



Federal Deposit Insurance Corporation • Center for Financial Research

WORKING PAPER SERIES

Regional Variation in Transaction Costs, Mortgage Rate Heterogeneity, and Mortgage Refinancing Behavior

Leonard Kiefer
Freddie Mac

Hua Kiefer
Federal Deposit Insurance Corporation

Tom Mayock
Virginia Tech

October 2023

FDIC CFR WP 2023-03

fdic.gov/cfr

The Center for Financial Research (CFR) Working Paper Series allows CFR staff and their coauthors to circulate preliminary research findings to stimulate discussion and critical comment. Views and opinions expressed in CFR Working Papers reflect those of the authors and do not necessarily reflect those of the FDIC or the United States. Comments and suggestions are welcome and should be directed to the authors. References should cite this research as a “FDIC CFR Working Paper” and should note that findings and conclusions in working papers may be preliminary and subject to revision.

Regional Variation in Transaction Costs, Mortgage Rate Heterogeneity, and Mortgage Refinancing Behavior *

Leonard Kiefer[†] Hua Kiefer[‡] Tom Mayock[§]

October 23, 2023

Abstract

Recent work has demonstrated that the U.S. mortgage market is characterized by significant heterogeneity in the interest rates that are offered to borrowers as well as mortgage refinancing behavior. In this study we contribute to the mortgage heterogeneity literature by providing the first systematic analysis of regional differences in transaction costs in the mortgage market. Using the Uniform Closing Database—a unique repository of loan-level closing cost information—we demonstrate that there is a tremendous amount of regional variation in transaction costs in the mortgage market, most of which is driven by differences in local mortgage stamp taxes and recording fees.

In the second part of our paper, we take up the question of how failing to account for such heterogeneity might affect studies of borrower behavior in the mortgage market. We do so through the lens of the failure-to-refinance literature on optimal refinancing activity. Accounting for rate and closing costs heterogeneity significantly reduces estimates of suboptimal refinancing behavior, particularly among borrowers with high-risk credit profiles and those living in states with high closing costs. Because regional variation in closing costs is driven by the state and municipal policies, our results suggest that local governments play a role in the pass through of monetary policy via the mortgage market that has not been previously documented. Our findings also provide a potential mechanism above and beyond

*The opinions expressed in this paper are those of the authors alone and do not necessarily reflect the views of the Federal Deposit Insurance Corporation, the United States of America, or Freddie Mac. The authors would like to thank Michael Lacour-Little, David Zhang, and participants at the 2023 AREUEA National Meetings and FDIC seminar for helpful comments.

[†]Freddie Mac, leonard_kiefer@freddiemac.com

[‡]Federal Deposit Insurance Corporation, hkiefer@fdic.gov

[§]Virginia Tech, tmayock@vt.edu

the home equity channel that could explain regional variation in refinancing activity and consumer spending during recoveries.

Keywords: Mortgage Lending, Household Finance

1 Introduction

Recent research suggests that heterogeneity is the rule as opposed to the exception in the mortgage market. For example, Hurst, Keys, Seru and Vavra (2016) find that default risk varies significantly across markets, several papers have documented significant differences in the interest rates paid by borrowers with nearly identical credit profiles (Woodward and Hall, 2012; Gurun, Matvos and Seru, 2016; Alexandrov and Koulayev, 2018; McManus, Liu and Yi, 2018; Bhutta, Fuster and Hizmo, 2020), and there is now evidence that refinancing behavior differs by region (Beraja, Fuster, Hurst and Vavra, 2019; Keys, Pope and Pope, 2016), credit profile (Keys, Pope and Pope, 2016), and demographic characteristics (Gerardi, Willen and Zhang, 2023; Andersen, Campbell, Nielsen and Ramadorai, 2020).

In this study we contribute to this mortgage heterogeneity literature by providing the first systematic analysis of regional differences in transaction costs in the mortgage market. Using the Uniform Closing Database (UCD), a unique repository of loan-level closing cost information, we show that the transaction costs—or “closing costs”—associated with refinancing a loan of a fixed amount vary tremendously across markets in the U.S.¹ As we discuss in more detail in Section 2, regional variation in closing costs is driven primarily by interstate—and in some cases *intrastate*—differences in the tax treatment of new mortgage originations. On net, these regional differences in tax policy create a sizable amount of variance in the transaction costs associated with refinancing a mortgage. For example, after netting out discount points, prepaid homeowner’s insurance, and prepaid escrow fees, the median cost of refinancing a conventional 30-year fixed-rate mortgage with an initial balance between \$190,000 and \$210,000 in our data ranged from a low of \$1,997 in Iowa to a high of \$4,957 in New York: nearly all of this difference is driven by taxes and government fees.²

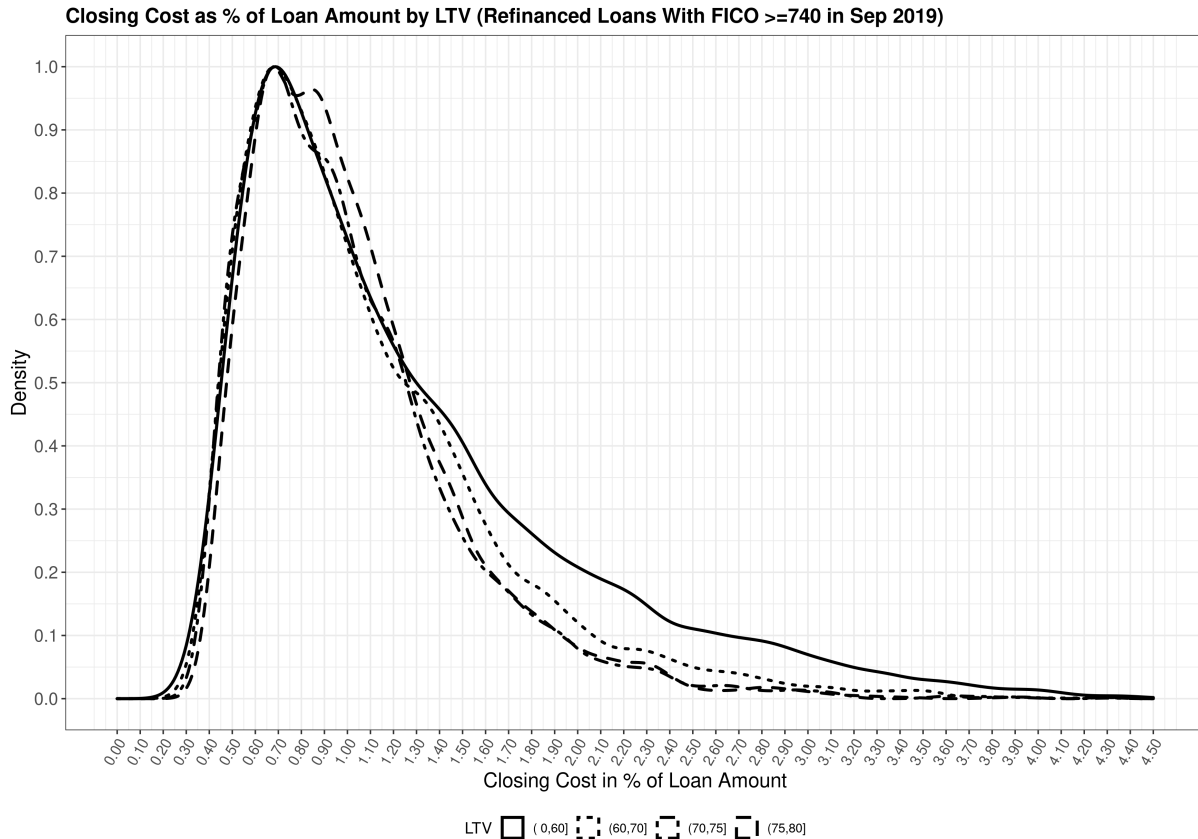
To preview the results of the first part of our analysis, we display in Figure 1 the empirical densities of closing costs expressed as a fraction of the initial principal balance for a subset of loans in our data.³ The variation in transaction costs exhibited by these densities is at odds with the assumption that closing costs scale only with loan balances that

¹When a borrower takes out a new loan on a property for which the annual insurance and property taxes will be paid via an escrow account, some fraction of these annual costs—typically referred to as “prepaids”—must be paid into the escrow account at the time of closing. While prepaids are included in the official Closing Disclosure document, they do not represent transaction costs in the traditional sense of the term: the insurance premium and property tax bill must be paid regardless of whether the borrower is taking out a new mortgage on the home. For this reason, we exclude prepaids when calculating closing costs in our analysis. It is worth emphasizing that even if prepaids were viewed as transaction costs, because prepaids reflect the amount of time between the loan’s origination date and the date on which property taxes and insurance premia must be paid, prepaids are not comparable even between otherwise identical loans that are originated at different points in time during the same calendar year.

²We focus on loans with a principal balance between \$190,000 and \$210,000 in this example as loans in this value range represent a significant fraction of new originations in our data.

³The subset of loans used to produce these densities were originated in September of 2019 to borrowers with FICO scores in excess of 740 in three different loan-to-value (LTV) ratios.

Figure 1: Closing Cost Densities for Refinance Loans Originated in September 2019



is frequently employed in studies of the mortgage market. This observation raises the following question: how does failing to account for closing cost heterogeneity impact studies of consumer behavior in the mortgage market? In the second part of our paper, we take this question up through the lens of the failure-to-refinance (FTR) literature pioneered by Agarwal, Driscoll and Laibson (2013) (“ADL” hereafter).⁴

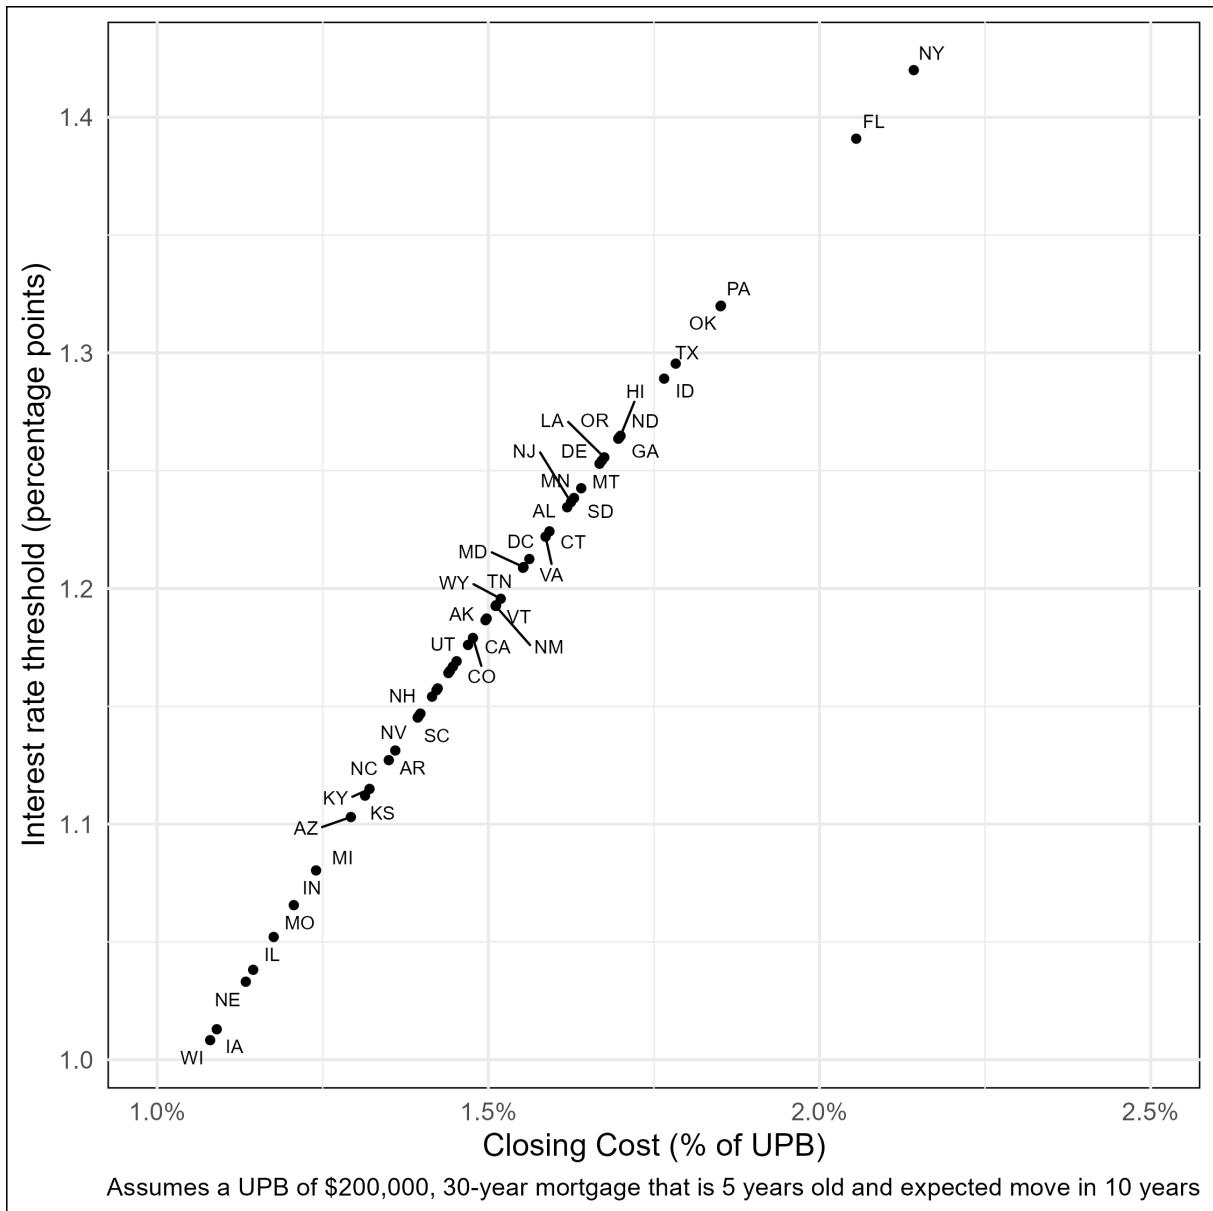
In the context of the FTR literature, failing to account for closing cost heterogeneity could result in estimates of suboptimal refinancing behavior (“FTR rates”) with a systematic upward bias in markets with high closing costs and a downwards bias in markets with low closing costs. To summarize the mechanics of this bias, in Figure 2 we display the minimum difference between the market rate and a borrower’s existing rate (the “interest rate threshold”) for it to be optimal for a hypothetical borrower with a \$200,000 30-year fixed-rate mortgage that expects to move in 10 years to refinance under the ADL framework.⁵ If our hypothetical borrower lived in a state with low closing costs such as Iowa,

