0.01***	-0.01***
				(0.001)	(0.001)
>4				0.01***	0.01***
				(0.002)	(0.001)
Log House Price	-0.48***	-0.48***	-0.51***	-0.51***	-0.37***
	(0.028)	(0.028)	(0.025)	(0.025)	(0.021)
Log House Price Squared	0.50***	0.50***	0.51***	0.51***	0.38***
	(0.029)	(0.029)	(0.025)	(0.025)	(0.022)
County-Year FE	✓	✓	✓	✓	✓
R2	0.33	0.33	0.26	0.27	0.35
Observations	29M	29M	29M	29M	29M

Panel B: Zipcode Market Conditions

	Zipcode Price Dispersion			
	(1)	(2)	(3)	(4)
Gini Index	0.01***			0.01***
	(0.001)			(0.001)
Population Density		-0.01***		-0.01***
		(0.003)		(0.002)
Vacancy Share			0.03***	0.03***
			(0.002)	(0.002)
Year FE	✓	✓	✓	✓
R2	0.02	0.02	0.08	0.09
Observations	276,079	276,079	276,079	276,079

Table 3: Mortgage Rejections and Zip House Price Dispersion

This table presents loan level regression results about mortgage rejections. The outcome variable in Panel A is an indicator that equals 100 if a loan is rejected and 0 otherwise. The outcome variable in Panel B is an indicator that equals 100 if a loan is rejected due to collateral reasons and 0 otherwise. In both panels, columns 1-3 report OLS results, and columns 4-6 report 2SLS results. The explanatory variable of interest is zipcode house price dispersion in columns 1-3 and is the predicted zipcode price dispersion in columns 4-6, all scaled by its standard deviation. IV construction is described in Section 3.2.1. IV first stage results are presented in Table A2. Borrower/Loan controls include zipcode house price, credit score and the squared term, log income, loan type, and loan to income ratio and its square term. The sample includes loan level observations from 2000 to 2017. Standard errors are clustered at county level. ***, **, * represent 1%, 5%, and 10% significance, respectively.

	OLS			2SLS		
	(1) Full	(2) Securitized	(3) Portfolio	(4) Full	(5) Securitized	(6) Portfolio
Panel A: Rejection						
Zip Price Dispersion	1.40***	1.42***	0.81***	2.31***	2.30***	2.22***
	(0.093)	(0.103)	(0.116)	(0.221)	(0.234)	(0.433)
Local Controls	✓	✓	✓	✓	✓	✓
County-Year FE	✓	✓	✓	✓	✓	✓
Lender-Year FE	✓	✓	✓	✓	✓	✓
Rejection Mean	15.9%	16.4%	16.2%	15.9%	16.4%	16.2%
R2	0.16	0.18	0.17	-	-	-
Observations	47M	34M	3.6M	47M	34M	3.6M
Underidentification t-stat				77.87	74.65	19.77
Underidentification p-value				0.00	0.00	0.00
Weak identification t-stat				30.75	29.81	13.74
Panel B: Rejection Due to Collateral						
Zip Price Dispersion	0.50***	0.54***	0.38***	0.79***	0.86***	0.73***
	(0.036)	(0.038)	(0.054)	(0.065)	(0.073)	(0.095)
Local Controls	✓	✓	✓	✓	✓	✓
County-Year FE	✓	✓	✓	✓	✓	✓
Lender-Year FE	✓	✓	✓	✓	✓	✓
Rejection due to Collateral Mean	2.0%	2.0%	2.3%	2.0%	2.0%	2.3%
R2	0.05	0.05	0.09	-	-	-
Observations	47M	34M	3.6M	47M	34M	3.6M
Underidentification t-stat				77.87	74.65	19.77
Underidentification p-value				0.00	0.00	0.00
Weak identification t-stat				30.75	29.81	13.74

Table 4: Price Dispersion and Cost Menu

This table presents loan-level regression results about the "cost menu", that is, interest rates controlling for LTPs. The outcome variable is loan-level interest rate, in bps. Columns 1-3 present OLS results, and columns 4-6 present 2SLS results. Columns 1 and 4 use the full sample. Columns 2 and 5 use securitized conventional loans (i.e., non-FHA loans that are securitized). Columns 3 and 6 use portfolio conventional loans (i.e., non-FHA loans that are held on lenders' balance sheets). The explanatory variable of interest is zipcode house price dispersion in columns 1-3 and is the predicted zipcode price dispersion in columns 4-6, all scaled by its standard deviation. IV construction is described in Section 3.2.1. IV first stage results are presented in Table A2. Borrower and loan controls include log house price, FICO score, FICO squared, LTV, LTV squared, DTI, DTI-squared, and loan type. The sample includes loan level observations from 2000 to 2020. Standard errors are clustered at county level. ***, **, * represent 1%, 5%, and 10% significance, respectively.

	OLS			2SLS		
	(1) Full	(2) Securitized	(3) Portfolio	(4) Full	(5) Securitized	(6) Portfolio
Zip Price Dispersion	0.89***	1.31***	1.26***	1.42***	1.68***	3.03***
	(0.123)	(0.109)	(0.372)	(0.306)	(0.331)	(0.683)
LTP	0.69***	0.58***	1.89***	0.69***	0.58***	1.88***
	(0.090)	(0.039)	(0.186)	(0.090)	(0.039)	(0.184)
Borrower and Loan Controls	✓	✓	✓	✓	✓	✓
Origination Month FE	✓	✓	✓	✓	✓	✓
County-Year FE	✓	✓	✓	✓	✓	✓
R2	0.87	0.88	0.87	-	-	-
Observations	4.8M	2.3M	1.1M	4.8M	2.3M	1.1M
Underidentification t-stat				68.08	71.05	60.10
Underidentification p-value				0.00	0.00	0.00
Weak identification t-stat				23.17	20.77	18.35

Table 5: Property-Level House Price Dispersion and LTP

This table presents property-level regression results on the relationship between price dispersion and mortgage LTPs. Columns 1-3 present OLS results. Columns 4-6 present IV results. In all columns, the outcome variable is the loan level loan-to-sale-price ratio. The explanatory variable of interest in columns 1-3 is property-level house price dispersion, scaled by its standard deviation, and is the predicted price dispersion in columns 4-6. IV construction is described in Section 3.2.1. IV first stage results are presented in Table A2. Controls include the transaction price of the property, mortgage type, mortgage term, and resale indicator. The sample includes property transaction level observations from 2000 to 2020. Standard errors are clustered at county level. ***, **, * represent 1%, 5%, and 10% significance, respectively.

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
Price Dispersion	-0.43***	-0.21***	-0.23***	-1.25***	-1.30***	-1.30***
	(0.042)	(0.036)	(0.033)	(0.118)	(0.108)	(0.104)
Controls	✓	✓	✓	✓	✓	✓
Transaction Date FE	✓	✓	✓	✓	✓	✓
County-Year FE		✓	✓		✓	✓
Lender-Year FE			✓			✓
R2	0.34	0.36	0.40	-	-	-
Observations	28M	28M	28M	28M	28M	28M
Underidentification test statistic				175.44	162.02	165.68
Underidentification test p-value				0.00	0.00	0.00
Weak identification test statistic				189.39	202.32	197.81

Table 6: Ex-Post Performance

This table analyzes the relationship between price dispersion and the ex-post performance of mortgage loans. Columns 1 and 4 use full sample. Columns 2 and 5 use securitized conventional loans (i.e., non-FHA loans that are securitized). Columns 3 and 6 use portfolio conventional loans (i.e., non-FHA loans that are held on lenders' balance sheets). The outcome variable is an indicator for default, which is equal to 100 if the loan defaults in two years since origination, and 0 otherwise. The explanatory variable of interest is zipcode house price dispersion in columns 1-3 and is the predicted zipcode price dispersion in columns 4-6, all scaled by its standard deviation. Non-reported controls include house price, loan type, and the squared-terms of FICO, DTI, and LTV. The sample includes all loans originated from 2000 to 2018. Since we need at least two-year performance to define *default*, we remove loans originated after 2018 from the full sample for this analysis. Standard errors are clustered at county level. ***, **, * represent 1%, 5%, and 10% significance, respectively.

	OLS			2SLS		
	(1) Full	(2) Securitized	(3) Portfolio	(4) Full	(5) Securitized	(6) Portfolio
Zip Price Dispersion	0.21***	0.15***	0.19***	0.03	-0.05	0.06
	(0.054)	(0.036)	(0.063)	(0.119)	(0.117)	(0.141)
Interest Rate	2.18***	2.39***	1.68***	2.18***	2.40***	1.69***
	(0.125)	(0.175)	(0.131)	(0.127)	(0.178)	(0.132)
FICO	-65.97***	-60.93***	-56.81***	-65.97***	-60.92***	-56.81***
	(0.703)	(1.042)	(1.639)	(0.702)	(1.040)	(1.642)
DTI	0.37***	-0.06	0.33***	0.36***	-0.07	0.32***
	(0.056)	(0.062)	(0.055)	(0.054)	(0.060)	(0.053)
LTV	-3.76***	-3.63***	-2.49***	-3.76***	-3.62***	-2.48***
	(0.155)	(0.152)	(0.166)	(0.153)	(0.149)	(0.166)
Origination Month FE	✓	✓	✓	✓	✓	✓
County-Year FE	✓	✓	✓	✓	✓	✓
Property & Loan Controls	✓	✓	✓	✓	✓	✓
Observation	4.3M	2.1M	0.9M	4.3M	2.1M	0.9M
R2	0.15	0.13	0.19	-	-	-
Underidentification test statistic				68.56	71.39	61.46
Underidentification test p-value				0.00	0.00	0.00
Weak identification test statistic				22.84	20.17	18.62

Table 7: Model Estimates

This table presents the estimates of model parameters. Panel A reports values of externally calibrated parameters. Panel B reports values of estimated parameters. The parentheses in Panel B show 95% confidence intervals, calculated from 200 bootstrapped moment samples, as we describe in Appendix H.2.1.

Panel A: Externally Calibrated Parameters

Description	Parameter	Value
Intertemporal elasticity of substitution	η	2
Wealth at time of home purchase	W_1	\$60,000
House price	P	\$200,000
Discount factor	β	0.96
	T	7
Maximum LTV parameter	ϕ	0.8

Panel B: Parameters Calibrated to the Data or through Moment Matching

Description	Parameter	Value
Appraisal Standard Deviation	$\sigma_1, ..., \sigma_{10}$	See Figure 7
Search cost	ζ	0.583
		(0.306, 1.016)
Appraisal Bias	b	0.086
		(0.086, 0.088)
Penalty rate on consumption	ψ	2.782
		(2.022, 3.794)
Marginal utility of next period consumption	u_2'	0.00243
		(0.00227, 0.00250)

Table 8: Automated Appraisals Counterfactuals

This table shows how different versions of automated appraisals influence market outcomes, counterfactually within our calibrated model. In Panel A, we show results assuming that automated appraisals simply remove human biases from appraisals. In Panel B, we show results assuming that human biases are removed, the appraisal mean is shifted within each σ-decile such that the probability of underappraisal is unchanged, and the variance of appraisals is halved, so underappraisals are on average smaller when they occur. In each panel, the first two columns show the percentage point changes in final loan-to-price ratios l_{final} and interest rates, relative to the baseline outcome, and the third column shows the net increase in mortgage failure probability. The last column shows the "compensating variation" in prices, relative to the status quo of human appraisers, needed to make the consumer indifferent to the shift to automated appraisers. Quantities in parentheses reflect 95% confidence intervals, calculated from 200 bootstrapped moment samples, as we describe in Appendix H.2.1.

Panel A: Removing Bias

	Loan-to-price Ratio (pp)	Interest Rate (pp)	Fail Probability (pp)	Required Compensating Price Change (%)
1st σ-bucket	-2.773	-0.052	10.54	-5.347
	(-3.044, -2.367)	(-0.057, -0.043)	(9.188, 12.398)	(-6.875, -3.914)
5th σ-bucket	-2.616	-0.049	11.964	-5.786
	(-2.827, -2.226)	(-0.053, -0.04)	(10.633, 13.905)	(-7.463, -4.216)
10th σ-bucket	-2.36	-0.044	13.561	-6.262
	(-2.533, -2.035)	(-0.048, -0.037)	(12.145, 15.44)	(-8.126, -4.536)

Panel B: Reducing Variance

	Loan-to-price Ratio (pp)	Interest Rate (pp)	Fail Probability (pp)	Required Compensating Price Change (%)
1st σ-bucket	0.166	0.003	-0.713	0.364
	(0.062, 0.241)	(0.001, 0.005)	(-0.752, -0.66)	(0.201, 0.595)
5th σ-bucket	0.105	0.002	-1.079	0.5
	(0.052, 0.229)	(0.001, 0.004)	(-1.119, -0.997)	(0.278, 0.81)
10th σ-bucket	0.109	0.002	-1.6	0.694
	(0.023, 0.197)	(0, 0.004)	(-1.684, -1.496)	(0.39, 1.086)

Appendix for Online Publication

A Measurement

A.1 f_c and g_c Functions

In order to estimate price dispersion, we need to model prices as a flexible function of characteristics. We do this using generalized additive models, which are a class of flexible nonparametric models; Wood (2017) describes the theory of GAMs. We use the mgcv package in R to implement the GAMs. We use this class of functions because, in our simulations, they provide a better fit to house prices than standard high-order polynomials.