⁴Throughout this paper we use FTR rate to measure the proportion of in-the-money borrowers who fail to refinance.

⁵For example, if a borrower currently has a loan with carrying an interest rate of 4 percent, if the ADL interest rate threshold is 1.1, the market interest rate must fall to 2.9 percent for it to be optimal to refinance.

it is optimal to refinance in this example when the rate difference is roughly 100 basis points. If this borrower lived in New York, on the other hand, interest rates would have to fall by more than 140 basis points relative to the borrower's existing rate to justify refinancing. Ignoring heterogeneity in closing costs would thus lead us to systematically overestimate suboptimal refinancing activity in New York while systematically underestimating suboptimal behavior in Iowa. By similar logic, if we fail to account for variation in the mortgage rates offered to buyers with different credit profiles, the FTR rates would be biased upward for low-risk borrowers who are offered below-average rates and biased downwards for high-risk borrowers.

Figure 2: Interest Rate Threshold for Optimal Refinancing Rule



In our analysis of refinancing behavior, we use the UCD records and a servicing database of residential mortgages between 2005 and 2019 to adjust the aforementioned ADL interest rate threshold to account for heterogeneity in both closing costs and interest rates to calculate an “adjusted FTR” rate. To get a sense of how failing to account for heterogeneity impacts studies of optimal refinancing behavior, we benchmark the adjusted FTR against an “unadjusted FTR” measure that does not account for variation in mortgage rates and transaction costs.

There are three primary takeaways from this FTR analysis, which we present in Section 4. First, when we aggregate across all loans, the adjusted FTR is typically lower than the unadjusted FTR; this finding implies that, in the aggregate, failing to account for rate and closing cost heterogeneity will overstate the amount of suboptimal refinancing behavior in the U.S. Second, when we study refinancing behavior in subgroups based on credit scores and loan-to-values (LTVs), we find that accounting for mortgage rate heterogeneity significantly reduces the estimated FTR incidence among borrowers with high LTVs and low credit scores. The high FTR rates reported for high-risk borrowers in previous work thus appear to be explainable, in part, by such borrowers facing higher mortgage rates that make refinancing less attractive.⁶ Lastly, when we study refinancing behavior on a state-by-state basis, we find that accounting for variation in closing costs significantly lowers the FTR rate in states with high closing costs, such as New York and Florida. Accounting for closing cost differences across states thus appears to be particularly important when conducting any regional comparisons of refinancing behavior.

Regarding the existing literature, our analysis of the UCD data is most closely related to a very small set of studies on closing costs in the U.S. mortgage market. To our knowledge this entire literature is comprised of just four studies. Shroder (2007) uses a small sample of loans guaranteed by the Federal Housing Administration that were originated in 1997 to show that the “Good Faith Estimate” of total lending and title fees that lenders were required to provide under the Real Estate Settlement and Procedures Act is an unbiased estimate of the actual fees paid by borrowers. Woodward (2008) also focuses on FHA lending and uses mortgages originated in 2001 to demonstrate that even in a pool of relatively homogeneous loans, origination fees, title fees, and real estate agent fees can vary significantly across borrowers. In a similar study, Feinburg, Kuehn, McKernan, Wissoker and Zhang (2012) provides evidence of significant heterogeneity in title insurance charges for FHA loans in five U.S. cities.⁷ Most recently, Acolin and Walter (2020) use the closing cost variables in the 2019 Home Mortgage Disclosure Act (HMDA) data to conduct a descriptive analysis of how mortgage origination costs vary with borrower and

⁶Accounting for interest rate heterogeneity from the possibility that borrowers with low credit scores and high LTVs are charged higher interest rates in the Keys, Pope and Pope (2016) measures would reduce the estimated failure-to-refinance rate among borrowers with low credit scores and high LTVs while raising the FTR rate among borrowers with high credit scores and low LTVs.

⁷“Title charges” in Feinburg et al. (2012) refer to the cost associated with purchasing title insurance, such as attorney fees and insurance premia. Title charges are only one component of total closing costs, and they do not include mortgage taxes that drive the regional variation in transaction costs that we document in the UCD data.

loan characteristics such as loan-to-value ratios, the presence of a government guarantee, and a borrower’s race. Notably, none of these studies addresses the regional heterogeneity in closing costs that is the focus of our analysis.

The papers closest to the optimal-refinancing analysis in the second part of our paper are Keys, Pope and Pope (2016) and Andersen et al. (2020). Using a sample of U.S. mortgages active in December 2010, Keys, Pope and Pope (2016) find that, based on the ADL optimal-refinancing threshold, the overall FTR rate for U.S. mortgages is between 31.1 percent and 41.2 percent.⁸ Importantly, this paper documents significant heterogeneity in the FTR rate among subpopulations, even in a sample that is limited to low-risk borrowers. For example, among the low-risk population, the FTR rate ranges from 12.3 percent for borrowers in the highest credit-score quartile to 20.1 percent among borrowers in the lowest credit-score quartile.⁹ The incidence of suboptimal refinancing activity also varies significantly across the LTV distribution, with the FTR ranging from 17.5 percent in the lowest LTV quartile to 23.4 percent in the highest LTV quartile.¹⁰ Beyond credit risk characteristics, Keys, Pope and Pope (2016) also provide evidence that the FTR rate exhibits significant regional variation across states, ranging from 8 percent—9 percent in Mississippi, Wisconsin, and Minnesota to roughly 30 percent in Florida, Oklahoma and New York.

Andersen et al. (2020) also use the ADL framework to quantify the prevalence of suboptimal refinancing behavior, but their study focuses on Danish borrowers with fixed-rate mortgages between 2009 and 2017. The authors find that while 43.7 percent of borrowers in the sample should refinance, only 7.6 percent actually did. The implied FTR of 36.1 percent in this study is very similar to the 31.1 percent to 41.2 percent FTR range reported in Keys, Pope and Pope (2016).¹¹ To gain a better understanding of what is driving this inaction in refinancing behavior, Andersen et al. (2020) estimate a modified version of the ADL model that allows for borrower-level heterogeneity in the optimal-refinancing threshold. The results from this exercise imply that household level observables—such as the borrower’s age, residential location, and measures of financial literacy—play an important role in explaining variation in the propensity to refinance and thus any measures of the failure to refinance.

Beyond the FTR literature, our findings also contribute to recent work studying the heterogeneous impact of monetary policy. Beraja et al. (2019) argue that regional variation in

⁸Keys, Pope and Pope’s (2016) sample is comprised of fixed-rate loans that are not guaranteed by the Federal Housing Administration or U.S. Department of Veterans Affairs. Additionally, the sample is limited to loans collateralized by owner-occupied single-family units that had a balance of at least \$75,000 as of December 2010.

⁹Keys, Pope and Pope (2016) classify low-risk borrowers as those with loan-to-value ratios below 90 and credit scores in excess of 680.

¹⁰The overall FTR rate in the low-risk borrower sample in this study is 31.1 percent.

¹¹In related work, Agarwal, Rosen and Yao (2016) find that even among borrowers who refinance, many borrowers either wait too long to refinance or refinance at a suboptimal interest rate. Interestingly, when they analyze a sample of borrowers who refinanced repeatedly in their data, the authors find evidence that borrowers make smaller errors if they have previously refinanced.

home equity resulted in geographic heterogeneity in refinancing activity during the Federal Reserve’s quantitative easing program. Our closing cost analysis suggests an additional mechanism beyond the home equity channel that could drive such regional heterogeneity in refinancing activity. Specifically, inter-state variation in the taxation of new mortgage originations, and thus closing costs, alters the financial incentives to refinance for otherwise identical borrowers. In conjunction with Beraja et al.’s (2019) work, our results thus imply that the tax treatment of mortgage debt by state and local governments might play an important role in how monetary policy is transferred to the real economy. Our FTR analysis is also related to recent work (Gerardi, Willen and Zhang, 2023) on heterogeneity in refinancing activity across borrowers of different races. Minority borrowers typically have lower credit scores and higher loan-to-value ratios and we document significant interest rate variation across these dimensions. Our work thus suggests that some of the racial variation in refinancing activity may be driven by interest rate variation across credit risk segments as well as geographic variation in the concentration of minority borrowers. We address this possibility directly in the final portion of our analysis.

The remainder of the paper is structured as follows. Section 2 introduces the UCD closing-cost data and provides a detailed analysis of the geographic variation in these costs. Section 3 discusses the data and methodology that we use to account for mortgage rate heterogeneity in analysis of refinancing behavior. Section 4 introduces the FTR framework, the results of which are reported in Section 5. Section 6 discusses the implications of our findings for quantifying FTR rates by race and ethnicity, and Section 7 concludes.

2 Closing Cost Heterogeneity

There are significant up-front costs that borrowers typically must pay to refinance their mortgage. These costs include loan origination charges, discount points the borrower may choose to pay to lower their interest rates, taxes and other government fees, and fees for an assortment of services that the borrower either may or may not be permitted to shop for.¹² Additionally, borrowers may need to prepay a portion of their interest, mortgage insurance, property taxes, and homeowner’s insurance at closing. These costs can easily add up to thousands of dollars, an amount significant enough to tilt the financial calculus on whether or not refinancing is optimal.

As discussed above, previous work on consumer behavior in the mortgage market has been based on an assumption—likely driven by a lack of loan-level closing cost data—that closing costs vary only with the mortgage balance but not with a borrower’s geography. The first contribution of our analysis is demonstrating that this assumption does not hold

¹²Under current mortgage lending regulations, borrowers are permitted to shop for appraisals, property inspections, and title insurance. Other services associated with closing the loan may be chosen by the lender.

true in the context of the U.S. mortgage market. We provide evidence of this assertion using data from the Uniform Closing Database (UCD). The UCD is a loan-level database comprised of all loans that were purchased by Fannie Mae and Freddie Mac between January of 2018 and May of 2021 and constitutes roughly 60 percent of the total U.S. mortgage market during the sample period. The UCD records are incredibly rich and report information on virtually every field that is recorded in the loan closing documents.