We implement a two-stage regression using general additive model (GAM) on a county level. Instead of a high order polynomial, GAM implements cubic spline basis (or tensor product for multivariates) to fit the regressors. Therefore, to avoid overfitting, we first throw out counties with less than 400 observations. In order to estimate the GAM, there needs to be sufficient variation in characteristics; thus, we only keep counties with at least 10 unique values of each of the following characteristics: geographic information (latitude and longitude), year built, square footage, and transaction date. We also normalize the months, latitude, and longitude, building square feet, and year built. Furthermore, we winsorize geographic information, year built and building square feet.

We then estimate the following generalized additive model:

$$
\begin{aligned}
f_c\left(x_i, t\right) = {} & h_c^{f,latlong}\left(t, lat_i, long_i\right) + h_c^{f,sqft}\left(t, sqft_i\right) + \\
& h_c^{f,yrbuilt}\left(t, yrbuilt_i\right) + h_c^{f,bedrooms}\left(t, bedrooms_i\right) + h_c^{f,bathrooms}\left(t, bathrooms_i\right)
\end{aligned}
$$

The functions $h_c^{f,latlong}$, $h_c^{f,sqft}$, and $h_c^{f,yrbuilt}$ are tensor products of 5-dimensional cubic splines in their constituent components: hence, for example, the $h_c^{f,latlong}\left(t, lat_i, long_i\right)$ is a three-dimensional spline tensor product, with a total of $5^3 = 125$ degrees of freedom. To combat overfitting, the spline terms also includes a shrinkage penalty term on the second derivative of the spline functions, with the smoothing penalty determined through generalized cross-validation. The functions $h_c^{f,bedrooms}$ and $h_c^{f,bathrooms}$ interact dummies for a given house having 1, 2, 3 or more bedrooms and 1, 2, 3 or more bathrooms respectively with cubic spline basis in time.

The functional form for $g_c(x_i, t)$ in (2) is exactly analogous to $f_c(x_i, t)$:

$$g_c(x_i, t) = h_c^{g,latlong}(t, lat_i, long_i) + h_c^{g,sqft}(t, sqft_i) + h_c^{g,yrbuilt}(t, yrbuilt_i) + h_c^{g,bedrooms}(t, bedrooms_i) + h_c^{g,bathrooms}(t, bathrooms_i)$$

A.2 Repeat-Sales Estimation and Results

One possible concern regarding our analysis is that our measure of value uncertainty relies heavily on our hedonic model (1) for house prices. To alleviate this concern, in this appendix, we construct an alternative measure of value dispersion using a repeat-sales model. We estimate the following regression specification:

$$p_{it} = \eta_{kt} + \mu_i + \epsilon_{it} \tag{A1}$$

where i indexes properties, k indexes counties, and t indexes months. Equation (A1) is a repeat-sales model for house prices: log prices p_{it} are determined by county-month fixed effects η_{kt}, time-invariant house fixed effects μ_i, and a mean-zero error term ϵ_{it}. Specification (A1) thus models log house prices as following parallel trends, plus error terms: if house A sells for twice the price of house B in June of 2011, house A should sell for twice as much as house B in June of 2017, and any deviation from this is attributed to the error term ϵ_{it}.

There are two additional concerns with measuring idiosyncratic dispersion using a repeat-sales specification. First, the number of data points used to estimate each house fixed effect is very low; thus, the estimated residuals $\hat{\epsilon}_{it}^2$ will tend to be larger for houses which are sold more times, because the house fixed effect γ_i is estimated more precisely. Second, (A1) implicitly assumes that idiosyncratic price dispersion does not depend on the house holding period; a concern is that there is a idiosyncratic price dispersion behaves partially like a random walk, so the error terms may be systematically larger for houses that are sold less frequently.[29] To alleviate the concern that our estimates of $\hat{\epsilon}_{it}^2$ are mechanically driven by sale frequency and time-between-sales, we purge $\hat{\epsilon}_{it}^2$ of any variation which can be explained by tbs_i and $sales_i$. First, we filter to houses sold at most four times over the whole sample period, with

[29] Note that Giacoletti (2021) and Sagi (2021) show that a large component of idiosyncratic dispersion does not scale with holding period, for both residential and commercial real estate transactions.

estimated values of $\hat{\epsilon}^2_{it}$ below 0.25. We then run the following regression, separately for each county:

$$\hat{\epsilon}^2_{it} = h_k\left(sales_i, tbs_i\right) + \zeta_{it} \tag{A2}$$

$h_k\left(sales_i, tbs_i\right)$ interacts a vector of $sales_i$ dummies with a fifth-order polynomial in tbs_i. The residual $\hat{\zeta}_{it}$ from this regression can be interpreted as the component of the house's price variance which is not explainable by $sales_i$ and tbs_i. We then add back the mean of $\hat{\epsilon}^2_{it}$ within county k:

$$\hat{\epsilon}^2_{TBSadj,it} = \hat{\zeta}_{it} + E_k\left[\hat{\epsilon}^2_{it}\right] \tag{A3}$$

$\hat{\epsilon}^2_{TBSadj,it}$ can be interpreted as the baseline estimates, $\hat{\epsilon}^2_{it}$, nonparametrically purged of all variation which is explainable by a smooth function of $sales_i$ and tbs_i. We then project $\hat{\epsilon}^2_{TBSadj,it}$ onto house characteristics and time, as in (2) in the main text, and take the predicted values as our house-level measure of idiosyncratic price dispersion, which we will call $\hat{\sigma}^2_{RS,it}$.

In comparison to the hedonic model, the repeat-sales model in (A1) is able to capture observable and unobservable features of houses that have time-invariant effects on house prices. Moreover, house fixed effects allow us to capture time-invariant house quality components in a fully nonparametric way, alleviating concerns that the specific functional form we use in (1) is driving our results. A weakness of specification (A1) are that it is unable to capture any features of houses which have time-varying effects on house prices.

Figure A1 shows a binscatter of $\hat{\sigma}^2_{RS,it}$ against our baseline estimates $\hat{\sigma}^2_{it}$. There is a very strong positive relationship. The repeat-sales and hedonic methodologies for measuring house value uncertainty are econometrically quite different; the fact that they produce very correlated results at the house level suggests that both measurement strategies are picking up fundamental value uncertainty among properties, rather than simply reflecting misspecification in the model we use for house prices.

Next, we repeat our regression specifications utilizing $\hat{\sigma}^2_{RS,it}$ as our measure of house price dispersion. Table A9 shows the results; all of our baseline results continue to hold, using $\hat{\sigma}^2_{RS,it}$ as our measure of house price dispersion.

B Data Cleaning

Corelogic tax & deed data. We clean the datasets using a number of steps. First, we use only arms-length new construction sales or resales of single-family residences, which are not foreclosures, which have non-missing sale price, date, APN, and county FIPS code in the Corelogic deed data, and which have non-missing year built and square footage in the Corelogic tax data. We use only data from 2000 onwards, as we find that Corelogic's data quality is low prior to this date. Even after throwing out pre-2000 data, we find that some counties have very low total sales for early years, suggesting that some data is missing. To address this, we manually filter out some early county-years for which the total number of sales is low.

We also filter out "house flips", as well as instances where reported sale price seems anomalous. If a house is ever sold twice within a year, we drop all observations of the house. Most of these kinds of transactions appear to be either flips, which are known to be a peculiar segment of the real estate market (Bayer et al., 2020; Giacoletti and Westrupp, 2017), or duplication bugs in the data, where a single transaction is recorded twice or more. To filter for potentially anomalous prices, if we ever observe a property whose annualized appreciation or depreciation is above 50% for any given pair of sales, we drop all observations of the property. Finally, if a house is ever sold at a price which is more than 5 times higher or lower than the median house price in the same county-year, we drop all observations of the house from our dataset.

Our model of prices involves a fairly large number of parameters, so we filter to counties with a fairly large number of house sales in order to precisely estimate the model. Thus, we filter to counties with at least 1,000 house sales remaining, and with at least 10 sales per month on average, after applying the filtering steps described above.

Corelogic LLMA data. We filter to only purchase loans, excluding refinancing loans. As in the Corelogic Deed data, we calculate the loan-to-price ratio as the mortgage loan amount, divided by the house transaction price. We dropped observations with empty property zipcode, FICO score, initial interest rate, mortgage amount, origination date, sale price, and back-end ratio. We divide the market into conforming and non-conforming loans, using a flag provided by corelogic. We dropped all observations with balloon loans, and with loan to price ratio > 100. We kept observations with full documentation and fixed interest rates. We dropped observations with outliers. Specifically, we dropped all observations lower than the 1st percentile and higher than the 99th percentile with respect to loan to price and initial interest rate.

C Drivers of Idiosyncratic Price Dispersion

We discuss a number of factors and theoretical forces that may drive dispersion, and explain why these theories have similar implications for mortgage credit provision.

Information asymmetry. Lenders of secured loans must be concerned about adverse selection. This is especially the case in the consumer credit market, where houses and used cars, for example, have diverse characteristics, some of which are difficult to measure, and homeowners have better information about these characteristics (Kurlat and Stroebel, 2015; Stroebel, 2016). Houses with more hard-to-measure characteristics tend to have higher value uncertainty. Thus, lenders who lend against houses with higher value uncertainty may worry more about adverse selection because the owners have more information advantage about the house than the lenders.

Search frictions. The housing search literature has argued that house transaction prices are not determined in a fully competitive and frictionless market. Prices appear to depend not only on house characteristics: the transaction price of a house appears to be causally influenced by characteristics of the buyer and seller. Sellers who are more patient achieve higher sale prices, by setting higher list prices and keeping houses on the market for longer; this has been shown using instruments for seller patience, such as homeowners' equity position (Genesove and Mayer, 1997; Guren, 2018) and homeowners' nominal losses since purchase (Genesove and Mayer, 2001). Dispersion in different buyers' values for the same house may also drive house price dispersion: using data from Norwegian housing auctions, Anundsen et al. (2020) shows that the standard deviation of the ratio between buyers' bid prices and appraisal values is approximately 7.9%. Other factors, such as the experience of the realtor selling the house, also appear to affect house sale probabilities and prices (Gilbukh and Goldsmith-Pinkham, 2024).

Other factors. We also note that there are other possible housing market frictions which generate price dispersion. The literature has studied many different models, such as random search (Wheaton, 1990), directed search (Albrecht et al., 2016), and price posting (Guren, 2018).

We do not take a stance on this paper on the particular theoretical microfoundation of price dispersion, since it is not crucial for studying the effects of dispersion on credit provision. Price dispersion decreases credit provision by increasing lenders' expected losses upon foreclosure, and by making appraisals noisier and thus appraisal constraints more binding. Both effects occur regardless of the particular theoretical microfoundation of prices dispersion.

If we observed all characteristics of houses that market participants observed, and our functional forms for house prices were fully flexible, our measurement strategy would fully filter out the effects of house characteristics, capturing only price dispersion generated by housing market frictions. In practice, in addition to frictional price dispersion, our estimates are likely to be confounded by two main factors. First, our estimation cannot account for the effects of house characteristics unobserved in our data, but observed by market participants. Second, our functional forms in (1) may not be flexible enough to capture the true conditional expectation function; model misspecification will thus contribute to our estimates of price dispersion. Both of these effects serve as confounds we would like to filter out from our analysis, since if lenders use the correct price model with the full set of observables, frictional price dispersion should affect mortgage lending decisions, but not errors attributable to unobservables or model misspecification.

We believe these confounds are unlikely to drive our main results, for the following reasons. We observe a rich set of characteristics, which are essentially all the features that mortgage lenders observe for houses. A limitation of our data is that we only have time-invariant characteristics and do not observe renovations and time variation in house characteristics. However, Giacoletti (2021) uses data on remodeling expenditures for houses in California and finds quantitatively small effects on estimated price: accounting for renovations decreases the estimated standard deviation of returns by only around 2% of house prices.

D Implications of Controlling for House Prices

We briefly discuss the implications of controlling for house prices in our main empirical specifications. Conceptually, when we control for prices, β is identified by comparing, for example, zipcode A to zipcode B, which have similar average prices, but B has higher price dispersion. However, in most models, price dispersion affects the level of house prices: if two zipcodes have similar house quality, but one has higher price dispersion, average prices should be lower in the high-dispersion zipcode. Thus, in our example, if zipcode A and B have identical average prices, zipcode B should have *higher* average house quality than zipcode A.

Regressing LTP on price dispersion controlling for prices, and not controlling for house quality, makes sense in a model in which the distribution of house prices captures all features of houses that are relevant for lenders' decision problem. We construct a simple model which links lending decisions to the mean and variance of house prices in Appendix G.2. Since lenders do not directly interact with the house they lend against, in principle they should only care about house characteristics to the extent that they change the level or dispersion of house prices. Controlling for house prices alleviates the possibility that lenders may have a preference to have systematically higher or lower LTPs for high-priced houses. It is not necessary to control for house quality in addition to house prices, since any two houses with the same mean and variance of house prices are equivalent to lenders, regardless of the particular characteristics of the two houses.

A simpler reason why controlling for prices should not substantially matter for our results is that, when house prices are higher, *loan size* should increase, but it is not obvious whether *loan-to-price* ratios should increase or decrease. In most lending models, such as our model in Appendix G.2, the overall level of prices has no effect on loan-to-price ratios. To test that our results are not driven by this choice, in Table A10, we estimate our main specifications without controls for prices; all results are qualitatively and quantitatively similar to the main text.

E Economic Magnitude of the Effects on Loan Size

In this appendix section, we illustrate that a small percentage change in the down payment requirement can lead to material impacts on homeownership, considering the low level of annual savings by the marginal home buyers in the US. To do this, we assume that young people between age 25 to 35 with income levels in the bottom quartile among their peers in the same state are the marginal home buyers facing down payment constraints; and assume that they save 20% of their annual income.

We calculate the additional down payments required based on the estimated impact of price dispersion on LTP in column 6 of Table 5, house prices and price dispersions of all transacted houses recorded in the Corelogic Deeds data, and households incomes obtained from the 2015 US census microdata. In particular, for each transacted house, we calculate its "excess price dispersion" ($\Delta\sigma$) as its price dispersion relative to the bottom decile of house price dispersion in the US. We use the bottom decile house price dispersion in the US as a convenient proxy for the lowest price dispersion a house could achieve in the US. We then find the state average excess price dispersion and house prices and use them to calculate the additional down payment needed for households living in each state:

$$\Delta\text{Down Payment}_s = 1.3\% \times \frac{\Delta\sigma_s}{0.11} \times \text{House Price}_s \tag{A4}$$

where $\frac{\Delta\sigma_s}{0.11}$ is the state-average excess price dispersion expressed in terms of the number of standard deviations and 1.3% is the estimated impact of one standard deviation higher price dispersion on LTP.

Finally, we express the impact of price dispersion on loan size as the share of annual savings required to cover the additional down payments, $\Delta\text{Down Payment}_s$. We approximate the level of annual savings of marginal home buyers in each state using 20% of the annual income level of individuals between age 25 to 35 in a state.

In Figure A3, we plot $\Delta\text{Down Payment}_s$ divided by our approximated annual savings of the marginal home buyers for each state. As shown in this figure, the required additional down payment exceeds twice the annual savings in several states and are above 80% of annual savings in most states. This simple numerical example implies that the impact of price dispersion on loan size plausibly delays homeownership by a half to a full year.

F Robustness Checks

F.1 Bunching Below Conforming Loan Limits

We test whether the effect of price dispersion on mortgage LTP and cost menu is driven by home buyers lowering the loan-to-price ratio to be eligible for securitization with the participation of government-sponsored enterprises (GSEs). Specifically, conforming mortgages must be below the conforming loan limits, which vary across regions and time. Conforming loans are much easier to sell than non-conforming loans, also known as jumbo loans, because of the participation of GSEs. GSEs insure default risks of loans they purchase and securitize, providing subsidized credit to GSE mortgage borrowers.

We test if our main findings are robust to the sub-sample of house transactions with sale prices below local conforming loan limits. These house transactions are not subject to the concern about bunching below conforming loan limit as the transaction prices are already below the conforming loan limit. Table A4 reports the results. The results show that our main finding is not driven by home buyers' incentive to keep their loan amount below the conforming loan limit. Among houses with prices below the conforming loan limit, houses with higher price dispersion are financed with smaller loans given the same interest rates than houses with lower price dispersion. The result holds in both OLS and IV settings.

F.2 Lender Market Power

Our main results in Tables 3, 4, and 5 are unlikely to be driven by lender market power because our empirical analysis exploits within county-year variation. If lenders essentially compete at the county level, buyers within a given county-year likely face roughly the same degree of lender market power, so within-county-year variation in outcomes we observe is unlikely to be driven by lender market power.

However, a concern is that only a smaller number of lenders are willing to lend against houses with high value uncertainty, giving these lenders market power to extract rents from home buyers of such houses but not from other home buyers purchasing houses with lower value uncertainty. To address this concern, we first investigate whether the market segment for houses with high value uncertainty indeed has fewer lenders than the low-value uncertainty market segment, within the same county-year.

In Table A5, we examine the correlation between price dispersion and lender HHI (columns 1-4) as well as the number of lenders (columns 5-8). We divide each county-year into four market segments based on the price dispersion of the underlying properties in Corelogic Deeds transaction records and collapse the sample into *county-year-price dispersion bin* level. In this sample, we have four observations in each county-year that indicate the average price dispersion, the number of lenders, and the lender HHI index for each price dispersion bin in a given county-year. We drop county-year-price dispersion bins with less than 50 transactions. We find that there is not a statistically significant positive correlation between market power and house price dispersion within county-years. In fact, in some specifications, lender market power appears to be negatively correlated with house price dispersion: we observe more lenders lending against houses with higher price dispersion, yielding lower lender concentration in the higher price dispersion segments within a county-year. Therefore, although suggestive, the evidence is not consistent with the hypothesis that lenders are able to earn monopolist rents on buyers of high-dispersion houses but not on buyers of low-dispersion houses in the same county-year.

To further isolate the effect of collateral value uncertainty from the effect of lender market power, in Tables A6-A8, we conduct subsample analyses for houses located in zipcodes with more lenders willing to lend to versus houses located in zipcodes with fewer lenders willing to lend to. In column 1 and 4 of these tables, we construct zipcode lender HHI index and include it as a control. In Tables A6 and A8, we further include lender-county-year fixed effects.[30] The inclusion of lender-county-year fixed effects allows us to compare houses financed by the same lender in a given county-year.

The results suggest that lender market power indeed reduces credit supply at the extensive margin: *within* the same county-year, mortgage rejection rates are higher in zipcodes with

[30] In the HMDA and the Corelogic Deeds records, we observe lender identities. But we do not observe lender identity in the Corelogic LLMA dataset. Thus, we cannot control for lender-related information for the cost menu analysis in Table 4.

more concentrated mortgage supply, and loan sizes tend to be smaller in these zipcodes. However, the effects of price dispersion on loan rejection, loan sizes, and interest rates remain significant and have similar magnitudes in these tables after controlling for mortgage supply concentration.

We then classify zipcodes as high or low concentration based on their lender HHI indexes compared to the median value among all zipcodes within a county-year. Columns 2 and 5 report the subsample analysis results using transactions in low HHI zipcodes, and columns 3 and 6 report the subsample analysis results using transactions in high HHI zipcodes. In all these tables, we confirm that the effects of price dispersion on loan rejection, loan sizes, and interest rates are quantitatively similar across zipcodes with different lender concentration levels. The results suggest that the main findings of this paper — the effects of price dispersion on mortgage provision — are not driven by the correlation between price dispersion and lender market power within a county-year.

G Proofs and Supplementary Material for Section 4

G.1 Microfounding the Penalty Cost Parameter ψ

This appendix constructs a microfoundation for the "penalty cost" parameter ψ, which implies that increases in down payments caused by under-appraisals decrease consumption more than one-for-one. We do a simple calculation to illustrate that the penalty cost can be fairly large in reasonable models. Suppose an agent lives for T periods, and maximizes discounted CRRA utility over consumption:

$$\sum_{t=1}^{T} \beta^t \frac{c_t^{1-\eta} - 1}{1-\eta}$$

$$s.t. \; a_{t+1} + c_t = y_t + a_t \left(1 + r\right)$$

Income y_t is exogeneous and nonrandom. As is standard in the lifecycle literature, we set $\eta = 2$. We set $\beta = 0.95, r = \frac{1}{\beta} - 1$, so that the optimal solution without uncertainty involves consuming equal amounts in every time period. We set $T = 10$, so a time period can be thought of as representing a year, and consumers can be thought of as have 10 years to save for a home purchase at time T. We set $y_t = 10$ for each period.

We compare two cases. The first is an *anticipated* shock to income in period T, whose realization is known in period 1. The anticipated shock can be thought of as the homebuyer choosing a lower target loan size: since she plans to make a larger down payment, she can consumption-smooth for this in advance. The second is an *unanticipated* shock, whose realization is only known in period T. This can be thought of as the homebuyer targeting a large loan size and anticipating that under-appraisals may force her to borrow less than the target loan size. This kind of shock is more costly because the consumer can consumption-smooth the first kind of shock in expectation, but cannot condition her consumption on the under-appraisal. We will show that the second kind of shock decreases total utility more than the former.

For both cases, we suppose that $y_T = 10$ and $y_T = 0$ with equal probability, and $y_t = 10$ for all periods $t \neq T$. In the anticipated case, we assume y_T is known when the buyer makes consumption decisions in earlier periods. Thus, to solve this problem, we simply solve a zero-uncertainty finite-horizon dynamic program for the consumer for each value of y_T, and then take the average lifetime value at $t = 0$ from each case. In the unanticipated case, the consumer's value function in period $T - 1$ is the average of her value if $y_t = 10$ and if $y_t = 0$. The rest of the consumer's problem can be solved with standard backwards induction. We solve both cases using the standard endogenous gridpoint method for solving lifecycle problems.

We compare the consumer's lifetime value in both the anticipated and unanticipated income decrease cases to the baseline case where $y_t = 10$ for all time periods. In the anticipated case, lifetime value drops by 0.0361, whereas in the unanticipated case lifetime value drops by 0.050. Hence, under these parameter settings, an unanticipated shock is roughly 40% more costly, in utility terms, than an anticipated shock of the same magnitude, due to the inability to condition early-period consumption on the realization of the shock. Thus, unanticipated shocks to consumption can have much larger effects on utility than equally sized anticipated shocks. Our consumption penalty parameter ψ is a reduced-form way to capture this effect.

G.2 Microfounding the Mortgage Rate Menu

In this appendix, we construct a microfounded model showing how mortgage interest rates depend on targeted loan size and price dispersion. We assume mortgage rates arise from competition between profit-maximizing lenders. Suppose that, once a homebuyer has purchased the house with a mortgage, the buyer will default on the mortgage at rate δ. If the buyer defaults, the lender incurs a proportional cost Pc to foreclose the house, reflecting foreclosure discounts and other hassle costs of foreclosing. The foreclosure price is a function of the initial transaction price and a random component, ϵ_F, which has standard deviation σ_F that depends on idiosyncratic price dispersion. Thus, the final recovery value is as follows:

$$F = P\left(1 - c + \epsilon_F\right) \tag{A5}$$

Thus, for a non-recourse mortgage, lender's expected loss conditional on default is:[31]

$$Loss = E\left[P\left(l - \max\left[l, 1 - c + \epsilon_F\right]\right)\right] = PE\left[\max\left[0, l - \left(1 - c + \epsilon_F\right)\right]\right] \tag{A6}$$

Lender's expected loss is increasing in σ_F because the lender can recover at most l and bears the cost when the foreclosure price is less than l.[32] Thus, when the variance of the foreclosure price is larger, the lender's expected losses on loans is higher.

Now, suppose lenders have cost of funds ρ, and let r represent the mortgage interest rate. Lenders' profit if buyers do not default is $Pl\left(r - \rho\right)$. In a competitive equilibrium, the menu of interest rates and loan size must be set such that the lender will break even on any mortgage-rate pair:

$$Pl\left(1 - \delta\right)\left(r - \rho\right) = \delta PE\left[\max\left[0, l - \left(1 - c + \epsilon_F\right)\right]\right] \tag{A7}$$

The LHS of (A7) is lenders' expected profit, which is the product of mortgage size l, the repayment probability $(1 - \delta)$, and the mortgage spread $(r - \rho)$. The RHS is lenders' expected

[31] Mortgages are recourse in some states, but wage garnishment and other methods for collecting debt from buyers after the house has been sold are expensive, and buyers cannot be collected from if they file Chapter 7 bankruptcy.

[32] We assume that if the borrower defaults, it happens before Period 2. This assumption is reasonable because buyers are more likely to default in early stage when they have less equity in the house. If we relax this assumption, the loss function will be as follows, which will result in similar results: $Loss = E\left[P\left(l(1+\rho) - \max\left[l(1+r), 1 - c + \epsilon_F\right]\right)\right]$

losses conditional on default, multiplied by the default probability δ.

Expression (A7) defines a menu of (l, r) pairs available to buyers. As we increase idiosyncratic price variance, thus increasing the variance of prices upon foreclosure σ_F, the menu of (l, r) pairs shifts to be worse for the borrower. Formally, when ϵ_F is normally distributed, the RHS of (A7) is always increasing in σ_F.[33] Thus, holding l fixed, increasing σ_F must cause r to increase. This rationalizes our observations in Figure 5 and Table 4. Expression (18) in the main text can be thought of as a linear approximation to this menu.

G.2.1 Mortgage Rate Menu Calibration

We next do a simple calibration, to show that this microfoundation can also quantitatively rationalize the relationships between interest rates, loan size, and price dispersion observed in the data. Essentially, in the calibration, we will group the data into buckets with different default rates δ. We will estimate σ_F based on price dispersion in the data, and we will choose a foreclosure discount c to minimize the distance between the model and data interest rate menus. We will then show that the fitted model, optimizing over a single parameter, can fit the empirical relationships between loan size l, price dispersion σ_F, and interest rates r, simultaneously for many levels of default rates.