The first step in documenting closing cost heterogeneity requires calculating the costs a borrower faces when refinancing a loan. Towards that end, we use the UCD’s closing document fields to calculate a total closing cost measure, which we denote CC_{sq} . For each mortgage s on a property located in state q , we calculate CC_{sq} by summing over the following UCD fields: origination charges net of discount points; services that the borrower did not shop for; services that the borrower did shop for; taxes and other government fees, including transfer taxes and recording fees; and any fees listed in the “Other Fees” field.¹³

It is worth noting that in addition to the components that we include in CC_{sq} , the UCD also reports information on prepayments into escrow accounts at the time of origination as well as discount points paid by the borrower to reduce the interest rate on the loan. If a borrower opts to have their servicer pay property taxes and homeowner’s insurance premia on their behalf, he or she will typically need to prepay some portion of the property tax bill and insurance into an escrow account at the time a loan is originated.¹⁴ Because a borrower must make tax and insurance payments regardless of whether the loan is refinanced, prepayments into escrow accounts that are reported in the UCD are not transaction costs in an economic sense, and for this reason, we do not include escrow prepayments in CC_{sq} . For similar reasons, we also do not include discount points in our origination charge measure as borrowers are free to vary the amount of points they pay at closing to receive higher or lower mortgage rates.¹⁵

After limiting the UCD records to conventional, 30-year fixed-rate refinance mortgages to match the loans used elsewhere in our analysis, we were left with more than 5 million loan-level closing-cost observations. We report key summary statistics for this sample in

¹³The origination fees in our closing cost measure include any application fees as well as the fee that the lender charges for underwriting the loan. Lenders are required to allow borrowers to shop for some services that are used in the loan closing process, such as property inspections and closing agents. If borrowers elect to shop for these services, the cost of the services will be listed under the “Services the Borrower Did Shop For” field. If borrowers instead elect to let the lender select the provider of these services, any associated fees will be reported under the “Services the Borrower Did Not Shop For” field. The “Other Fees” field can contain payments for services such home inspections and title insurance. If a loan was originated without discount points, we define origination charges for that loan as the total upfront origination charges reported in the UCD. If a borrower purchased discount points, we define origination charges as the total upfront origination charges in the UCD less the discount points.

¹⁴The amount of funds that must be paid into escrow at origination depend on when the loan is originated. If property taxes are not due for another year, for example, the borrower would only have to prepay a small fraction of the total tax bill when the loan is closed. Alternatively, if the tax bill is due shortly after closing, the required prepaid taxes into escrow would be roughly equal to the full property tax bill.

¹⁵See Stanton and Wallace (1998) for a formal analysis of mortgage point choice.

Table 1: Closing Cost Summary Statistics

	Cost Type		
	Total Net Cost	Origination Charges	Taxes and Government Fees
Mean	\$2,844	\$2,181	\$265
Median	\$2,990	\$1,510	\$140
Standard Deviation	\$1,136	\$1,905	\$374
25 th Percentile	\$2,372	\$1,030	\$69
75 th Percentile	\$3,701	\$2,913	\$281
Observations	5,443,056	5,443,056	5,443,056

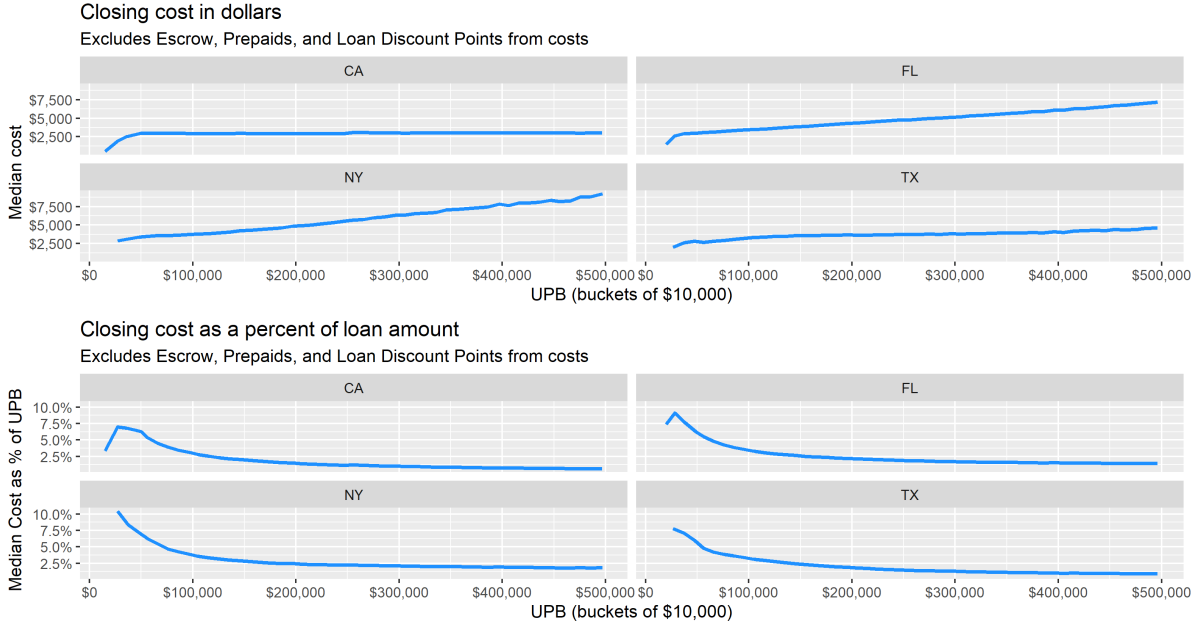
This table summarizes the net closing cost and its two main components. The sample is comprised of conventional, 30-year, fixed-rate, refinance mortgages reported in the UCD that closed between January 1, 2018 and May 5, 2021. Total Net Cost is defined as the sum of: origination charges net of discount points; services that the borrower did not shop for; services that the borrower did shop for; taxes and other government fees, including transfer taxes and recording fees; and any fees listed in the “other fees” field. Discount points are not included in our Origination Charges measure.

Table 1. The average and median total net closing costs in our sample are \$2,844 and \$2,990, respectively. As evidenced by Figure 3, which reports average total net closing costs, origination charges, and taxes and government fees on a state-by-state basis, pooling loans across states obscures a considerable amount of geographic variation how closing costs are determined for a particular loan.

An analysis of the sub-components of total closing cost measures reveals that much of this regional heterogeneity is driven by inter-state differences in tax policies. Taxes and government fees for refinancing are very low in much of the U.S. For example, the average borrower in the state with the lowest closing costs in our sample (Iowa) paid only \$84 in taxes and fees. In two of the largest states by population (New York and Florida), the fee and tax component of closing costs exceeds \$1,000 on average. Florida’s high closing costs are driven by a combination of two separate taxes: a documentary stamp tax of 35 cents per \$100 of the mortgage’s outstanding balance at origination (*Florida Statute 201.08*, 2020) and an intangible tax of 20 cents per \$100 of the mortgage balance (*Florida Statute 199.133*, 2020).

The mortgage tax landscape is even more complex in New York, where the mortgage recording tax varies by county and loan amount. The minimum recording tax in state of New York is 1.05 percent of the mortgage balance, and borrowers in New York City could

Figure 4: Closing Costs by Loan Balance for Select States



This figure displays the median closing costs as well as the median of closing costs as a fraction of the unpaid principal balance for loans of different balance amounts for select states in our sample.

between UPB and total closing costs. That said, to estimate the closing costs that a borrower would face when refinancing, we first use the UCD data to estimate simple regression models of the following form on a state-by-state basis

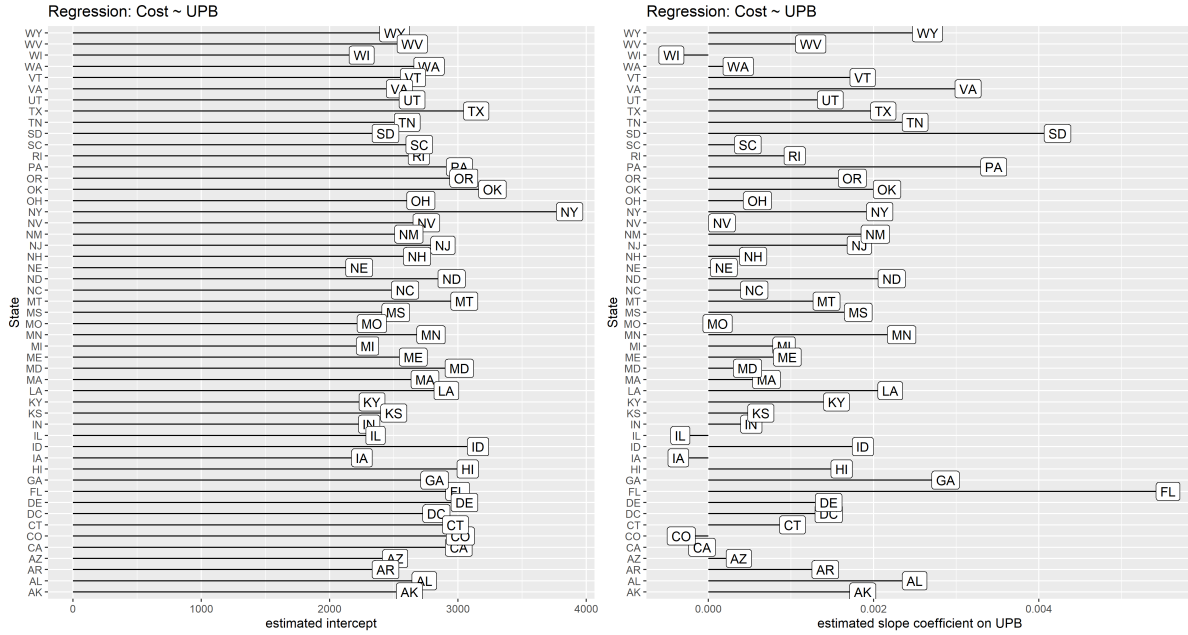
$$CC_{sq} = \alpha_q + \beta_q UPB_{sq} + \epsilon_{sq} \quad (1)$$

where UPB_{sq} denotes the outstanding balance for loan s in state q at the time when it was refinanced and ϵ_{sq} is the error term. Note that because we estimate this model on a state-by-state basis, we allow for regional variation in both the fixed costs of refinancing, as captured by α_q , and the relationship between total closing costs and UPB_{sq} , as captured by β_q . We summarize the results of these regressions in Figure 5, which reports the estimated slope and intercept from Equation 1 for each state in our sample.

Our FTR analysis in Section 4 uses the estimated regression parameters summarized in Figure 5 in conjunction with the property location and outstanding principal balance to construct our preferred estimate of the closing costs that a borrower would face if they opted to refinance.¹⁷

¹⁷In the context of Equation 1, Agarwal, Driscoll and Laibson (2013) and Keys, Pope and Pope (2016) assumed that $\alpha_q = 2000$ and $\beta_q = 0.007905$ for all loans, regardless of the state in which they were originated. This value for β_q was based on a closing-cost expression derived in Appendix A of Agarwal, Driscoll and Laibson (2013).

Figure 5: Closing Cost Regression Parameters by State



This figure reports the estimated slope and intercept from Equation 1 for each state in our UCD sample.