We restrict the sample to all portfolio loans. We first group the data into four FICO score bins, Excellent (800-850), Very Good (740-799), Good (670-739), and Fair (580-669), indexed by f. We split each FICO score bin into high- and low-dispersion counties, indexed by d, and also split loans into LTP bins, from 60-65, 65-70, up to 80. For each FICO score bucket f, dispersion case d, and LTP bin l, we estimate average residualized interest rates r_{fld} in our sample of loans. Since the level of r_{fld} is meaningless after residualization, we normalize by subtracting the mean rate $\bar{r}_f$ within each FICO bucket f:

$$\tilde{r}_{fld} = r_{fld} - \bar{r}_f \tag{A8}$$

Since we normalize within FICO buckets, we preserve the relationships between $\tilde{r}_{fld}$, loan size l, and price dispersion d within each FICO bucket. The residuals $\tilde{r}_{fld}$ are essentially the

[33] Note that the RHS of (A7) is equal to δ times the value of a European call option on $l - (1 - c + \epsilon_F)$ with strike 0; the value of such a call option is always increasing in volatility.

points in the interest rate menu of Figure 5, separate for each of the four FICO buckets.

Next, we describe how we simulate value of model-predicted interest rate menu points $\tilde{r}^{model}_{fld}(c)$, given the foreclosure discount c. We assume that ϵ_F is normally distributed, with mean 0 and variance σ_F. In each FICO score bin, we calculate a homogeneous value of δ as the average delinquency rate across all loans. To determine σ_F in the high- and low-dispersion areas, we calculate the average repeat-sales residual, as described in Appendix A.2, separately for high-dispersion and low-dispersion counties.[34] We find $\sigma_F = 0.0941$ for low-dispersion areas, and $\sigma_F = 0.131$ for high-dispersion counties. Given δ, σ_F, and loan size l, for any value of the foreclosure discount c, we can calculate the interest rate spread $r^{model}_{fld} - \rho$ using (A7):

$$r^{model}_{fld} - \rho = \frac{\delta E\left[\max\left[0, l - (1 - c + \epsilon_F)\right]\right]}{l(1-\delta)} \tag{A9}$$

where the expectation on the RHS of (A9) can be analytically calculated, since we assumed ϵ_F is normally distributed. We can then calculate the model counterpart of the interest rate residuals (A8), by subtracting the mean interest rate in each FICO bucket f:

$$\tilde{r}^{model}_{fld}(c) = r(l, \delta, c, \sigma_F) - \frac{\sum_l \sum_f r(l, \delta, c, \sigma_F)}{\sum_l \sum_f 1} = \\ \left(r(l, \delta, c, \sigma_F) - \rho\right) - \frac{\sum_l \sum_f r(l, \delta, c, \sigma_F) - \rho}{\sum_l \sum_f 1} \tag{A10}$$

Note that (A10) implies that $\tilde{r}^{model}_{fld}$ does not depend on the choice of ρ, so we set an arbitrary value of ρ in calculating $\tilde{r}^{model}_{fld}(c)$. We then choose a value of the foreclosure discount c through generalized method of moments, to minimizes the squared distance between the

[34] We use repeat-sales residuals to estimate σ_F, rather than the hedonic model residuals in the main text, because repeat-sales are closer to the thought experiment in the collateral recovery model. We are interested in, when a house forecloses, how variable its price is relative to its purchase price, which is captured in a repeat-sales specification. If a house has large errors in the hedonic model, but not the repeat-sales model – that is, a house has persistently high values relative to its characteristics – this does not affect the variability of the house price relative to loan value upon foreclosure, so this should not be included in ϵ_F.

data residuals $\tilde{r}_{fld}$, and the model residuals $\tilde{r}_{fld}^{model}$:

$$c^* = \arg\min_c \sum_l \sum_f \sum_d w_{fd} \left(\tilde{r}_{fld} - \tilde{r}_{fld}^{model} \right)^2$$

where, we set the weights w_{fld} equal to the inverse of the standard deviation of residuals $\tilde{r}_{fld}$ within each FICO and dispersion bucket; this is useful since, without weights, the errors in the low-FICO buckets would dominate the GMM objective function, since rates are higher and more variable when FICO scores are lower.

Our GMM estimate of the foreclosure discount c^* is 0.2018. This is within the range of foreclosure discounts estimate in the literature; for example, Pennington-Cross (2006) estimate a foreclosure discount of 22%, and Zhou et al. (2015) estimate discounts ranging from 11% to 26%.

Figure A4 illustrates the fit of the model. In the top two panels, we show the data and model rate residuals, $\tilde{r}_{fld}$ and $\tilde{r}_{fld}^{model}$, on the y-axis, against the LTP on the x-axis, separately for low-dispersion (top left) and high-dispersion (top right) areas. Different colors represent different credit score bins. In the data, the interest rate menu is steeper when FICO scores are lower; the model is able to quantitatively match this feature of the data, with some errors from the model-predicted interest rate menus being slightly too flat for low FICO bins. This shows that the collateral recovery model is able to quantitatively explain the relationship between interest rates and loan size.

To focus on the effect of price dispersion of credit, in the bottom panel of Figure A4, we show the difference in interest rates between high- and low-dispersion cases, for each FICO bucket and LTP; that is, each point in the bottom panel shows:

$$r_{lf,d=H} - r_{lf,d=L} \tag{A11}$$

This is the difference between interest rates in high-dispersion and low-dispersion areas. In other words, the solid green line in the bottom panel is equal to the difference between the solid green line in the top right panel (rates for high-dispersion areas in FICO bin 4) and the solid green line in the top left panel (rates for low-dispersion areas in FICO bin 4). In the data, (A11) is larger when FICO scores are lower: dispersion affects mortgage credit more when default rates are higher. The model lines are very close to the data lines in Figure A4, implying that the model produces a surprisingly good fit of the relationship between default rates, and the relationship of price dispersion with mortgage interest rates: we are able to match the average level of each of the lines, as well as the slope for the green line, representing the lowest FICO scores.

Thus, we have shown that the interrelationships between interest rate residuals, LTP, default rates, and price dispersion in the portfolio segment of our data are quantitatively consistent with a simple collateral recovery model, under realistic parameter settings. The simple model fits the data surprisingly well, given that we only optimize a single parameter, the foreclosure discount c, in the model fitting.

G.3 Appraiser Incentives

This appendix constructs a microfounded model of appraiser behavior, which rationalizes our assumptions on how appraisers bias appraisal prices in (19) of Subsection 4.1. Our model is essentially a special case of Calem et al. (2021). The model also shares some similarities with Conklin et al. (2020), but does not model competition between appraisers. Our model is simplified and disregards some stylized facts shown in the literature: for example, we rule out the possibility that house prices are renegotiated downwards when appraisals fall below sale prices, a phenomenon which is analyzed in Fout et al. (2021).

From (12), the max loan the borrower can take out is:

$$L_{max} = \phi \max(P, A)$$

Suppose that the house appraiser receives utility χL_{max} if the loan size is L_{max}; that is, the appraiser receives some side benefit χ, for every unit they can increase the borrower's max loan size by. This could capture, for example, possible repeat business incentives to produce high appraisals, relationships with lenders (Eriksen et al., 2019), and other such forces.

We also assume that appraisers have some convex cost of biasing appraisals. If the "true" raw appraisal price is A_{raw}, and the appraiser generates appraisal A, then the appraiser incurs a cost:

$$c\left(A, A_{raw}\right) = \gamma\left(A - A_{raw}\right)^2 \tag{A12}$$

This cost is a reduced-form way to capture the fact that it is more costly for appraisers to generate larger distortions in appraisal prices. The literature has documented that appraisers have a number of methods to shift appraisal prices, such as misreporting certain house attributes (Eriksen et al., 2024) and changing the weights on comparable sales used to calculate appraisals (Eriksen et al., 2019). Appraisers would have to misreport attributes or shift weights more to bias appraisals by larger amounts, which may be more costly to the appraiser in terms of legal and reputational risk, or psychological costs.

Appraisers thus solve:

$$\max_A U_{appr}\left(A\right) = \chi L_{max}\left(A\right) - \gamma\left(A - A_{raw}\right)^2 \tag{A13}$$

The optimization problem in (A13) has three distinct regions. First, if $A_{raw} > P$, then the appraiser cannot increase L_{max}; it is thus optimal to set $A = A_{raw}$.

Second, suppose A_{raw} is very low. Conjecture that the optimal A is below P, so that the first-order condition for optimality holds:

$$\chi\frac{\partial L_{max}}{\partial A} = 2\gamma\left(A - A_{raw}\right)$$

This gives $A - A_{raw} = \frac{\chi\phi}{2\gamma}$. Define $b \equiv \frac{\chi\phi}{2\gamma P}$. We then have:

$$A - A_{raw} = bP$$

Third, suppose that:

$$P\left(1 - b\right) \leq A_{raw} \leq P$$

In this range, we have that:

$$\frac{\partial U_{appr}}{\partial A} > 0 \; \forall A < P$$

Hence, it is optimal for the appraiser to set A=P.

We have thus shown that the appraiser's optimal appraisal A^* satisfies:

$$A^* = \begin{cases} A_{raw} + bP & A_{raw} \leq (1-b)\,P \\ P & (1-b)\,P < A_{raw} \leq P \\ A_{raw} & P < A_{raw} \end{cases}$$

which is exactly (19) in the main text.

G.4 Proof of Theorem 1

Conditional on the appraisal value a, the buyer can choose to proceed with the loan and purchase the property (continue), or renege on the offer and search for a new house and loan (renege). Let the value of each option, with loan size l and appraisal a, be respectively $V\,(a, l, continue)$ and $V\,(a, l, renege)$. The maximized value at any a and l is:

$$V\,(a,l) \equiv \max\left[V\,(a,l,continue)\,, V\,(a,l,renege)\right] \tag{A14}$$

We proceed to characterize $V\,(a, l, continue)$ and $V\,(a, l, renege)$.

G.4.1 Characterizing $V\,(a, l, continue)$

If the buyer proceeds with appraisal a, her utility is:

$$V\,(a,l,continue) = \frac{c_1^{1-\eta} - 1}{1-\eta} + \beta^T u_2' c_2 \tag{A15}$$

From (15) and (16) in the main text, we have:

$$c_1 = \underbrace{W_1 - P\,(1-l)}_{Targeted\ consumption} - \underbrace{\psi P \max\left(0, l - \phi a\right)}_{Appraisal\ shortfall} \tag{A16}$$

$$c_2 = -\left(1 + r\,(l)\right)^T P\left(l - \max\left[0, l - \phi a\right]\right) \tag{A17}$$

where, as we discussed in the main text, we have set $W_2 = 0$, since second-period wealth only linearly shifts buyers' utility and does not interact with any of the buyer's decisions. In words, (A16) states that the buyer's consumption in period 1 is equal to her targeted consumption $W_1 - P(1-l)$, minus an "appraisal shortfall" term $\max(0, l - \phi a)$. If $a < \frac{l}{\phi}$, then the buyer must decrease her borrowing from l to ϕa; this decreases her period-1 consumption by $l - \phi a$, multiplied by the price, and the penalty term $\psi > 1$. Since the final loan size l_{final} is smaller, this also decreases the amount that the buyer must pay back in period 2 by $\left(1 + r(l)\right)^T P \max[0, l - \phi a]$. Substituting (A16) and (A17) into (A15), we have:

$$\begin{aligned} V(a, l, continue) = & u_1\left(W_1 - P(1-l) - \psi P \max(0, l - \phi a)\right) \\ & - u_2' \beta^T \left(1 + r(l)\right)^T Pl \\ & + u_2' \beta^T \left(1 + r(l)\right)^T P \max[0, l - \phi a] \end{aligned} \tag{A18}$$

where, $u_1(c) \equiv \frac{c^{1-\eta} - 1}{1-\eta}$. Recall that, in (21), we defined:

$$\omega(a, l) \equiv u_1\left(W_1 - P(1-l) - \psi P \max(0, l - \phi a)\right) + u_2' \beta^T \left(1 + r(l)\right)^T P \max[0, l - \phi a]$$

Using this definition, we have:

$$V(a, l, continue) = -u_2' \beta^T \left(1 + r(l)\right)^T Pl + \omega(a, l) \tag{A19}$$

G.4.2 Characterizing $V(a, l, renege)$

If the buyer reneges, she receives:

$$V(a, l, renege) = -\beta^T u_2' \zeta P + E_a\left(V(a, l)\right) \tag{A20}$$

In words, she pays a cost ζP in period 2 consumption, which costs $-\beta^T u_2' \zeta P$ in utility terms. She then returns to the beginning of the game, and thus receives the expectation of $V(a, l)$ over uncertainty in a. Expanding $E_a\left(V(a, l)\right)$, we have:

$$E_a\left(V(a, l)\right) = \int_0^\infty \max\left(V(a, l, continue), V(a, l, renege)\right) dF_a(a) \tag{A21}$$

Now, note that $V(a, l, renege)$, is independent of a, whereas from (A18), $V(a, l, continue)$ is increasing in a. Thus, there is some cutoff value $\bar{a}(l)$, such that continuing is optimal for all $a > \bar{a}(l)$. At the boundary $\bar{a}(l)$, continuing and reneging have equal value:

$$V(\bar{a}, l, renege) = V(\bar{a}, l, continue) \tag{A22}$$

Substituting for $V(\bar{a}, l, continue)$ using (A19), we have:

$$V(\bar{a}, l, renege) = -\beta^T \left(1 + r(l)\right)^T u_2' Pl + \omega(\bar{a}, l)$$

Substituting into (A21), we have:

$$\begin{aligned} &E_a\left(V(a, l)\right) = \\ &\qquad \int_0^\infty \max\Big(-\beta^T \left(1 + r(l)\right)^T u_2' Pl + \omega(a, l), -\beta^T \left(1 + r(l)\right)^T u_2' Pl + \omega(\bar{a}, l)\Big) dF_a(a) \end{aligned}$$

$$E_a\left(V(a, l)\right) = -\beta^T \left(1 + r(l)\right)^T u_2' Pl + \int_0^\infty \max\left(\omega(a, l), \omega(\bar{a}, l)\right) dF_a(a) \tag{A23}$$

Substituting into (A20), we have:

$$\begin{aligned} &V(a, l, renege) = \\ &\qquad - \beta^T u_2' \zeta P - \beta^T \left(1 + r(l)\right)^T u_2' Pl + \int_0^\infty \max\left(\omega(a, l), \omega(\bar{a}, l)\right) dF_a(a) \end{aligned} \tag{A24}$$

G.4.3 Solving For $\bar{a}$

Having characterized $V(a, l, renege)$ and $V(a, l, continue)$, we now solve for $\bar{a}$. Plugging in expressions for $V(\bar{a}, l, renege)$ and $V(\bar{a}, l, continue)$ into (A22), we have:

$$\begin{aligned} &- \beta^T \left(1 + r(l)\right)^T u_2' Pl + \omega(\bar{a}, l) = \\ &\qquad - \beta^T u_2' \zeta P - \beta^T \left(1 + r(l)\right)^T u_2' Pl + \int_0^\infty \max\left(\omega(a, l), \omega(\bar{a}, l)\right) dF_a(a) \end{aligned}$$

Rearranging, and deleting the shared term $\beta^T \left(1 + r(l)\right)^T u_2' Pl$, we have:

$$\omega(\bar{a}, l) = -\beta^T u_2' \zeta P + \int_0^\infty \max\left(\omega(a, l), \omega(\bar{a}, l)\right) dF_a(a) \tag{A25}$$

This is (20) of Theorem 1. Equation (A25) characterizes $\bar{a}(l)$. In words, the LHS of (A25) is the period-1 utility from continuing with the appraisal $\bar{a}$, suffering the cost from under-appraising. The RHS is the expected value from reneging, which is the utility cost $-\beta^T u_2' \zeta P$, plus the expected period-1 utility from drawing a new appraisal. At $\bar{a}$, these must be equal.

We can rearrange (A25) to:

$$\int_{a > \bar{a}} \left(\omega(a, l) - \omega(\bar{a}, l)\right) dF_a(a) = \beta^T u_2' \zeta P \tag{A26}$$

Since ω is increasing in a, the LHS of (A26) is strictly decreasing in $\bar{a}$, hence for any parameters, there is at most one value of $\bar{a}$ which solves (A26). Note also that (A26) shows that the optimal $\bar{a}$ must satisfy:

$$\bar{a} < \frac{l}{\phi}$$

that is, the optimal cutoff $\bar{a}$ must be low enough that it constrains the amount that can be borrowed. To see this, note that from (21), we have:

$$\omega(a, l) = u_1\left(W_1 - P(1 - l)\right) \;\; \forall a > \frac{l}{\phi}$$

That is, when $a > \frac{l}{\phi}$, so the appraisal is high enough that it does not constrain borrowing, then $\omega(a, l)$ is constant in a. As a result,

$$\int_{a > \bar{a}} \left(\omega(a, l) - \omega(\bar{a}, l)\right) dF_a(a) = 0 \; \forall \bar{a} \geq \frac{l}{\phi}$$

Hence, the LHS of (A26) is 0 for all $\bar{a} > \frac{l}{\phi}$; the RHS is positive, so it can never be optimal to set $\bar{a} > \frac{l}{\phi}$.

G.4.4 Optimal Loan Choice

Repeating (A23), we have that, given the optimal appraisal cutoff $\bar{a}(l)$, the expected value attained by the buyer, in expectation over uncertainty in a, is:

$$E\left(V\left(\bar{a}(l), l\right)\right) = -\beta^T \left(1 + r(l)\right)^T u_2' Pl + \int_0^\infty \max\left(\omega(a, l), \omega\left(\bar{a}(l), l\right)\right) dF_a(a) \quad \text{(A27)}$$

The buyer picks l to maximize (A27); this is (22).

G.5 Comparative Statics: Optimal Loan Choice

To do comparative statics, we will apply the envelope theorem to the optimization framing of the buyer's choice problem. Define:

$$\Gamma(l) \equiv E\left(V\left(\bar{a}(l), l\right)\right)$$

We can write Γ as:

$$\Gamma(l) = \max_{\bar{a}} \left[\int_{\bar{a}}^\infty \left[-\beta^T \left(1 + r(l)\right)^T u_2' Pl + \omega(a, l)\right] dF_a(a) + F_a(\bar{a}) \left[\Gamma(l) - P\beta^T u_2' \zeta\right]\right] \quad \text{(A28)}$$

In words, the buyer receives $-\beta^T \left(1 + r(l)\right)^T u_2' Pl + \omega(a, l)$ in the range $[\bar{a}, \infty]$ where the buyer continues, and $\Gamma(l) - P\beta^T u_2' \zeta$ in the range $[0, \bar{a}]$ where she reneges. In this framing, since $\bar{a}$ is chosen optimally given any l, we have:

$$\frac{\partial}{\partial \bar{a}} \max_{\bar{a}} \left[\int_{\bar{a}}^\infty -\beta^T \left(1 + r(l)\right)^T u_2' Pl + \omega(a, l)\, dF_a(a) + F_a(\bar{a}) \left[\Gamma(l) - P\beta^T u_2' \zeta\right]\right] = 0$$

Hence, the envelope theorem applies; we have:

$$\frac{d\Gamma}{dl} = \frac{\partial}{\partial l}\int_{\bar{a}^*}^{\infty} -\beta^T \left(1 + r\left(l\right)\right)^T u_2' P l + \omega\left(a, l\right) dF_a\left(a\right) + F_a\left(\bar{a}^*\right)\left[\Gamma\left(l\right) - P\beta^T u_2'\zeta\right]$$

Now, we can write $\Gamma\left(l\right)$ substituting for $\omega\left(a, l\right)$ using (21), to get:

$$\begin{aligned}\Gamma\left(l\right) = &\max_{\bar{a}}\int_{\bar{a}}^{\infty}\left[-\beta^T\left(1 + r\left(l\right)\right)^T u_2' P l\right.\\ &+ \beta^T\left(1 + r\left(l\right)\right)^T u_2' P \max\left[0, l - \phi a\right]\\ &\left.+ u_1\left(W_1 - P\left(1 - l\right) - \psi P \max\left(0, l - \phi a\right)\right)\right] dF_a\left(a\right)\\ &+ F_a\left(\bar{a}\right)\left[\Gamma\left(l\right) - P\beta^T u_2'\zeta\right]\end{aligned} \tag{A29}$$

Now, note that:

$$l - \max\left[0, l - \phi a\right] = \min\left[l, \phi a\right]$$

Hence, we can write (A29) as:

$$\begin{aligned}\Gamma\left(l\right) = &\max_{\bar{a}}\int_{\bar{a}}^{\infty}\left[-\beta^T\left(1 + r\left(l\right)\right)^T u_2' P \min\left(l, \phi a\right)\right.\\ &\left.+ u_1\left(W_1 - P\left(1 - l\right) - \psi P \max\left(0, l - \phi a\right)\right)\right] dF_a\left(a\right)\\ &+ F_a\left(\bar{a}\right)\left[\Gamma\left(l\right) - P\beta^T u_2'\zeta\right]\end{aligned} \tag{A30}$$

Differentiating with respect to l, we have:

$$\begin{aligned}\frac{d\Gamma}{dl} = &\frac{\partial}{\partial l}\left[\int_{\bar{a}}^{\infty} -\beta^T\left(1 + r\left(l\right)\right)^T u_2' P \min\left(l, \phi a\right)\right.\\ &\left.+ u_1\left(W_1 - P\left(1 - l\right) - \psi P \max\left(0, l - \phi a\right)\right) dF_a\left(a\right)\right] + F_a\left(\bar{a}^*\right)\frac{d\Gamma}{dl}\end{aligned} \tag{A31}$$

$$\begin{aligned}\frac{d\Gamma}{dl}\left(1 - F_a\left(\bar{a}^*\right)\right) = &\frac{\partial}{\partial l}\left[\int_{\bar{a}}^{\infty} -\beta^T\left(1 + r\left(l\right)\right)^T u_2' P \min\left(l, \phi a\right)\right.\\ &\left.+ u_1\left(W_1 - P\left(1 - l\right) - \psi P \max\left(0, l - \phi a\right)\right) dF_a\left(a\right)\right]\end{aligned} \tag{A32}$$

Now, we can separately analyze the RHS, in the under-appraisal region $a \in \left[\bar{a}, \frac{l}{\phi}\right]$ and the over-appraisal region $a \in \left[\frac{l}{\phi}, \infty\right]$. In the over-appraisal region, we have $\min(l, \phi a) = l$ and $\max(0, l - \phi a) = 0$, hence:

$$\frac{\partial}{\partial l}\int_{\frac{l}{\phi}}^{\infty} -\beta^T \left(1 + r(l)\right)^T u_2' P \min(l, \phi a) + u_1 \left(W_1 - P(1-l) - \psi P \max(0, l - \phi a)\right) dF_a(a) =$$

$$\frac{\partial}{\partial l}\int_{\frac{l}{\phi}}^{\infty} -\beta^T \left(1 + r(l)\right)^T u_2' P l + u_1 \left(W_1 - P(1-l)\right) dF_a(a) =$$

$$\left(1 - F_a\left(\frac{l}{\phi}\right)\right)\left(\underbrace{-T\beta^T \left(1 + r(l)\right)^{T-1} r'(l)\, u_2' P l}_{Rate\ Change} + \underbrace{P u_1' \left(W_1 - P(1-l)\right) - P\beta^T \left(1 + r(l)\right)^T u_2'}_{Consumption\ Smoothing}\right)$$

$$\underbrace{-\left[-\beta^T \left(1 + r(l)\right)^T u_2' P l + u_1 \left(W_1 - P(1-l)\right)\right] f\left(\frac{l}{\phi}\right)}_{Nuisance\ Term} \quad \text{(A33)}$$

The "rate change" term in (A33) represents the increase in interest payments in period 2 from increasing $r(l)$. The "consumption smoothing" term represents gains from more effectively smoothing consumption over the two periods. The intuition is that, if the house over-appraises, targeting a larger loan allows the buyer to borrow more, smoothing consumption, and gaining on the margin the gap between period-1 and period-2 marginal utilities. The "nuisance term" will cancel once we consider the under-appraisal region.

In the underappraisal region, we have $\min(l, \phi a) = \phi a$ and $\max(0, l - \phi a) = l - \phi a$, hence:

$$\frac{\partial}{\partial l}\int_{\bar{a}}^{\frac{l}{\phi}} -\beta^T \left(1+r\left(l\right)\right)^T u_2' P \min\left(l, \phi a\right) + u_1 \left(W_1 - P\left(1-l\right) - \psi P \max\left(0, l-\phi a\right)\right) dF_a\left(a\right) =$$
$$\frac{\partial}{\partial l}\int_{\bar{a}}^{\frac{l}{\phi}} -\beta^T \left(1+r\left(l\right)\right)^T u_2' P\phi a + u_1 \left(W_1 - P\left(1-l\right) - \psi P\left(l-\phi a\right)\right) dF_a\left(a\right) =$$
$$\int_{\bar{a}}^{\frac{l}{\phi}} \underbrace{-T\beta^T \left(1+r\left(l\right)\right)^{T-1} r'\left(l\right) u_2' P\phi a}_{Rate\ Change} + \underbrace{\left(1-\psi\right) P u_1' \left(W_1 - P\left(1-l+\psi\left(l-\phi a\right)\right)\right)}_{Under-Appraisal\ \ Penalty} dF_a\left(a\right)$$
$$+ \underbrace{\left[-\beta^T \left(1+r\left(l\right)\right)^T u_2' P\phi\left(\frac{l}{\phi}\right) + u_1\left(W_1 - P\left(1-l\right) - \psi P\left(l-\phi\left(\frac{l}{\phi}\right)\right)\right)\right] f\left(\frac{l}{\phi}\right)}_{Nuisance\ \ Term} \tag{A34}$$

The "rate increase" term is analogous to (A33). The intuition behind the "under-appraisal penalty" term is that, if the house eventually under-appraises, targeting a larger loan does not increase the eventual borrowing amount, but increases the size of any under-appraisal, causing the buyer to have to pay a penalty $\psi - 1 > 0$ of the incremental loan amount. The "nuisance term" simply cancels with the corresponding term from (A33) once we add the two components.