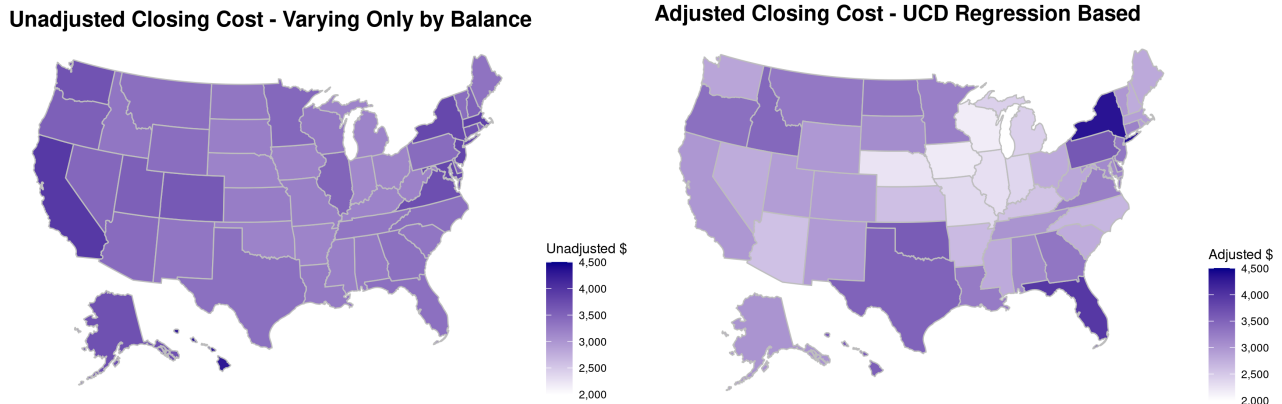
To get a sense of how ignoring state-level heterogeneity in closing costs might affect studies of borrower refinancing behavior, we took the loan-month observations from our McDash data sample¹⁸ and estimated transaction costs in a manner identical to previous studies in the U.S (“Unadjusted Closing Costs”).¹⁹ Then, we used the regression parameters from Equation 1 estimated using the UCD data to estimate closing costs on the same set of loans (“Adjusted Closing Costs”). Figure 6 displays the state-by-state comparisons of these two closing-cost measures, with darker shades of purple denoting higher average closing costs.

Unsurprisingly, there is little geographic heterogeneity in closing cost estimates when we assume away regional variation in the relationship between loan balances and closing costs: by assumption, all of the variation in the left panel of this figure is driven by differences in loan amounts. When we account for geographic heterogeneity in closing costs, however, a very different story emerges: the adjusted closing costs for mortgages in some states like Florida and New York are markedly higher on average, while the adjusted transaction costs in states like Iowa and Nebraska are significantly lower. All else equal, this simple analysis implies that failing to account for regional differences in closing costs could significantly overstate the amount of suboptimal refinancing behavior in states with

¹⁸The McDash data sample is the main data source underlying our FTR calculation in Section 4.

¹⁹Across all states, this sample is comprised of roughly 829 million loan-months observations. Section 4 describes the construction of the sample in detail.

Figure 6: Geographic Variation in Unadjusted and Adjusted Closing Costs



This figure displays the state-by-state averages of our two closing-cost measures: Unadjusted Closing Costs and Adjusted Closing Cost, with darker shades of purple denoting higher average closing costs. The loan sample underlying this figure is from our McDash data used for the FTR analysis.

high taxes and fees on mortgage originations, such as Florida and New York.

3 Interest Rate Heterogeneity

Previous mortgage refinance research has assumed that all borrowers could potentially refinance at the same interest rate, regardless of the borrower’s credit profile. This common-rate assumption is at odds with the emerging literature on mortgage rate heterogeneity which finds that rates vary significantly both across and within borrower risk segments (Woodward and Hall, 2012; Gurun, Matvos and Seru, 2016; Alexandrov and Koulayev, 2018; McManus, Liu and Yi, 2018; Bhutta, Fuster and Hizmo, 2020). We use loan-level origination records from Black Knight’s McDash servicing database to account for this heterogeneity. In addition to reporting “dynamic” information on mortgages such as a loan’s delinquency status, outstanding balance, and borrowers’ refreshed credit scores, the McDash database also includes static variables on loan and borrower characteristics as of a loan’s origination date—such as the credit score, LTV, and interest rate used to underwrite the loan.

We use the static component of the McDash records to estimate the rate at which a loan would likely be refinanced. The first step in this estimation process is defining loan segments based on LTV and credit scores that align with the typical segmentation schemes used in mortgage pricing rate sheets.²⁰ We then limit the McDash data to fixed-rate first-lien *refinance* loans on single-family owner-occupied homes that satisfied the following

²⁰The credit score ranges that we use to define our segments are: (0, 640), [640, 660), [660, 680), [680, 700), [700, 720), [720, 740), and [740, 850). The LTV ranges that we use to define our segments

restrictions: the loan had a 30-year term; the loan was not a cash-out refinance; the loan was not guaranteed by the FHA or VA; and the loan did not have exotic characteristics such as negative amortization periods, balloon payments, or prepayment penalties. Next, to align the origination population with the sample that we use in the FTR analysis described in Section 4, we remove mortgages from the data with original loan balances that were below \$75,000 or greater than the relevant conforming loan limit.

After imposing these data restrictions, we were left with a sample of more than 6 million distinct refinance mortgage originations. We then used the static McDash variables to calculate the median mortgage rates for newly originated loans on a month-by-month basis for each LTV-FICO segment.²¹ These median rates, which we denote as i_{rt} , serve as our estimate of the rate that a loan in segment r as of month t would receive upon refinancing. The loans in the McDash data are drawn from across the credit spectrum, and in the aggregate the McDash data has covered between 52 and 70 percent of the total U.S. market since 2005. That said, i_{rt} should thus serve as a good approximation to the rate a borrower in a particular risk segment would expect to receive if decided to refinance.

Consistent with previous work, we find that mortgage rates vary considerably both within and across segments. For example, consider Figure 7, which depicts box plots for each of the risk segments for loans originated in September of 2019 that satisfy the restrictions discussed above.²² In this plot, within each FICO-LTV segment the upper and lower boundaries of the rectangle represent the 25 percent and 75 percent quantiles of the rate distribution, respectively, while the segment-level mean is depicted by a red dot and the segment-level median is depicted by the horizontal black line.

In this particular month, for loans in the lowest risk segments with LTVs less than 75 and FICO scores in excess of 740, the median interest rate on originations was 3.625 percent. In contrast, in segments with LTVs in excess of 90 and FICO scores below 700, all of the segment specific rates were 4 percent, with rates in the highest risk segments exceeding 4.125 percent. The overall mean interest rate in the September 2019 sample was 3.8 percent, and the reported rate in Freddie Mac’s Primary Mortgage Market Survey (PMMS) in that month ranged from 3.49 percent to 3.64 percent. The variation in rates summarized in Figure 7 suggests that if we assume that all borrowers can refinance at either the average rate in the market or at the rate reported in the PMMS rate, we will seriously understate the true mortgage rates facing high-risk borrowers who might consider

are: (0, 60], (60, 70], (70, 75], (75, 80], (80, 85], (85, 90], (90, 95], and (95, 100].

²¹For example, one segment is defined as borrowers with a credit score greater than 740 and an LTV less than 60 taking out a 30-year fixed-rate mortgage. To construct loan pricing measures for this segment in January of 2019, we identify all 30-year fixed-rate refinance mortgages in the McDash records that satisfy our data restrictions that were originated in January of 2019 to borrowers with credit scores greater than 740 and LTVs less than 60. We then take the median of these interest rates to construct i_{rt} .

²²Rate variation across risk segments was evident in every month in our sample. We chose to focus on the variation in one month here to simplify the exposition. Box plots of rate variation by risk segment for every month in our sample are available upon request.

refinancing, while the converse is true for very low-risk borrowers. Much like our analysis of closing cost variation, our mortgage rate analysis implies that our FTR analysis should account for the sizable and systematic heterogeneity in interest rates across different risk segments.

4 Failure to Refinance Framework

Whether or not a borrower should refinance a mortgage is a complex decision that depends on a host of variables. The starting point for all FTR analyses is adopting a framework that maps all of these factors into a decision rule that can be used to determine when it is in the borrower's interest to refinance his or her mortgage. Following previous work on the topic, we use the methodology established in Agarwal, Driscoll and Laibson (2013) for making this determination.

Before outlining the optimal refinancing rule, we must introduce some notation. Let M_{st} denote the outstanding balance on loan s as of month t , let i_{st} denote the prevailing nominal mortgage rate at which loan s could be refinanced as of month t , and define \bar{i}_s as the existing rate on loan s . If x_{st} is defined as

$$x_{st} = i_{st-1} - \bar{i}_s \quad (2)$$

it is financially advantageous for borrowers to refinance their mortgage when x_{st} falls below a critical level denoted x_{st}^* . Agarwal, Driscoll and Laibson (2013) show that x_{st}^* can be approximated as


$$x_{st}^* \approx -\sqrt{\frac{\sigma \kappa_{st}}{M_{st}(1-\tau)}} \sqrt{2(\rho + \lambda_{st})} \quad (3)$$

where σ is the annualized standard deviation of the mortgage interest rate, ρ is the discount rate, τ is the marginal tax rate, and κ_{st} is the cost of refinancing loan s . λ_{st} denotes the expected rate of exogenous mortgage prepayment and is defined as

$$\lambda_{st} = \mu + \frac{\bar{i}_s}{\exp(\bar{i}_s \Gamma_{st}) - 1} + \pi \quad (4)$$

where Γ_{st} is the remaining life of the existing mortgage in years for loan s as of month t , π is the inflation rate, and μ is the exogenous relocation hazard rate. Note that λ_{st} captures only the exogenous factors that reduce the expected value of the real mortgage obligation, such as relocation, principle repayment, and inflation; mortgage terminations due to refinancing are endogenous and are not reflected in λ_{st} .

Figure 7: Variation in Mortgage Rates by LTV-FICO Segment for September 2019 Refinance Originations



This figure summarizes the variation of mortgage rates by FICO-LTV buckets using the McDash sample for loans originated in September of 2019. In this plot, within each FICO-LTV segment the upper and lower boundaries of the rectangle represent the 25 percent and 75 percent quantiles of the rate distribution, respectively, while the segment-level mean is depicted by a red dot and the segment-level median is depicted by the horizontal black line.

In this framework, it is optimal to refinance a mortgage if

$$i_{st} - \bar{i}_s < -\sqrt{\frac{\sigma \kappa_{st}}{M_{st}(1 - \tau)}} \sqrt{2(\rho + \lambda_{st})} \quad (5)$$

To study how heterogeneity in mortgage rates and closing costs influences FTR measures, we construct two different estimates of the fraction of borrowers who failed to refinance optimally in each month between January of 2005 and December of 2019. We produce the first measure, which we refer to as the “*Unadjusted FTR*,” using the methodology of Agarwal, Driscoll and Laibson (2013) under the assumption that closing costs only vary with loan balance and that all borrowers receive the same interest rate when refinancing. We then contrast the Unadjusted FTR measures against an “*Adjusted FTR*” that also uses Agarwal, Driscoll and Laibson’s (2013) optimal-refinancing rule but allows for geographic variation in fixed and variable closing costs as well as heterogeneity in mortgage rates across credit risk segments.

We use the dynamic variables in the McDash data to create both the Unadjusted and Adjusted FTR shares. To construct these measures, we first restrict the dynamic McDash records to loan-month observations for owner-occupied first-lien 30-year fixed-rate mortgages on single-family homes that have a loan age of at least two months and have loan balances that are at least \$75,000 but less than the conforming loan limits. We also remove from the sample loans that are guaranteed by the Federal Housing Administration (FHA) or the Department of Veterans Affairs (VA), as well as loans that have exotic characteristics, such as negative amortization schedules, balloon payments, or prepayment penalties. Lastly, we limit the data to loans that are reported as being current as of the reporting month.²³ These data filters align closely with those used in previous work and will thus allow for a comparison between our benchmark Unadjusted FTR results and those reported in previous work.²⁴

Following the existing literature, our Unadjusted FTR measure is constructed based on the assumption that in a particular month all loans in the sample could potentially be refinanced at the same mortgage rate reported in the monthly PMMS survey and that the monetary cost of refinancing at this rate is equal to 0.79 percent of the loan’s outstanding

²³A loan is classified as being current as of a reporting month if the dynamic McDash indicate that the loan is not delinquent nor is the loan in any part of the foreclosure process.

²⁴For example. Keys, Pope and Pope (2016) limit their data to fixed-rate mortgages loans collateralized by single-family owner-occupied homes that had an outstanding balance of at least \$75,000 and were not in foreclosure proceedings as of December of 2010. They also exclude from their sample any loans that were guaranteed by the Federal Housing Administration or the Department of Veterans Affairs. In a robustness check, the authors further limit the sample to loans that had a FICO score in excess of 680, had a current LTV of less than 90, and did not have a missed payment recorded in the servicing records; these limitations were intended to remove from the sample households that could not have refinanced even if they wanted to because of credit conditions at the time.

balance plus \$2,000.²⁵ That is, in the context of Equation 2 and Equation 3, when we construct our Unadjusted FTR measure we assume the following values for the interest rate and closing-cost function parameters

$$\begin{aligned} i_{st} &= i_t^{PM} \\ \kappa_{st} &= 2000 + 0.007905M_{st} \end{aligned} \quad (6)$$

where i_t^{PM} is the rate reported in the PMMS in month t .