Combining (A33) and (A34), we have:

$$\frac{\partial}{\partial l}\int_{\bar{a}}^{\infty} u_1\left(W_1 - P\left(1-l\right) - \psi P \max\left(0, l-\phi a\right)\right) dF_a\left(a\right) =$$
$$\left(1-F_a\left(\bar{a}\right)\right)\left(-T\beta^T\left(1+r\left(l\right)\right)^{T-1} r'\left(l\right) u_2' P l\right) +$$
$$\left(1-F_a\left(\frac{l}{\phi}\right)\right) P\left(u_1'\left(W_1 - P\left(1-l\right)\right) - \beta^T\left(1+r\left(l\right)\right)^T u_2'\right) -$$
$$\int_{\bar{a}}^{\frac{l}{\phi}} \left(1-\psi\right) P u_1'\left(W_1 - P\left(1-l+\psi\left(l-\phi a\right)\right)\right) dF_a\left(a\right) \tag{A35}$$

Finally, combining (A35) with (A32), we have:

$$\frac{d\Gamma}{dl} = \left(-T\beta^T\left(1+r\left(l\right)\right)^{T-1}r'\left(l\right)u_2'Pl\right) + \\ \frac{1}{\left(1-F_a\left(\bar{a}^*\right)\right)}\left[\left(1-F_a\left(\frac{l}{\phi}\right)\right)P\left(u_1'\left(W_1-P\left(1-l\right)\right)-\beta^T\left(1+r\left(l\right)\right)^T u_2'\right) - \\ \int_{\bar{a}}^{\frac{l}{\phi}}\left(1-\psi\right)Pu_1'\left(W_1-P\left(1-l+\psi\left(l-\phi a\right)\right)\right)dF_a\left(a\right)\right]$$

Setting $\frac{d\Gamma}{dl}$ to 0 and rearranging, we can write the FOC for optimal loan choice as:

$$\underbrace{\left(1-F_a\left(\frac{l}{\phi}\right)\right)P\left(u_1'\left(W_1-P\left(1-l\right)\right)-\beta^T\left(1+r\left(l\right)\right)^T u_2'\right)}_{Consumption\ Smoothing} = \\ \underbrace{\left(1-F_a\left(\bar{a}^*\right)\right)\left(T\beta^T\left(1+r\left(l\right)\right)^{T-1}r'\left(l\right)u_2'Pl\right)}_{Rate\ Change} + \\ \underbrace{\int_{\bar{a}}^{\frac{l}{\phi}}\left(1-\psi\right)Pu_1'\left(W_1-P\left(1-l+\psi\left(l-\phi a\right)\right)\right)dF_a\left(a\right)}_{Under-Appraisal\ Penalty} \tag{A36}$$

The LHS of (A36) captures the effect of increasing loan size on consumption smoothing. If the house eventually appraises successfully, increasing targeted loan size by a dollar moves consumption from period 2, where marginal utility is lower, to period 1, where it is higher. The RHS captures the two costs of increasing l: first, the interest rate paid increases; second, conditional on under-appraisal, increasing l does not change the final loan size, but increases the consumption penalty from under-appraisal, since under-appraisals are larger. Hence, at the optimal choice of l, the LHS is positive: the buyer would prefer to increase loan size slightly, to shift consumption from period 2 to period 1, but is deterred from doing so by the rate change and under-appraisal penalty effects.

H Proofs and Supplementary Material for Section 5

H.1 Sensitivity Analysis

Our calibration takes strong stances on a number of parameters; a natural question is how sensitive our quantitative conclusions are to these input values. To address this, we re-calibrate our model varying various inputs, and show how results change. We consider sensitivity to three parameters: η, the coefficient of relative risk aversion in the first period; T, the assumed duration of the mortgage; and β, the annual discount rate. We test the results from setting η equal to 1.5 or 3; from T equal to 5 or 10; and from β equal to 0.94 or 0.98. In each case, we re-estimate the model, rerunning the moment matching procedure for the baseline moments, as well as all counterfactuals.

Table A11 shows how these alternative settings influence our parameter estimates. Other than the bias parameter b, the precise parameter estimates are somewhat sensitive to the calibrated values. However, Table A12 then shows how counterfactual outcomes vary across these alternative settings: counterfactual outcomes in fact vary surprisingly little across the different parameter settings. The first column shows the "compensating variation" for consumers in the 5th σ-decile, as in Subsection 6.1. The second and third columns show the change in mortgage failure probability, and compensating variation in prices, under the version of the automated appraiser counterfactual in Panel A of Table 8, where we assume automated appraisals remove human biases, but do not change appraisal variance. The fourth and fifth columns show the change in failure probability, and compensating variation in prices, under the counterfactual in Panel B of Table 8, in which automated appraisals are calibrated to maintain the same failure probability, but reduce the variance by half. The fail probability changes are very stable across specifications; the "compensating variation" numbers vary, but by less than a factor of 2, across all specifications we have tried.

We interpret these results as showing that, while the values of our estimated parameters vary somewhat with calibrated parameters, the estimated counterfactual quantities – fail probability changes and "compensating variation" amounts – are in fact linked to our input moments in a manner which is not very sensitive to the choices of calibrated model parameters.

H.2 Analytical Expressions for Variances of Empirical Moments

In order to derive our inverse-variance moment error weights for our GMM parameter estimates, as well as to resample moments for our parametric bootstrap, we require estimates of the sample variances of each of the moments we use as inputs to our GMM procedure; we derive analytical expressions for these variances here.

Failure probabilities. Transaction failures are binary outcomes, so we think of the number of transaction failures within each σ-decile as binomial, where we estimate the mean to be $\hat{p}_i$. A consistent estimator of the variance of $\hat{p}_i$ is thus:

$$\sqrt{\frac{\hat{p}_i\left(1-\hat{p}_i\right)}{n}}$$

Appraisal deviations. Calculating the standard error of appraisal deviations is somewhat more involved, and we construct an approximate upper bound for the variance. We can write the appraisal deviation, (25), as:

$$Appr\hat{D}ev_i = p_i^{under} E\left[1 - a_i \mid a_i < 1\right] \tag{A37}$$

Expression (A37) is a function of the empirical CDF of appraisals, $\hat{F}_i\left(a\right)$, as:

$$Appr\hat{D}ev_i = \int_0^1 \left(1-a\right) d\hat{F}_i\left(a\right) \tag{A38}$$

We can use the "layer cake" representation of (A38) to facilitate standard error calculation. Note that:

$$\frac{d}{da}\left(1-a\right)\hat{F}_i\left(a\right) = -\hat{F}_i\left(a\right) + \left(1-a\right)\hat{f}_i\left(a\right)$$

Hence, by the fundamental theorem of calculus, we have:

$$0 = \left(1-a\right)\hat{F}_i\left(a\right)|_0^1 = \int_0^1 \left(1-a\right) d\hat{F}_i\left(a\right) - \int_0^1 \hat{F}_i\left(a\right) da$$

where we used that appraisals cannot be negative, so $\hat{F}_i(0) = 0$. This implies that:

$$\hat{ApprDev_i} = \int_{-\infty}^{1} (1-a)\, d\hat{F}_i(a) = \int_{0}^{1} \hat{F}_i(a)\, da \tag{A39}$$

The RHS of expression (A39) facilitates calculating the standard error of $\hat{F}_i(a)$, since the RHS is more directly a function of the empirical CDF $\hat{F}_i(a)$. Approximating (A39) as a Riemann sum, we have:

$$\hat{ApprDev_i} \approx \sum_{a<1} \hat{F}_i(a)\, \Delta_a \tag{A40}$$

for some grid of points a and interval lengths Δ_a. Now, (A40) expresses the estimator $\hat{ApprDev_i}$ as a sum of the empirical CDF $\hat{F}_i(a)$ on a grid of points a. This allows us to upper-bound the standard deviation of $\hat{ApprDev_i}$, as follows. For any given a, $\hat{F}_i(a)$ is a Bernoulli random variable, with standard deviation asymptotically equal to:

$$SD\left(\hat{F}_i(a)\right) = \sqrt{\frac{\hat{F}_i(a)\left(1-\hat{F}_i(a)\right)}{n}} \tag{A41}$$

Now, since (A40) is a sum, its standard deviation is maximized when all $\hat{F}_i(a)$ variables, for different values of a, are perfectly correlated. We will thus construct an upper bound of the standard deviation of $\hat{ApprDev_i}$ by assuming perfect correlation between $\hat{F}_i(a)$ at different values of a. While the correlation is obviously not perfect in practice, since $\hat{ApprDev_i}$ is based mostly on tail observations of a, the correlations between $\hat{F}_i(a), \hat{F}_i(\tilde{a})$ will all be positive, and will be fairly high when a and $\tilde{a}$ are close to each other.

Under the perfect-correlation assumption, we have that:

$$SD\left(\sum_{a<1} \hat{F}_i(a)\, \Delta_a\right) = \sum_{a<1} SD\left(\hat{F}_i(a)\right) \Delta_a \tag{A42}$$

Combining expressions (A41) and (A42), an estimator for the standard deviation of $\hat{ApprDev_i}$

is:

$$SD\left(Appr\hat{Dev}_i\right) \approx \sum_{a<1} \sqrt{\frac{\hat{F}_i\left(a\right)\left(1-\hat{F}_i\left(a\right)\right)}{n}}\Delta_a \tag{A43}$$

In words, (A43) is just the integral of the Bernoulli standard deviation formula, over a. We thus estimate $SD\left(Appr\hat{Dev}_i\right)$ by evaluating the RHS of (A43) on a fine grid of a values.

Appraisal variance, $\hat{\sigma}_{a,i}$. From (23) in the main text, we use the sample standard deviation formula from over-appraisals data to estimate σ for each bin. We can then estimate the standard deviation of the sample variance, in general, as:

$$\sqrt{\frac{1}{n}\left(\mu_4 - \frac{n-3}{n-1}\sigma^4\right)}$$

where μ_4 is the 4th central moment of appraisals, which we calculate in each σ-decile simply using the unbiased right-tail data, as:

$$\mu_4 = \sqrt{E\left[\left(a_i-1\right)^4 \mid a_i > 1\right]}$$

Then, using the delta method, the standard deviation of $\hat{\sigma}_{a,i}$ is asymptotically:

$$\frac{1}{2\sigma}\sqrt{\frac{1}{n}\left(\mu_4 - \frac{n-3}{n-1}\sigma^4\right)}$$

l_{final}-to-σ regression coefficient. For the variance of this moment, we use simply the regression standard error, from Column 3 of Table 5.

H.2.1 Parametric Bootstrap

Using the standard deviations of our moments, we generate 200 bootstrap samples of the failure probabilities, the appraisal deviations $Appr\hat{Dev}_i$, the loan size-regression coefficient moment, as well as $\hat{\sigma}_{a,i}$ values; in each sample, we draw each parameter independently from a normal distribution with mean equal to the full-sample value of the moment, and standard deviation equal to the expressions derived above. We thus assume independence and normality of moments in this procedure. Normality should hold asymptotically as N is large. Since the moments for different σ-bins come from different counties, the assumption of independence across σ-deciles is arguably justified. The assumption of independence within σ-deciles is stronger; however, estimating correlations of these moments empirically is difficult since, for example, the fail probability and appraisal deviation moments come from different datasets, which we could not find an internally consistent way to resample from. Note also that the manual weights we impose on the different moment errors in our estimation process play no role in our bootstrap resampling procedure.

I Appendix Figures and Tables

Figure A1. Repeat-Sales Estimates and Hedonic Estimates

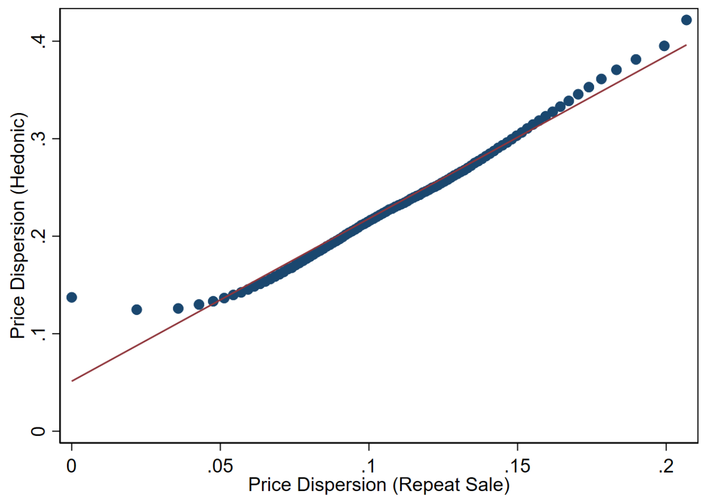

Notes: This figure compares our repeat-sale estimates and our hedonic estimates of price dispersion. The points are results from a binned scatterplot, where each data point is a single property. The x-axis shows repeat-sale estimates of price dispersion, described in Appendix A.2, and the y-axis shows our hedonic-regression estimates used in the main analysis. The sample includes property-level observations from 2000 to 2020 from the Corelogic Deeds.

Figure A2. County Level House Price Dispersion and LTP

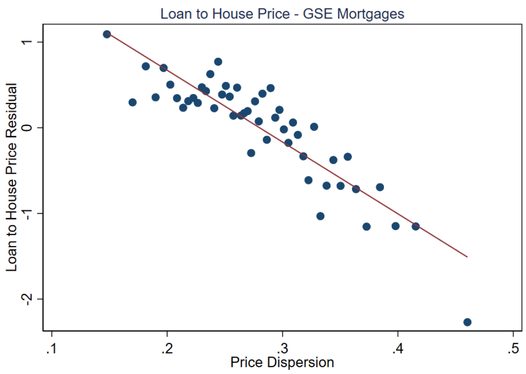

(a) GSE

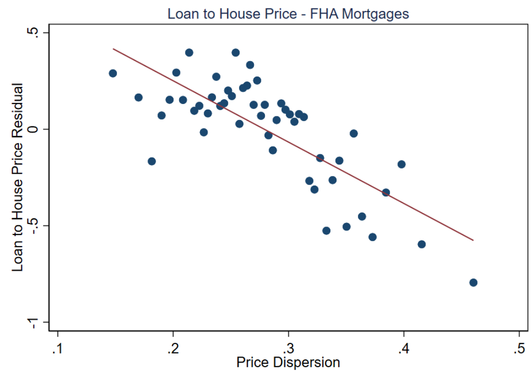

(b) FHA

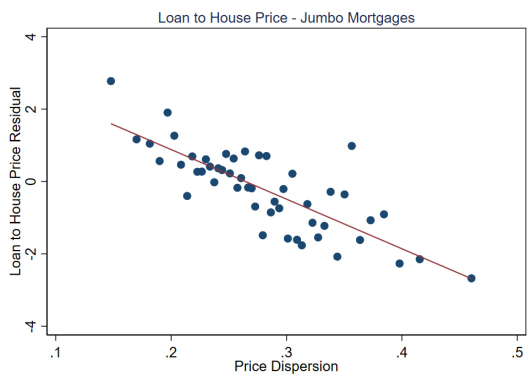

(c) Jumbo

Notes: This figure shows the correlation between county level house price dispersion and residualized county average LTP. Panels a-c plot GSE loans, FHA loans, and jumbo loans, respectively. The sample includes annual county observations from 2000 to 2020. County house price dispersion is estimated using Corelogic Deeds records. Mortgage data are from Corelogic LLMA.

Figure A3. Additional Down Payment to Annual Savings Ratio

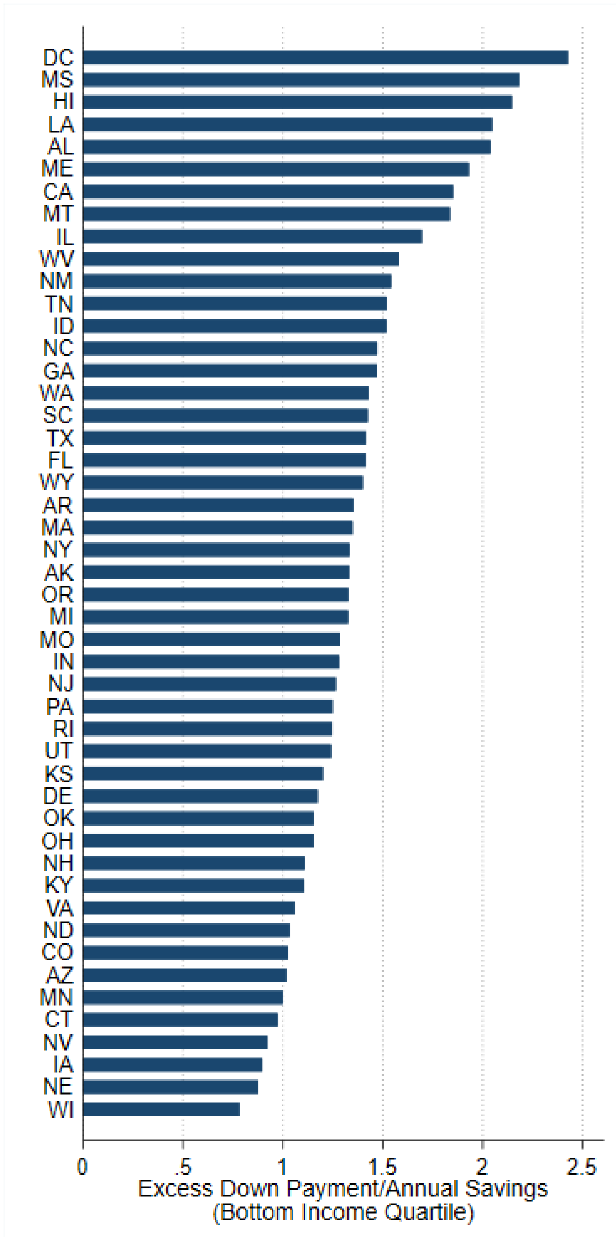

Notes: This figure presents additional down payment to annual savings ratio by state. To calculate additional down payment associated with house price dispersion, we use the bottom decile house price dispersion in the US as a benchmark for the lowest price dispersion a house could achieve in the US. The calculation of additional down payment is based on price dispersion of transacted houses in each state relative to the benchmark and the estimated coefficient in column 6 of Table 5, as described in Appendix E. We assume that young people between age 25 to 30 save 20% of their annual income levels reported in the Census. We then take the average annual saving level of young people with income levels in the bottom quartile among their peers in the same state, obtained from 2015 U.S. Census Microdata.

Figure A4. Rate Menu Model Fit

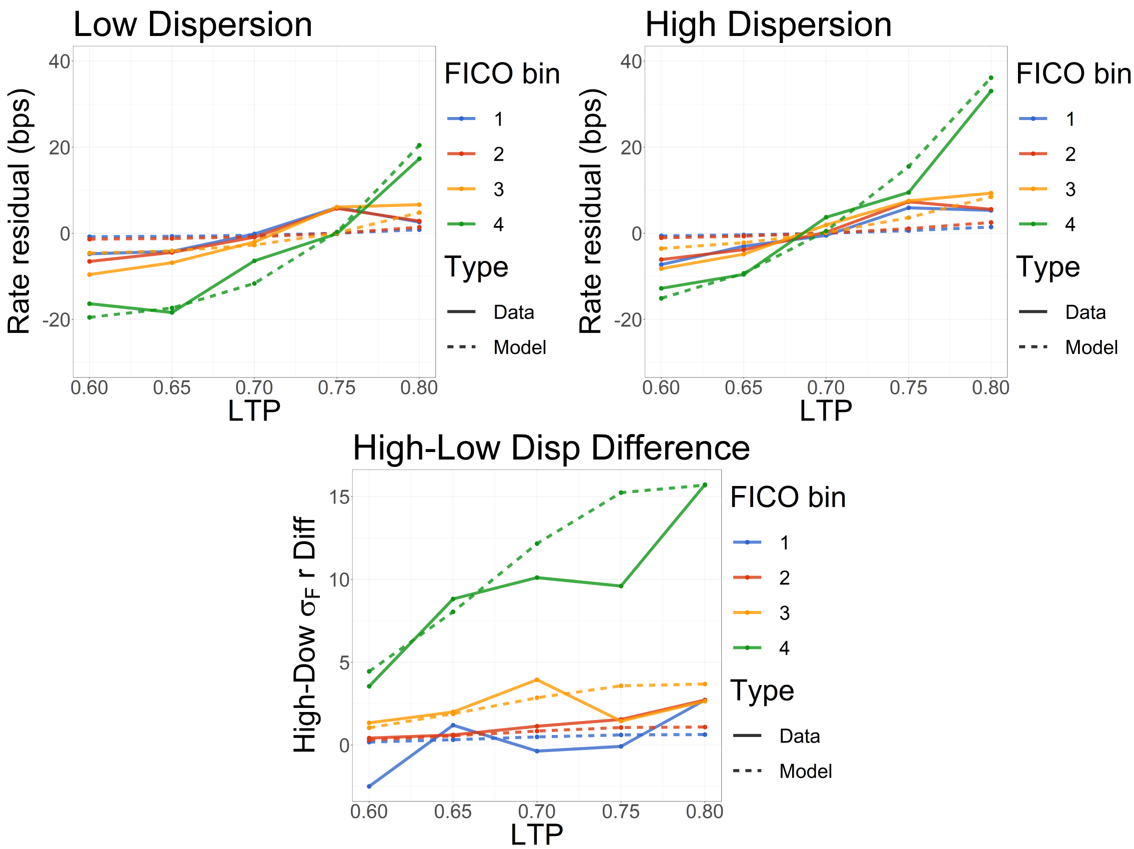

Notes: This figure shows how well our calibration in Appendix G.2 is able to fit the rate menu in the data. The top two panels show empirical interest rate residuals $\tilde{r}_{fld}$ (solid lines), from (A8), and model-predicted rate residuals $\tilde{r}^{model}_{fld}(c)$ (dashed lines), from (A10), in the fitted model. LTP ratios are shown on the x-axis, and different FICO buckets are shown as different colors. The top left plot shows results for low-dispersion areas, and the top right plot shows results for high-dispersion areas. The bottom plot shows the differences $r_{lf,d=H} - r_{lf,d=L}$ in the data (solid) and in the model (dashed). In other words, each line in the bottom panel is the difference between the corresponding line in the top right panel (the high-dispersion menu) and the line in the top left panel (the low dispersion menu).

Table A1: Mortgage Rejection Reasons

This table presents loan level regression results about mortgage rejection reasons. We restrict the sample to only rejected loans, and estimate Specification 4 using various rejection reason indicators as the outcome variables. The explanatory variable of interest is zipcode house price dispersion, scaled by its standard deviation. Borrower/loan controls include zipcode house price, credit score and the squared term, log income, loan type, and loan to income ratio and its squared term. The sample includes loan level observations from 2001 to 2017. Standard errors are clustered at county level. ***, **, * represent 1%, 5%, and 10% significance, respectively.

	(1) Collateral	(2) Down Payment	(3) Debt-to-Income	(4) Employment	(5) Credit Score
Panel A: OLS					
Zip Price Dispersion	1.34*** (0.138)	-0.07*** (0.021)	-0.31*** (0.062)	-0.12*** (0.021)	-0.35*** (0.052)
Local Controls	✓	✓	✓	✓	✓
County-Year FE	✓	✓	✓	✓	✓
Lender-Year FE	✓	✓	✓	✓	✓
R2	0.16	0.10	0.18	0.05	0.24
Panel B: IV					
Zip Price Dispersion	2.45*** (0.229)	-0.02 (0.054)	-0.28 (0.183)	-0.11*** (0.039)	-0.95*** (0.166)
Local Controls	✓	✓	✓	✓	✓
County-Year FE	✓	✓	✓	✓	✓
Lender-Year FE	✓	✓	✓	✓	✓
Underidentification t-stat	64.49	64.49	64.49	64.49	64.49
Underidentification p-value	0.00	0.00	0.00	0.00	0.00
Weak identification t-stat	44.27	44.27	44.27	44.27	44.27
Sample mean	12.23	4.90	17.08	2.76	21.12
Observations	7.5M	7.5M	7.5M	7.5M	7.5M

Table A2: IV Relevance Condition

This table tests the relevance condition for our instruments. The outcome variable is house price dispersion, scaled by its standard deviation. The explanatory variables are the five instruments, introduced in Section 3.2. The three columns correspond to column 1 in Table 3, column 1 in Table 4, and column 3 in Table 5. Standard errors are clustered at county level. ***, **, * represent 1%, 5%, and 10% significance, respectively.

	Price Dispersion		
	(1) Rejection Sample	(2) Cost Menu Sample	(3) LTP Sample
IV: Square Footage	0.02	0.02	0.17***
	(0.027)	(0.025)	(0.007)
IV: Number of Bedrooms	0.11***	0.08	0.07***
	(0.043)	(0.049)	(0.004)
IV: Number of Bathrooms	0.18***	0.18***	0.05***
	(0.033)	(0.033)	(0.006)
IV: Building Age	0.07***	0.05***	0.08***
	(0.020)	(0.015)	(0.008)
IV: Geo-coordinates	0.05***	0.06***	0.06***
	(0.015)	(0.015)	(0.010)
Controls	✓	✓	✓
Origination Month FE		✓	
Transaction Date FE			✓
County-Year FE	✓	✓	✓
Lender-Year FE	✓		✓
R2	0.50	0.49	0.29
Observations	47M	4.8M	28M

Table A3: IV Balance Test

This table presents the balance test results. The underlying sample contains zipcode-year level observations. The outcome variables in column 1 and 2 are log transaction prices obtained from the Corelogic Tax and Deeds data, in columns 3 and 4 are FICO score of transacted mortgages in the Corelogic LLMA data, and in columns 5 and 6 are log applicant income obtained from the HMDA data. Columns 1, 3, and 5 report the OLS results, in which the explanatory variable is the raw price dispersion measures. Columns 2, 4, and 6 report the 2SLS results, in which the raw price dispersion is instrumented using the five IVs introduced in the main text. In all the columns, we include year fixed effects. Standard errors are clustered at county level. ***, **, * represent 1%, 5%, and 10% significance, respectively.

	Log Price		FICO		Log Income	
	(1) OLS	(2) IV	(3) OLS	(4) IV	(5) OLS	(6) IV
Price Dispersion	-0.20***	-0.00	-0.21***	0.10	-0.04***	0.11**
	(0.018)	(0.065)	(0.027)	(0.107)	(0.012)	(0.043)
Year FE	✓	✓	✓	✓	✓	✓
R2	0.12	-	0.15	-	0.04	-
Observation	129,003	129,003	129,003	129,003	129,003	129,003

Table A4: Robustness Tests — Not about Bunching ($\frac{SalePrice}{ConformingLimit} < 1$)

This table presents robustness test for bunching below conforming limit. The outcome variable is loan-to-price ratio. We restrict the sample to house transactions with non-missing mortgage interest rates from Corelogic Deeds and further restrict the sample to houses whose transaction price is smaller than the local conforming loan limit. Standard errors are clustered at county level.