For each of the loan-month observations that satisfy the data restrictions discussed above, we use Equation 3 to define the optimal refinancing threshold under the unadjusted parameter assumptions - indicated by a superscript U - as²⁶

$$x_{st}^{U*} \approx -\sqrt{\frac{\sigma(0.007905M_{st} + 2000)}{M_{st}(1 - \tau)}} \sqrt{2(\rho + \lambda_{st})} \quad (7)$$

in which case it is optimal to refinance if

$$i_t^{PM} - \bar{i}_s < x_{st}^{U*}. \quad (8)$$

If P_{st} is an indicator variable that is equal to one if loan s prepays between month $t - 1$ and t , then we classify a loan as failing to refinance ($FTR_{st}^U = 1$) under the unadjusted

²⁵Agarwal, Driscoll and Laibson's (2013) closing-cost function is defined as

$$\kappa_{st} = F_s + \gamma_s M_{st}$$

where M_{st} is the balance of loan s in month t and F_s denotes the fixed component of closing costs. γ_s captures the rate at which closing costs change with the balance of the newly originated loan and is defined as follows

$$\gamma_s = f_s \left[1 - \frac{\tau}{\theta + \rho + \pi} \left(\frac{1 - \exp(-(\theta + \rho + \pi)N)}{N} \frac{\rho + \pi}{\theta + \rho + \pi} + \theta \right) \right]$$

In the expression above, f_s denotes the discount points paid as a percentage of the loan's balance, θ denotes the expected arrival rate of an event such as a move or subsequent refinancing that allows for fully deducting the discount points, N denotes the number of years of the borrower's new mortgage term, and $\tau/\rho/\pi$ are defined the same as in Equations 3 and 4. Agarwal, Driscoll and Laibson (2013) and Keys, Pope and Pope (2016) assume that $\tau = 0.28$, $\rho = 0.05$, $\pi = 0.03$, $\theta = 0.2$, $F_s = 2000$, $N = 30$, and $f_s = 0.01$, resulting in the following closing-cost function

$$\kappa_{st} = 2000 + 0.007905M_{st}$$

²⁶When calculating λ_{st} and x_{st}^{U*} , we assume that $\tau = 0.28$, $\rho = 0.05$, $\pi = 0.03$, and $\mu = 0.1$. Additionally, we set $\sigma = 0.0109$. These assumed parameter values are consistent with the assumptions used to construct the optimal financing threshold in Agarwal, Driscoll and Laibson (2013) and Keys, Pope and Pope (2016).

assumptions as follows

$$FTR_{st}^U = \begin{cases} 1 & \text{if } i_t^{PM} - \bar{i}_s < x_{st}^{U*} \text{ and } P_{st} = 0 \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

If there are S active loans in the sample in month t , the Unadjusted FTR in month t , denoted \overline{FTR}_t^U , is defined as

$$\overline{FTR}_t^U = \frac{1}{S} \sum_{s=1}^S FTR_{st}^U \quad (10)$$

When classifying loan s as failing to refinance for our *Adjusted* FTR measure, we first assign the loan to a FICO-LTV segment using the updated FICO score and LTV measures that are reported for loan s in the dynamic McDash variables in month t . If the loan is assigned to risk segment r in month t , we assume that if the borrower were to refinance, they would do so at the median mortgage rate of *new originations* in month t and segment r in the McDash data that satisfy the data restrictions discussed in detail in Section 3; we let i_{rt} denote this median segment-specific interest rate.²⁷

In addition to allowing for across-segment interest-rate variation, when constructing the adjusted optimal-refinancing threshold, in lieu of assuming that the closing-cost function (κ_{st}) is the same for all borrowers throughout the U.S., we use the estimated state-specific parameters from Equation 1 to construct an adjusted closing-cost measure that allows for inter-state variation in κ_{st} . That is, for our Adjusted FTR measure, we assume the following values for the interest rate and closing costs

$$\begin{aligned} i_{st} &= i_{rt} \\ \kappa_{st} &= \hat{\alpha}_q + \hat{\beta}_q M_{st} \end{aligned} \quad (11)$$

where $\hat{\alpha}_q$ and $\hat{\beta}_q$ are the parameters for Equation 1 estimated using the UCD data for loan originations in state q .

The optimal refinancing threshold under the adjusted parameter assumptions - indicated by a superscript A - is thus

$$x_{st}^{A*} \approx -\sqrt{\frac{\sigma(\hat{\alpha}_q + \hat{\beta}_q M_{st})}{M_{st}(1-\tau)}} \sqrt{2(\rho + \lambda)} \quad (12)$$

where σ , τ , ρ , and λ are set at the same values used to construct the unadjusted optimal-

²⁷For a particular loan, changes in i_{rt} can be driven by migration between the credit risk segments over times as well as by the changes in the prevailing mortgage rate within these segments.

refinance boundary. Under the adjusted parameter assumptions it is optimal for a borrower to refinance if

$$i_{rt} - \bar{i}_s < x_{st}^{A*} \quad (13)$$

and we classify a borrower as failing to refinance ($FTR_{st}^A = 1$) in month t as follows

$$FTR_{st}^A = \begin{cases} 1 & \text{if } i_{rt}^A - \bar{i}_{st} < x_{st}^{A*} \text{ and } P_{st} = 0 \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

We define the Adjusted FTR in month t , denoted \overline{FTR}_t^A , as

$$\overline{FTR}_t^A = \frac{1}{S} \sum_{s=1}^S FTR_{st}^A \quad (15)$$


where again S is the number of active loans in the McDash data in month t .²⁸

5 Failure to Refinance Results

We report in Figure 8 the results of our FTR analysis based on four different calculation methods on a month-by-month basis. The Adjusted FTR and Unadjusted FTR series in this plot were calculated as described in Section 4. We calculated the “Rate-Adjusted Only” FTR using the segment-specific mortgage rates described in Section 3 in conjunction with the unadjusted closing-cost measures that ignore geographical heterogeneity, while the “Cost-Adjusted Only” measure uses the state-specific closing cost estimates from Section 2 while assuming that all borrowers would refinance at the common rate reported in the PMMS, regardless of their risk segment. For all four of the measures, we find that the FTR series are low when interest rates are flat or rising and the series rise rapidly when mortgage rates—as measured by the PMMS rate—begin to decline. These dynamics are to be expected as being able to refinance at a lower rate is a necessary though not sufficient condition for being classified as failing to refinance.

²⁸To account for the time gap between rate lock and closing, which is roughly 45 days, we lag i_{st} and $i_{r,t}$ by one month in our empirical analysis followed, as in (Keys, Pope and Pope, 2016).

Figure 8: Overall Failure to Refinance by Calculation Method



This figure displays the results of our FTR analysis using the McDash sample. The numbers reported here are based on four different calculation methods on a month-by-month basis. The Adjusted FTR and Unadjusted FTR series in this plot were calculated as described in Section 4. We calculated the “Rate-Adjusted Only” FTR using the segment-specific mortgage rates described in Section 3 in conjunction with the unadjusted closing-cost measures that ignore geographical heterogeneity, while the “Cost-Adjusted Only” measure uses the state-specific closing cost estimates from Section 2.

As discussed above, ignoring rate and closing cost heterogeneity could potentially impart downwards or upward bias to FTR measures. We find that when constructed using a pool of loans from all risk segments and geographies, the Unadjusted FTR measure consistently overestimates the amount of suboptimal refinancing behavior. On average, the Unadjusted FTR exceeded the Adjusted FTR by roughly 3 percentage points, and the Unadjusted FTR exceeded the FTR for 83 percent of the months in our sample. We summarize these findings in the first panel of Figure 8. Notably, the difference between the Unadjusted and Adjusted FTRs in levels varies significantly over time, with the differences particularly pronounced when the overall FTR measures are high. For example in 2012—the year in which FTR rates peaked in our sample—the Unadjusted FTR measure exceeded the Adjusted FTR measure by between 9.3 and 19.7 percentage points in any given month.

When overall FTR rates are low, on the other hand, the absolute differences between the Unadjusted and Adjusted series are very small. Some of this compression in the differences in levels is mechanical in nature. When mortgage rates are rising, fewer loans cross the optimal-refinancing boundary regardless of how mortgage rates and closing costs are estimated. The absolute difference between the Adjusted and Unadjusted FTR measures in a rising-rate environment will thus be small as both measures will be very close to zero. The limited absolute difference notwithstanding, there are sizable *proportionate* differences between the measures, even when overall FTR levels are low. The years 2006 and 2007 are a case in point: while the FTR measures are virtually indistinguishable in Figure 8 during this time period, the Unadjusted FTR exceeded the Adjusted FTR on average by 72 percent in these years. The sizable proportionate differences between the FTRs also extend beyond rising-rate periods as the Unadjusted FTR exceeded the Adjusted FTR by 40 percent in the average month in our sample. We interpret these findings as evidence that failing accounting for interest rate and transaction cost heterogeneity is important in absolute terms when mortgage rates are falling and in relative terms regardless of the interest rate environment.

Turning to the middle and bottom panels of Figure 8, we see that the Rate-Adjusted-Only FTR series is virtually identical to the Adjusted FTR series in the top panel, while the Cost-Adjusted Only series mimics the Unadjusted FTR series. The difference between our Adjusted and Unadjusted FTR measures in the full population of mortgages in our data is thus driven primarily by our correction for between-borrower variation in mortgage rates, while our closing-cost correction has a very limited impact on the estimated FTR rates in our pooled sample of loans.

While it is interesting that the Unadjusted FTR appears to be biased upward in a sample of loans drawn from all risk segments and U.S. states, our primary hypothesis concerns differences between the Unadjusted and Adjusted FTR in particular loan subpopulations. Specifically, we expect the Unadjusted FTR to systematically overstate the FTR ratio for high-risk borrowers who face interest rates higher than the PMMS rate and systematically understate the FTR ratio for low-risk borrowers who receive loan offers below the PMMS rate. Likewise, we expect our closing-cost correction to affect FTR measures most for

borrowers living in states where closing costs are significantly above or below the national average.

We test these hypotheses by alternately partitioning the McDash data by risk segment and geography and then comparing the Unadjusted and Adjusted FTRs. For our risk-segment analysis, we first calculate the default rate for loans in each of the LTV-FICO groups that we used to construct segment-specific interest rates.²⁹ We then rank-order the FICO-LTV segments by default risk and calculate the Adjusted and Unadjusted FTR on a segment-by-segment basis. We summarize these measures in the top panel of Figure 9, where Segment 1 denotes the highest risk segment, and Segment 56 denotes the lowest risk segment.

For both measures, the FTR rate falls in a nearly monotonic way as default risk decreases; the exceptions to this pattern, which are depicted by large spikes in the FTR plots, denote risk segments where the LTVs exceed 95.³⁰ Our finding that the Unadjusted FTR declines with credit risk is consistent with previous results in the literature (Keys, Pope and Pope, 2016). The fact that the Adjusted FTR also declines with risk indicates that a higher incidence of suboptimal refinancing activity among borrowers with higher LTVs and lower credit scores cannot be explained entirely by variation in mortgage rates and closing costs.