	OLS		2SLS	
	(1)	(2)	(3)	(4)
Price Dispersion	-0.21***	-0.17***	-1.11***	-0.95***
	(0.024)	(0.022)	(0.086)	(0.070)
Interest Rate	0.94***	0.72***	0.95***	0.73***
	(0.071)	(0.049)	(0.072)	(0.050)
Loan Controls	✓	✓	✓	✓
Origination Month FE	✓	✓	✓	✓
County-Year FE	✓		✓	
Lender-Year FE		✓		✓
R2	0.43	0.50	-	-
Observations	4M	4M	4M	4M
Underidentification test statistic			128.45	126.17
Underidentification test p-value			0.00	0.00
Weak identification test statistic			102.39	98.83

Table A5: Price Dispersion and Lender Market Power

This table presents the correlation between price dispersion and lender market power. We divide each county-year into four market segments based on the price dispersion of the underlying properties in Corelogic Deeds transaction records and collapse the sample into *county-year-price dispersion bin* level. In this sample, we have four observations in each county-year that indicate the average price dispersion, the number of lenders, and the lender HHI index for each price dispersion bin in a given county-year. We drop county-year-price dispersion bins with less than 50 transactions. The outcome variable in columns 1-4 is lender HHI index, ranging from 0 to 10000. The outcome variable in columns 5-8 is the number of lenders. Price Dispersion is expressed in terms of the number of standard deviations in columns 1, 3, 5, and 7. Standard errors reported in parentheses are clustered at county level.

	Lender HHI				Number of Lenders			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Price Dispersion	4.41		1.86		1.70***		0.70***	
	(11.933)		(8.597)		(0.476)		(0.120)	
Price Dispersion Quartiles (Benchmark Category: First Quartile)								
Second Quartile		-22.20***		-22.42***		0.25		0.37**
		(6.705)		(6.494)		(0.183)		(0.160)
Third Quartile		-24.56**		-19.58*		0.72***		0.79***
		(10.082)		(10.477)		(0.220)		(0.205)
Fourth Quartile		-16.15		-1.57		1.87***		2.11***
		(18.381)		(21.198)		(0.371)		(0.311)
Outcome Variable Mean	701	701	701	701	68	68	68	68
Controls	✓	✓	✓	✓	✓	✓	✓	✓
County-Year FE			✓	✓			✓	✓
Observations	17,045	17,045	17,015	17,015	17,045	17,045	17,015	17,015
R2	0.13	0.13	0.89	0.89	0.79	0.79	0.99	0.99

Table A6: Robustness Tests — Lender Market Power, Rejection

This table presents robustness tests related to lender market power for the loan rejection results in Table 3. The outcome variable in Panel A is an indicator for whether the loan application is rejected and in Panel B is an indicator for whether the loan application is rejected due to collateral reasons. In both panels, Columns 1-3 are OLS results, and columns 4-6 are 2SLS results. To isolate the effect of lender market power, we construct zipcode lender HHI index and include it as a control in columns 1 and 4. We then classify a zipcode as high (low) HHI if its lender HHI index is above (below) the median value among all zipcodes within a county-year. Columns 2 and 5 report the subsample analysis results using transactions in low HHI zipcodes, and columns 3 and 6 report the subsample analysis results using transactions in high HHI zipcodes. Standard errors are clustered at county level.

Panel A: Rejection

	OLS			2SLS		
	(1) Full Sample	(2) Low HHI	(3) High HHI	(4) Full Sample	(5) Low HHI	(6) High HHI
Price Dispersion	1.30***	1.29***	1.31***	2.42***	2.40***	2.43***
	(0.089)	(0.115)	(0.081)	(0.175)	(0.225)	(0.197)
Zip Lender HHI	0.21***			0.12**		
	(0.043)			(0.048)		
Loan Controls	✓	✓	✓	✓	✓	✓
Lender-County-Year FE	✓	✓	✓	✓	✓	✓
Observations	47M	29M	17M	47M	29M	17M
R2	0.21	0.22	0.22	-	-	-
Underidentification stat				80.40	85.02	60.68
Underidentification p-value				0.00	0.00	0.00
Weak identification stat				61.03	61.09	33.88

Panel B: Rejection Due to Collateral

	OLS			2SLS		
	(1) Full Sample	(2) Low HHI	(3) High HHI	(4) Full Sample	(5) Low HHI	(6) High HHI
Price Dispersion	0.49***	0.50***	0.50***	0.79***	0.79***	0.82***
	(0.035)	(0.042)	(0.034)	(0.061)	(0.081)	(0.064)
Zip Lender HHI	0.05***			0.03**		
	(0.009)			(0.012)		
Loan Controls	✓	✓	✓	✓	✓	✓
Lender-County-Year FE	✓	✓	✓	✓	✓	✓
Observations	47M	29M	17M	47M	29M	17M
R2	0.09	0.10	0.11	-	-	-
Underidentification stat				80.40	85.02	60.68
Underidentification p-value				0.00	0.00	0.00
Weak identification stat				61.03	61.09	33.88

Table A7: Robustness Tests — Lender Market Power, Cost Menu

This table presents robustness tests related to lender market power for the cost menu results in Table 4, using a subsample of house transactions with non-missing mortgage rate information recorded in the Corelogic Deeds and Tax records. Columns 1-3 are OLS results, and columns 4-6 are 2SLS results. To isolate the effect of lender market power, we construct zipcode lender HHI index and include it as a control in columns 1 and 4. We then classify a zipcode as high (low) HHI if its lender HHI index is above (below) the median value among all zipcodes within a county-year. Columns 2 and 5 report the subsample analysis results using transactions in low HHI zipcodes, and columns 3 and 6 report the subsample analysis results using transactions in high HHI zipcodes. Standard errors are clustered at county level.

	OLS			2SLS		
	(1) Full Sample	(2) Low HHI	(3) High HHI	(4) Full Sample	(5) Low HHI	(6) High HHI
Zip Price Dispersion	0.84***	1.25***	1.25***	1.27***	2.04***	2.04***
	(0.135)	(0.136)	(0.136)	(0.331)	(0.278)	(0.278)
Zip Lender HHI	-0.23***			-0.24***		
	(0.061)			(0.064)		
Loan Controls	✓	✓	✓	✓	✓	✓
Origination Month FE	✓	✓	✓	✓		
County-Year FE	✓	✓	✓	✓	✓	✓
Observations	4M	2.3M	2.3M	4M	2.3M	2.3M
R2	0.86	0.86	0.86	-	-	-
Underidentification stat				64.15	62.11	62.11
Underidentification p-value				0.000	0.000	0.000
Weak identification stat				20.43	23.77	23.77

Table A8: Robustness Tests — Lender Market Power, LTP

This table presents robustness tests related to lender market power for the loan size results in Table 5. Columns 1-3 are OLS results, and columns 4-6 are 2SLS results. To isolate the effect of lender market power, we construct zipcode lender HHI index and include it as a control in columns 1 and 4. We then classify a zipcode as high (low) HHI if its lender HHI index is above (below) the median value among all zipcodes within a county-year. Columns 2 and 5 report the subsample analysis results using transactions in low HHI zipcodes, and columns 3 and 6 report the subsample analysis results using transactions in high HHI zipcodes. Standard errors are clustered at county level.

	OLS			2SLS		
	(1) Full Sample	(2) Low HHI	(3) High HHI	(4) Full Sample	(5) Low HHI	(6) High HHI
Price Dispersion	-0.17***	-0.14***	-0.18***	-1.19***	-1.09***	-1.24***
	(0.031)	(0.034)	(0.031)	(0.117)	(0.130)	(0.108)
Zipcode Lender HHI	-0.11***			-0.12***		
	(0.018)			(0.015)		
Loan Controls	✓	✓	✓	✓	✓	✓
Transaction Date FE	✓	✓	✓	✓		
Lender-County-Year FE	✓	✓	✓	✓	✓	✓
N	19M	9M	9M	19M	9M	9M
R2	0.46	0.47	0.48	-	-	-
Underidentification stat				144.29	155.42	140.07
Underidentification p-value				0.00	0.00	0.00
Weak identification stat				190.55	177.35	183.64

Table A9: Property-Level House Price Dispersion and LTP - Repeat Sales

This table presents the results of property-level regressions with repeat sale sigma estimates. The outcome variable is loan-to-sale price ratio. The explanatory variable of interest is property-level house price dispersion estimated using repeat sales, scaled by its standard deviation. Controls include the mortgage rate, transaction price of the property, mortgage type, mortgage term, and resale indicator. The sample includes property transaction level observations from 2000 to 2020. Standard errors are clustered at county level. ***, **, * represent 1%, 5%, and 10% significance, respectively.

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
Price Dispersion	-0.77***	-0.36***	-0.36***	-1.67***	-1.54***	-1.41***
	(0.090)	(0.038)	(0.036)	(0.135)	(0.089)	(0.080)
Interest Rate	0.83***	0.91***	0.69***	0.88***	0.94***	0.71***
	(0.073)	(0.058)	(0.042)	(0.076)	(0.059)	(0.043)
Log House Price	-3.29***	-3.41***	-3.07***	-3.40***	-3.53***	-3.19***
	(0.112)	(0.119)	(0.119)	(0.113)	(0.110)	(0.109)
Loan Controls	✓	✓	✓	✓	✓	✓
Transaction Date FE	✓	✓	✓	✓	✓	✓
County-Year FE		✓	✓		✓	✓
Lender-Year FE			✓			✓
R2	0.44	0.47	0.54	-	-	-
Observations	3M	3M	3M	3M	3M	3M
Underidentification test statistic				84.09	93.32	93.69
Underidentification test p-value				0.00	0.00	0.00
Weak identification test statistic				28.77	46.36	47.66

Table A10: Property-Level House Price Dispersion and LTP — Without House Price as a Control

This table presents property-level regression results without house price as a control. Columns 1-2 present OLS results. Columns 3-4 present IV results. In all columns, the outcome variable is the loan level loan-to-sale price ratio. The explanatory variable of interest in columns 1-2 is property-level house price dispersion, scaled by its standard deviation, and is the predicted price dispersion in columns 3-4. Controls include mortgage type, mortgage term, and resale indicator. The sample includes property transaction level observations from 2000 to 2020. Standard errors are clustered at county level. ***, **, * represent 1%, 5%, and 10% significance, respectively.

	OLS		2SLS	
	(1)	(2)	(3)	(4)
Price Dispersion	-0.30***	-0.33***	-1.62***	-1.65***
	(0.055)	(0.046)	(0.197)	(0.183)
Controls	✓	✓	✓	✓
Transaction Date FE	✓	✓	✓	✓
County-Year FE	✓	✓	✓	✓
Lender-Year FE		✓		✓
R2	0.32	0.37	-	-
Observations	28M	28M	28M	28M
Underidentification test statistic			161.57	164.72
Underidentification test p-value			0.00	0.00
Weak identification test statistic			205.21	201.16

Table A11: Parameter Estimate Sensitivity to Input Moments

This table shows estimated parameters for the baseline model (row 1), along with 6 alternative specifications in which we consider low and high η, the coefficient of relative risk aversion in the first period (rows 2 and 3); low and high β, the annual discount rate (rows 4 and 5); and low and high T, the assumed duration of the mortgage (rows 6 and 7).

Case	ζ	ψ	b	u_2'
Baseline	0.583	2.782	0.086	0.00243
Low Eta	0.400	3.001	0.087	0.01078
High Eta	1.148	2.448	0.086	0.00012
Low Beta	0.644	2.764	0.087	0.00250
High Beta	0.449	2.514	0.087	0.00236
Low T	0.554	2.666	0.085	0.00234
High T	0.653	2.961	0.086	0.00248

Table A12: Counterfactual Estimate Sensitivity to Input Moments

This table shows estimated counterfactual magnitudes for the baseline model (row 1) along with 6 alternative specifications in which we consider low and high η, the coefficient of relative risk aversion in the first period (rows 2 and 3); low and high β, the annual discount rate (rows 4 and 5); and low and high T, the assumed duration of the mortgage (rows 6 and 7). Column 1 shows the "compensating variation" for consumers in the 5th σ-decile, as described in Subsection 6.1 in the main text. Column 2 and 3 show the change in mortgage failure probability, and compensating variation in prices, under the version of the automated appraiser counterfactual in Panel A of Table 8, where we assume automated appraisals remove human biases, but do not change appraisal variance. Column 4 and 5 show the change in failure probability, and compensating variation in prices, under the counterfactual in Panel B of Table 8, in which automated appraisals are calibrated to maintain the same failure probability, but reduce the variance by half.

		Removing Bias		Reducing Variance	
	(1) Required Compensating Price Change (%)	(2) Fail Prob (pp)	(3) Required Compensating Price Change (%)	(4) Fail Prob (pp)	(5) Required Compensating Price Change (%)
Baseline	-0.658	11.964	-5.786	-1.079	0.500
Low Eta	-0.540	11.443	-4.479	-1.023	0.349
High Eta	-0.948	12.717	-8.417	-1.072	0.890
Low Beta	-0.630	12.019	-5.601	-1.054	0.476
High Beta	-0.624	12.347	-5.511	-1.052	0.446
Low T	-0.626	11.895	-5.941	-1.104	0.534
High T	-0.740	12.002	-5.678	-1.065	0.498