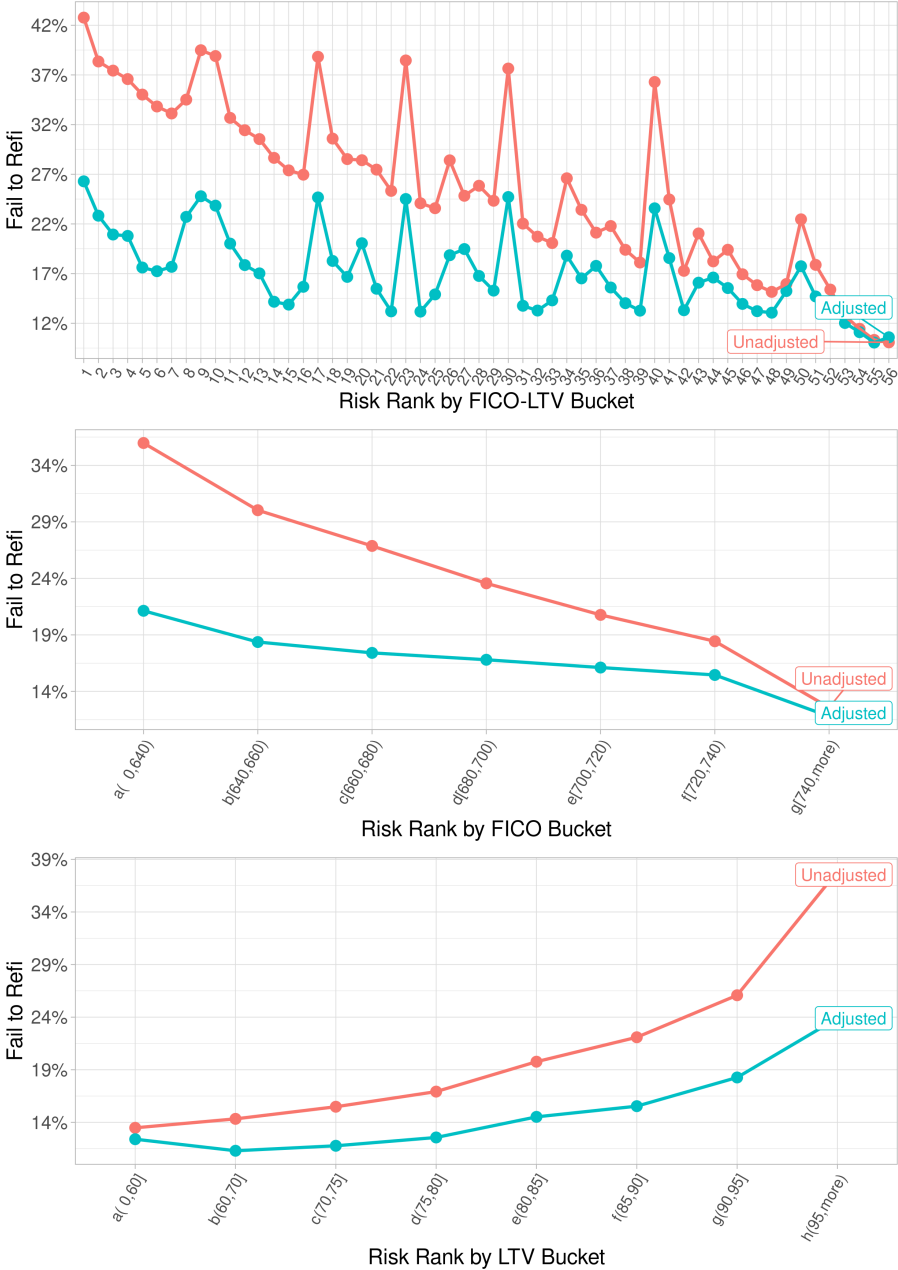
Consistent with our hypothesis, we find that the Unadjusted FTR measure exceeds the Adjusted FTR for higher risk loans that would likely only be able to refinance at rates significantly higher than that reported in the PMMS survey. Furthermore, the *difference* between the Unadjusted and Adjusted FTRs generally declines with default risk, a finding that is also consistent with the PMMS systematically overstating potential interest rate savings for high-risk borrowers while serving as an adequate proxy for the rates facing the average borrower.³¹ It is only in the lowest risk segments that the Adjusted FTR exceeds the Unadjusted FTR; while this can be interpreted as evidence that the PMMS understates the rates facing low-risk borrowers, the absolute magnitude of the bias in Unadjusted FTR for the low-risk segments is significantly smaller than the bias in high-risk segments. For instance, in the highest risk segment, the Unadjusted FTR (42.47 percent) was nearly 17 percentage points higher than the Adjusted FTR (25.99 percent), while in the lowest risk segment, the Unadjusted FTR (9.78 percent) was only 51 basis points lower than the Adjusted FTR (10.29 percent).

²⁹In this part of our analysis, we define a loan as being in default if it is 30 days-past-due or worse as of the observation date.

³⁰Borrowers with high LTVs will typically have difficulty finding a lender to refinance their loan. The spikes in the FTR series at these points thus reflect a combination of suboptimal refinancing behavior as well as the fact that such borrowers are credit constrained.

³¹Interestingly, the difference between the Unadjusted and Adjusted FTR measures also spikes in the segments with LTVs in excess of 95. This pattern may be driven by the fact that the PMMS rate “is based on first-lien prime conventional conforming home purchase mortgages with a loan-to-value of 80 percent” (*PMMS Frequently Asked Questions*, 2022), and rates on loans with LTVs in excess of 95 are typically much higher than the rates on loans with LTVs of 80.

Figure 9: Failure to Refinance by Default Risk



This figure reports the Adjusted and Unadjusted FTR on a segment-by-segment basis using the McDash sample. The FICO-LTV segments are rank-ordered by default risk with Segment 1 denoting the highest risk segment and Segment 56 denoting the lowest risk segment.

We also find strong support for our hypothesis when, in lieu of the LTV-FICO segments, we segment borrowers by LTV and FICO scores separately. We report the results of these analyses in the middle and bottom panels of Figure 9. As summarized in the middle panel, as FICO scores increase, and thus default risk decreases: both FTR measures decline; the Unadjusted FTR exceeds the Adjusted FTR; and the gap between the Adjusted and Unadjusted FTR declines. When we use LTV to segment credit risk in the bottom panel, the results are almost identical qualitatively to those from the FICO analysis.³² When combined with the LTV-FICO segment results in the top panel, we interpret these findings that regardless of how we quantify credit risk, the Unadjusted FTR measure is biased upward and the magnitude of that bias increases with credit risk.

Lastly, to investigate possible heterogeneity in the impact of our closing-cost correction on FTR measures, we calculate the Adjusted and Unadjusted FTRs on a state-by-state basis. We report these measures in Figure 10.³³ The FTR measures and the median closing costs for properties in a particular state as reported in the UCD data are displayed on the y-axis and x-axis, respectively. The state abbreviation in black font denotes the Adjusted FTR, while the state abbreviation in red denotes the Unadjusted FTR. The dashed red and black lines represent the fitted lines from a regression of the Unadjusted and Adjusted FTR measures on median closing costs, respectively.

On a state-by-state basis, the Unadjusted FTR exceeds the Adjusted FTR in all but three states (Iowa, Nebraska, and Wisconsin) in which the measures were virtually identical. Focusing on the left side of Figure 10, we see that in states with low closing costs, the upward bias associated with the Unadjusted FTR was typically quite small. For example, in Iowa and Wisconsin—the states with the lowest median closing costs—the Unadjusted FTRs were 12.4 percent and 12.35 percent, respectively. These figures were little changed after accounting for interest rate and closing cost heterogeneity, as the Adjusted FTR for Iowa was 12.67 percent and the Adjusted FTR for Wisconsin was 12.44 percent.


As median state-level closing costs increase, so does the upward bias of the Unadjusted FTR measure. For instance, the impact of accounting for closing cost heterogeneity was most pronounced in Florida and New York, the states with the highest median closing costs in the UCD data: the Unadjusted FTR in Florida (19.58 percent) and New York (18.97 percent) fell by more than 6 percentage points to 11.06 percent and 12.17 percent, respectively, following the heterogeneity corrections.

Figure 10 makes clear that when studying variation in FTR rates across regions, the net impact of accounting for heterogeneity in closing costs *and* interest rates is most pronounced in states where closing costs on a typical loan exceed the national average.

³²Note that while credit risk is declining as we move from left to right on the x-axis in the middle panel of Figure 9, credit risk is increasing as we move in the same direction in the bottom panel as default risk increases monotonically with LTV.


³³We calculated the FTR measures displayed in in Figure 10 by first calculating the FTRs for each state-month combination in our data. Then, we took the average of the monthly FTR measures within the state.

Figure 10: Failure to Refinance by State



This figure reports the Adjusted and Unadjusted FTRs on a state-by-state basis. The FTR measures along the y-axis are calculated using the McDash sample, and the median closing costs for properties in a particular state along the x-axis are reported in the UCD data. The state abbreviation in black font denotes the Adjusted FTR, while the state abbreviation in red denotes the Unadjusted FTR. The dashed red and black lines represent the fitted lines from a regression of the Unadjusted and Adjusted FTR measures on median closing costs, respectively.

Figure 11: Failure to Refinance by State



This figure reports the Cost-Adjusted Only FTR rate over time for four states using the McDash data: two states with the highest closing costs in our sample - Florida and New York, and their bordering states with significantly lower transaction costs - Georgia and Pennsylvania.

To isolate the impact of the closing-cost correction on cross-region comparisons of sub-optimal refinancing behavior, we calculated the Cost-Adjusted Only FTR rate over time for the two states with the highest closing costs in our sample (Florida and New York). To serve as a point of comparison, we also constructed the Cost-Adjusted Only FTR series for bordering states with significantly lower transaction costs (Georgia and Pennsylvania). We present these series in Figure 11. The Unadjusted and Cost-Adjusted Only series are virtually identical for Georgia and Pennsylvania over the full course of our sample. In Florida and New York, in contrast, while there is little difference between the series when interest rates are stable or rising, the Unadjusted FTR exceeds the Cost-Adjusted FTR when mortgage rates are declining. These results imply that—much like how ignoring interest rate variation imparts an upward bias to FTR measures for high-risk borrowers—failing to account for transaction-cost heterogeneity biases FTR measures upward for loans in states where closing costs are high, but only in declining-rate environments.

6 Implications for FTR Measures by Race and Ethnicity

Previous work has demonstrated that Black and Hispanic borrowers are less likely than White borrowers to refinance (Kelly, 1995; Clapp, Goldberg, Harding and LaCour-Little, 2001; Deng and Gabriel, 2006; Firestone, Van Order and Zorn, 2007; Kau, Fang and Munneke, 2019; Gerardi, Willen and Zhang, 2023). Virtually all of these papers have assumed—either explicitly or implicitly—that all borrowers face the same closing costs and could refinance at the same interest rate, regardless of their location or measures of creditworthiness.³⁴ On average, Black and Hispanic borrowers have lower credit scores and higher LTVs at origination relative to White borrowers. We demonstrated in Section 3 that low-FICO and high-LTV borrowers face interest rates that are significantly above the market-wide average. In the context of our FTR analysis, this implies that using the average mortgage rate to proxy for the rate that available to Black and Hispanic borrowers will systematically underestimate the true rate they could expect to realize and thus systematically overestimate the FTR rate for such borrowers. Furthermore, minority borrowers are not uniformly distributed across U.S. states, and the higher concentration of minorities in states like Florida with high closing costs could also bias FTR measures upward.

These facts suggest that previously documented racial disparities in refinancing may in part be a statistical artifact of heterogeneity in mortgage rates and closing costs. We investigate this possibility using the “core data” records from the National Mortgage Database (NMDB). The NMDB is a restricted-access mortgage database created and maintained jointly by the Federal Housing Finance Agency (FHFA) and the Consumer Financial Protection Bureau (CFPB). The construction of the core NMDB files begins with a 5 percent random sample of all closed-end first-lien mortgage files that were outstanding in Experian’s credit data repository as of January of 1998. The performance of each of these loans was then tracked until the mortgages terminated through prepayment, foreclosure, or maturity. For each quarter following January of 1998, a random 5 percent sample of mortgages that were *newly reported* to Experian were added to the NMDB files, with the performance of these loans also tracked through termination for loans that paid off or through the present time for loans that were still active as of the end of the most recent sampling period (Avery et al., 2021). The result of this sampling process is a 5 percent representative random sample of activity in the U.S. mortgage market that spans more than two decades.

The fields included in the NMDB include both variables reported as of the loan’s origi-

³⁴The exception is Gerardi, Willen and Zhang (2023). This paper utilizes updated credit scores and LTVs to construct a refinance option-value measure (“Call Option V2”) that accounts for risk-based pricing. This adjustment is similar to the rate adjustment that we make with the NMDB data in the analysis described below. While Gerardi, Willen and Zhang (2023) do not explicitly model regional variation in transaction costs, the impact that such variation would have on refinancing behavior is captured in state- and zip-code-level fixed effects in some of their specifications. Even after accounting for rate heterogeneity, they find that Black and Hispanic borrowers are less likely to refinance.

nation date (“static” variables) as well as variables that are updated over the life of the loan (“dynamic” variables). Among the variables included in the dynamic portion of the database are fields reporting the unpaid balance of the loan as of the reporting date, updated credit scores for the borrowers associated with the loan, the loan’s delinquency status, and information on whether the loan has been paid off; in our FTR analysis, we use the loan payoff field to classify borrowers as repaying the loan.

The information contained in the static portion of the database is incredibly rich and reports the Census tract in which the mortgaged property is located, the original loan balance, the number of borrowers on the loan, the value of the property at the time of origination, and a host of factors used in underwriting such as mortgage contract characteristics, the original interest rate, and borrowers’ credit scores and incomes. Most importantly for our study, unlike any other administrative mortgage servicing database, the NMDB files report the age, gender, race, ethnicity, and veteran status for each borrower on the loan.

To align our FTR-by-race analysis with the McDash results discussed above, we limit the NMDB records to conventional, 30-year, fixed-rate mortgages that were active between the first quarter of 2005 and the fourth quarter of 2019. As in the McDash analysis, we also exclude from the NMDB records: loans with exotic characteristics, such as balloon payments; loans without valid values for the variables used in our analysis; loans that were not current as of the observation date; and loans with an outstanding balance below \$75,000 as of the observation date.

We use the LTV and credit score reported in the static NMDB records to calculate risk-based mortgage rates over time in a very similar manner to the rate-estimation process described in Section 3, with two notable exceptions. First, the NMDB records report a borrower’s Vantage Score, while the McDash records report FICO scores, so we construct our NMDB risk segments using Vantage Scores. Second, the dynamic NMDB data is reported quarterly, in contrast to the McDash records which are updated monthly, so we define the adjusted mortgage rate for each Vantage Score-LTV segment as the median rate of the new originations in that segment on a quarter-by-quarter basis.³⁵

After constructing the segment-specific origination rates, we use the updated Vantage Score and LTV fields in dynamic NMDB records to assign a risk-based mortgage rate to each observation.³⁶ We then use the regression parameters estimated using the UCD data to estimate a loan’s closing costs based on the state in which the property is located.³⁷ We construct adjusted and Unadjusted FTR rates for the overall population as well as by race and ethnicity using the methodology described in Section 4. Lastly, to get a sense

³⁵The Vantage Score and LTV cutoffs that we use to create the risk segments in the LTV data are identical to those used to create the segments for the McDash data. FICO and Vantage scores are on the same scale and range from 300 to 850.

³⁶As in our analysis of the McDash data, when assigning a segment-level interest rate in the NMDB analysis, we use the lagged rate in that segment from the previous quarter.

³⁷See Section 2 for details on the closing cost estimation.

of the relative importance of the interest-rate and closing-cost corrections, we calculate the FTR rate using a segment-specific interest rate but the standard closing-cost estimate that ignores geographic heterogeneity in transaction costs (“FTR Rate-Adjusted Only”) as well as the FTR rate where we adjust for heterogeneity in closing costs but ignore interest rate variation (“FTR Cost-Adjusted Only”).

We summarize key variables and FTR rates by borrower race and ethnicity as well as for the overall NMDB sample in Table 3. Consistent with expectations based on previous research, in the NMDB sample average Vantage Scores are lower, and average LTVs are higher, for Black and Hispanic borrowers relative to Asian and White borrowers. Accounting for regional heterogeneity reduces the average closing-cost estimates for borrowers of all races, with the adjustment most pronounced for Asian and White borrowers. Relative to the unadjusted PMMS rate, the average adjusted mortgage rates are higher for borrowers of all races. The difference between the adjusted and unadjusted rates was most pronounced for Black (28 basis points) and Hispanic borrowers (25 basis points); these borrowers also had the highest average adjusted mortgage rates in our sample.

The Unadjusted FTR rates throughout the sample for Black (25 percent) and Hispanic (20 percent) borrowers are significantly higher than those of Asian (12 percent) and White (13 percent) borrowers. While our approach to analyzing refinance behavior by race differs from the methodologies utilized in previous work, the racial differences in FTR rates that we document are generally consistent with previous work on racial differences in mortgage prepayment (Kelly, 1995; Clapp et al., 2001; Deng and Gabriel, 2006; Firestone, Van Order and Zorn, 2007; Kau, Fang and Munneke, 2019; Gerardi, Willen and Zhang, 2023). Adjusting for closing cost and interest rate heterogeneity had a limited impact on the results for Asian and White borrowers, reducing the FTR rates for such borrowers by 2 percent and 3 percent, respectively. The differences between the Unadjusted and Adjusted FTR rates for Black and Hispanic borrowers, in contrast, were larger: the heterogeneity adjustment reduced the FTR rate for Black borrowers by 6 percentage points and reduced the FTR rate for Hispanic borrowers by 4 percentage points.

The net impact of the heterogeneity adjustment is a reduction in racial differences in FTR rates. For example, the Black-White Adjusted-FTR difference (9 percentage points) is 33 percent lower than the Black-White Unadjusted-FTR gap (12 percentage points), while the heterogeneity adjustment reduced the Hispanic-White FTR gap by 17 percent from 7 to 6 percentage points. Our findings thus imply that failing to account for systematic differences across borrowers in interest rates and closing costs can systematically overestimate racial gaps in refinancing behavior. Accounting for such heterogeneity, however, does not completely eliminate racial FTR gaps, a finding that is consistent with the results in Gerardi, Willen and Zhang (2023).³⁸

³⁸While similar in spirit, our methodology to quantifying racial differences in mortgage repayment behavior is different from that of Gerardi, Willen and Zhang (2023). Our FTR measures are, in essence, the percentage of borrowers who do not refinance conditional on the refinance option being “in the money.” Gerardi, Willen and Zhang (2023), in contrast, model racial differences in repayment behavior conditional on the value of the option, which may be positive or negative.

Table 3: Mean Values of Key Variables in NMDB Analysis

	All	Asian	Black	Hispanic	Non-Hispanic
	Borrowers			White	White
Vantage Score at Origination	742.16	751.22	706.74	724.08	745.45
Updated Vantage Score	735.99	748.86	698.71	716.26	739.30
Updated LTV	67.53	63.30	73.75	70.12	67.17
Unadjusted Closing Costs	3514.88	3892.20	3398.27	3503.27	3498.96
Adjusted Closing Costs	3072.92	3161.81	3154.16	3179.83	3052.40
Unadjusted Mortgage Rate	4.66	4.54	4.71	4.66	4.67
Adjusted Mortgage Rate	4.87	4.73	4.99	4.91	4.87
Unadjusted FTR	0.14	0.12	0.25	0.20	0.13
Adjusted FTR	0.11	0.10	0.19	0.16	0.10
FTR Rate-Adjusted Only	0.10	0.09	0.19	0.15	0.09
FTR Cost-Adjusted Only	0.15	0.13	0.26	0.21	0.14
Observations	37,223,789	1,973,045	1,937,798	2,753,259	30,559,687

This table summarizes key variables of our NMDB FTR sample, which is comprised of conventional, 30-year, fixed-rate mortgages reported in the NMDB that were active between the first quarter of 2005 and the fourth quarter of 2019. The sample excludes: loans with exotic characteristics, such as balloon payments; loans without valid values for the variables used in our analysis; loans that were not current as of the observation date; and loans with an outstanding balance below \$75,000 as of the observation date. These data restrictions are identical to those used to conduct the FTR analysis with the McDash data. The calculation of the unadjusted closing costs follows the existing literature and does not directly account for rate and closing cost heterogeneity; while the calculation of the adjusted closing costs adjusted for the rate and closing cost heterogeneity. The Unadjusted Mortgage Rate is rate reported in the PMMS. The Adjusted Mortgage Rate is the average rate on newly originated fixed-rate 30-year refinance loans in the pricing bin that corresponds to the borrower’s updated Vantage score and updated LTV.

Comparing the Unadjusted FTR measures with the “FTR Rate-Adjusted Only” and “FTR Cost-Adjusted Only” measures in Table 3, we see that the impact of the closing cost-adjustment in isolation *raises* the FTR rates by about 1 percentage point. The decline in racial gaps under the Adjusted FTR measure discussed above is thus driven by the interest rate correction. This finding is consistent with the analysis summarized in Section 5 where the closing-cost correction had little impact on FTR measures constructed using a pooled sample of loans from all U.S. states.³⁹

In addition to constructing the FTR measures aggregated over the full course of our sample, we also calculated FTR rates on a quarterly basis. Figure 12 displays the Unadjusted, Adjusted, Rate-Adjusted Only, and Cost-Adjusted Only series for borrowers of all races,

³⁹As discussed in Section 5, accounting for closing-cost heterogeneity had an important impact on FTR measures across states. It could thus be the case that the cost correction plays a more important role when loans are pooled by borrower race-state combinations.

and Figure 13 displays the Adjusted and Unadjusted FTR series by borrower race and ethnicity. The FTR dynamics summarized in Figure 12, which displays the FTR series for the NMDB sample, are very similar to the FTR series based on the McDash records shown in Figure 8. In both samples the heterogeneity adjustments typically reduce FTR measures, and the difference between the Adjusted and Unadjusted FTR series—which is driven primarily by the interest rate correction—is greatest when interest rates are declining.⁴⁰

The FTR series by race and ethnicity displayed in Figure 12 largely follow the same pattern as those of the overall sample. For borrowers of all races, FTR rates are very low when interest rates are flat or rising and rise significantly when rates begin to decline. While the Adjusted FTRs are almost always lower than the Unadjusted FTRs for borrowers of all races throughout the sample, the impact of the heterogeneity adjustment is most pronounced for Black and Hispanic borrowers when interest rates are declining rapidly. For example, when rates were declining rapidly in the third quarter of 2011, the Unadjusted FTR for Black and Hispanic borrowers exceeded the Adjusted FTR by more than 18 percentage points. In that same quarter, the heterogeneity adjustments were significantly smaller for Asian and White borrowers, reducing the Asian FTR rate by 8 percentage points and the White FTR rate by 10 percentage points.⁴¹ Our findings imply that while failing to account for the higher interest rates faced by Black and Hispanic borrowers imparts upward bias to FTR rates for such borrowers in all time periods, this bias is particularly acute when interest rates are declining rapidly.

7 Conclusion


In this paper we provide the first systematic investigation of regional heterogeneity in transaction costs in the U.S. mortgage market. The primary takeaway from this analysis is that there is a tremendous amount of variation across U.S. states in the relationship between loan amounts and closing costs. We also demonstrate that nearly all this variation is attributable to differences in mortgage taxation and government fees. The heterogeneity in closing costs that we document is thus driven by the decisions of state and in some cases municipal policymakers.

The magnitude of the regional heterogeneity in closing costs is economically significant. Based on the estimated parameters from Equation 1, it costs \$3,708 more to take out a \$500,000 mortgage in Florida (\$5,780), the state with the highest closing costs, than it

⁴⁰The readers is reminded that the McDash data is reported on a monthly basis while the NMDB data is reported on a quarterly basis. The FTR levels in Figure 12 and Figure 8 are thus not directly comparable.

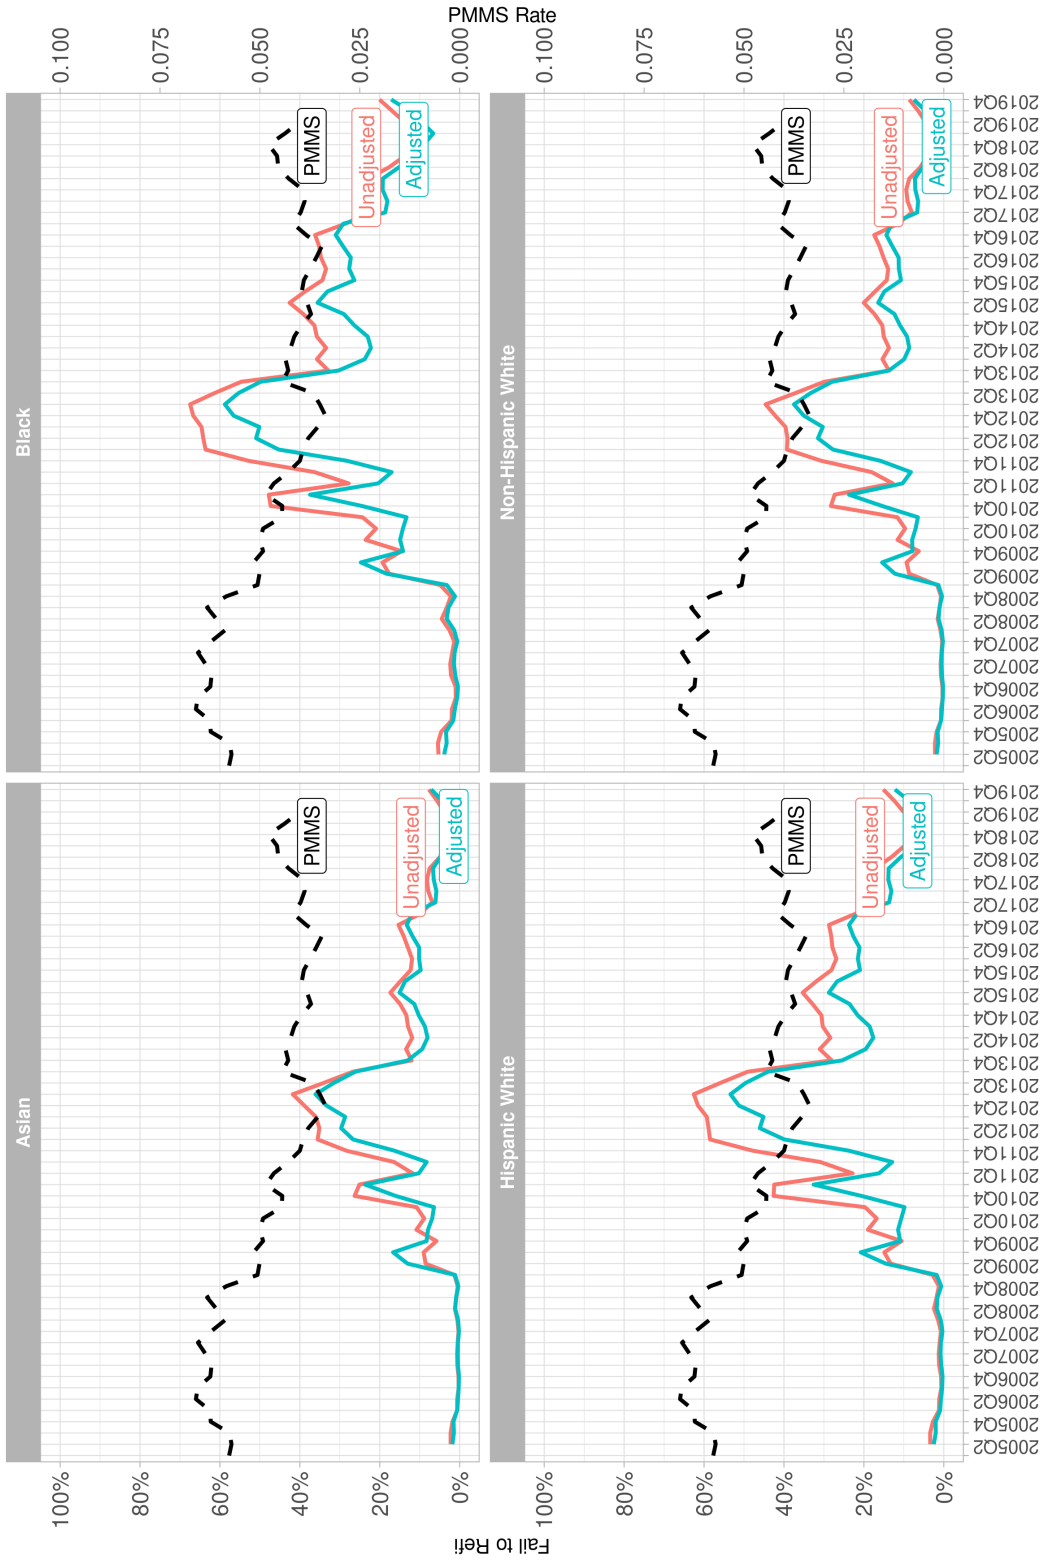
⁴¹In the third quarter of 2011, the Unadjusted (Adjusted) FTR rates for Black and Hispanic borrowers were 36.31 percent (17.02 percent) and 30.92 percent (12.91 percent). For Asian and White borrowers, the Unadjusted (Adjusted) rates in this quarter were 16.40 percent (8.26 percent) and 18.04 percent (8.23 percent), respectively.

Figure 12: FTR Measures for All NMDB Borrowers



This figure displays the Unadjusted, Adjusted, Rate-Adjusted Only, and Cost-Adjusted Only FTR rates for borrowers of all races using the NMDB sample.

Figure 13: FTR Measures by Race and Ethnicity



This figure displays the Unadjusted, Adjusted, Rate-Adjusted Only, and Cost-Adjusted Only FTR rates by borrower race and ethnicity using the NMDB sample.

costs to take out an identical loan in Iowa (\$2,071), which has the lowest closing costs. To put context around the magnitude of this difference, in their study of regional subsidies in the mortgage market, Hurst et al. (2016, p. 3009) produce estimates of the value of the transfers that occur between low- and high-risk states because of mispricing in the mortgage market: the 90th percentile of this transfer distribution was \$780. The Florida-Iowa difference in transaction costs is thus nearly five times larger than the implicit mortgage-rate subsidy to borrowers in the highest-risk markets in the U.S.⁴² Given the magnitude of the regional differences in transaction costs, our results provide a mechanism above and beyond the home equity channel discussed in Beraja et al. (2019) that could explain the documented geographic variation in refinancing activity and consumer spending during recoveries. Because variation in closing costs is largely attributable to differences in mortgage stamp and recording fees, our results suggest that state and local governments may be able to speed economic recovery when interest rates are declining by lowering such taxes.

After establishing these stylized facts about closing costs, we investigate how ignoring heterogeneity—in both closing costs as well as interest rates—could potentially impact studies of refinancing behavior using the failure-to-refinance (FTR) framework of Agarwal, Driscoll and Laibson (2013). This analysis reveals that failing to account for rate and cost heterogeneity can bias estimates of suboptimal behavior in the mortgage market in important ways. Furthermore, the nature of this bias depends on the borrower subpopulations under consideration. For example, when studying refinancing behavior by a borrower’s race, ethnicity, and risk profile, we find that ignoring mortgage rate heterogeneity creates a significant upward bias in the estimated incidence of suboptimal behavior for minority borrowers as well as borrowers with high-risk credit characteristics, such as high LTVs and low credit scores. Accounting for closing cost heterogeneity, on the other hand, has the most significant impact when studying regional differences in refinancing behavior. In light of the growing evidence of heterogeneity in both the mortgage market as well as the larger consumer credit market, these findings should serve as a warning that researchers must think carefully about how assuming away such heterogeneity might impact their results, especially when the focus of a study concerns between-group differences in behavior.

⁴²Interestingly, Hurst et al. (2016) find that many Florida borrowers receive large implicit transfers via low interest rates. One way of interpreting our findings is that Florida recaptures some of that subsidy via tax policies.

References

- Acolin, A. and R. Walter. 2020. “Loan Costs: Magnitude and Drivers of Variations.” *Working Paper* .
- Agarwal, S., J. Driscoll and D. Laibson. 2013. “Optimal Mortgage Refinancing: A Closed-Form Solution.” *Journal of Money, Credit and Banking* 45(4):591–622.
- Agarwal, S., R. Rosen and V. Yao. 2016. “Why Do Borrowers Make Mortgage Refinancing Mistakes?” *Management Science* 62(12):3494–3509.
- Alexandrov, A. and S. Koulayev. 2018. “No Shopping in the U.S. Mortgage Market: Direct and Strategic Effects of Providing Information.” *Consumer Financial Protection Bureau Office of Research Working Paper* .
- Andersen, S., J. Campbell, K. Nielsen and T. Ramadorai. 2020. “Sources of Inaction in Household Finance: Evidence from the Danish Mortgage Market.” *American Economic Review* 110(10):3184–3230.
- Avery, R., M. Bilinski, A. Clement, T. Critchfield, I. Keith, I. Mohamed, S. Patrabansh, C. Redmer, J. Schultz and R. Sullivan. 2021. “National Survey of Mortgage Originations Technical Documentation.” Federal Housing Finance Agency NMDB Technical Report 1.
- Beraja, M., A. Fuster, E. Hurst and J. Vavra. 2019. “Regional Heterogeneity and the Refinancing Channel of Monetary Policy.” *The Quarterly Journal of Economics* 134(1):109–183.
- Bhutta, N., A. Fuster and A. Hizmo. 2020. “Paying Too Much? Price Dispersion in the U.S. Mortgage Market.” *Federal Reserve Board Working Paper 2020-062* .
- Clapp, J., G. Goldberg, J. Harding and M. LaCour-Little. 2001. “Movers and Shuckers: Interdependent Prepayment Decisions.” *Real Estate Economics* 29(3):411–450.
- Deng, Y. and S. Gabriel. 2006. “Risk-Based Pricing and the Enhancement of Mortgage Credit Availability Among Underserved and Higher Credit-Risk Populations.” *Journal of Money, Credit and Banking* 38(6):1431–1460.
- Dhen, A. 2021. “CEMA Loans: Everything You Need to Know.” *Rocket Mortgage* . <https://www.rocketmortgage.com/learn/cema> Accessed 6/30/2021.
- Feinburg, R., D. Kuehn, S. McKernan, D. Wissoker and S. Zhang. 2012. “What Explains Variation in Title Charges? A Study of Five Large Markets.” U.S. Department of Housing and Urban Development Office of Policy Development and Research Report.
- Firestone, S., R. Van Order and P. Zorn. 2007. “The Performance of Low-Income and Minority Mortgages.” *Real Estate Economics* 35(4):479–504.

- Florida Statute 199.133*. 2020. The 2020 Florida Statutes. http://www.leg.state.fl.us/Statutes/index.cfm?App_mode=Display_Statute&Search_String=&URL=0100-0199/0199/Sections/0199.133.html Accessed 6/30/2021.
- Florida Statute 201.08*. 2020. The 2020 Florida Statutes. http://www.leg.state.fl.us/statutes/index.cfm?App_mode=Display_Statute&URL=0200-0299/0201/Sections/0201.08.html#1 Accessed 6/30/2021.
- Gerardi, K., P. Willen and David H. Zhang. 2023. “Mortgage Prepayment, Race, and Monetary Policy.” *Journal of Financial Economics* 147(3):498–524.
- Gurun, U., G. Matvos and A. Seru. 2016. “Advertising Expensive Mortgages.” *The Journal of Finance* 71(5):2371–2416.
- Hurst, E., B. Keys, A. Seru and J. Vavra. 2016. “Regional Redistribution through the U.S. Mortgage Market.” *American Economic Review* 106(10):2982–3028.
- Kau, J., L. Fang and H. Munneke. 2019. “An Unintended Consequence of Mortgage Financing Regulation—A Racial Disparity.” *The Journal of Real Estate Finance and Economics* 59:549–588.
- Kelly, A. 1995. “Racial and Ethnic Disparities in Mortgage Prepayment.” *Journal of Housing Economics* 4(4):350–372.
- Keys, B., D. Pope and J. Pope. 2016. “Failure to Refinance.” *Journal of Financial Economics* 122(3):482–499.
- McManus, D., L. Liu and Mingzhe. Yi. 2018. “Why Are Consumers Leaving Money on the Table?” Freddie Mac Insight Report. <http://www.freddiemac.com/fmac-resources/research/pdf/201804-Insight.pdf> Accessed 1/19/2021.
- PMMS Frequently Asked Questions*. 2022. The Freddie Mac Primary Mortgage Market Survey. <https://www.freddiemac.com/pmms/about-pmms> Accessed 10/27/2022.
- Prevost, L. 2013. “Reducing Refinancing Expenses.” *The New York Times*. <https://www.nytimes.com/2013/05/26/realestate/reducing-refinancing-expenses.html> Accessed 6/30/2021.
- Shroder, M. 2007. “The Value of the Sunshine Cure: The Efficacy of the Real Estate Settlement Procedures Act Disclosure Strategy.” *Cityscape* 9(1):73–91.
- Stanton, R. and N. Wallace. 1998. “Mortgage Choice: What’s the Point?” *Real Estate Economics* 26(2):173–205.
- Woodward, S. 2008. “A Study of Closing Costs for FHA Mortgages.” U.S. Department of Housing and Urban Development Report.

Woodward, S. E and R. Hall. 2012. "Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence." *American Economic Review* 102(7):3249–76.