A Demand Curve for Disaster Recovery Loans*

Benjamin L. Collier[†] & Cameron M. Ellis[‡]

October 10, 2023

Abstract

We estimate and trace a credit demand curve for households that recently experienced damage to their homes from a natural disaster. Our administrative data include over one million applicants to a federal recovery loan program for households. We estimate extensive-margin demand over a large range of interest rates. Our identification strategy exploits 24 natural experiments, leveraging exogenous, time-based variation in the program's offered interest rate. Interest rates meaningfully affect consumer demand throughout the distribution of rates. On average, a 1 percentage point increase in the interest rate reduces loan take-up by 26%. We find a large impact of applicants' credit quality on demand and evidence of monthly payment targeting. Using our estimated demand curve and information on program costs, we find that the program generates an average social surplus of $2,900 per borrower.

JEL Codes: **D12, G51, Q54**

Keywords: **Credit Demand, Household Finance, Financial Constraints, Climate Risk, Public Policy**

*Acknowledgments: We thank Benjamin Keys, Anthony DeFusco, Allen Berger, Asaf Bernstein, Ruiyuan Chen, Tatyana Deryugina, Meghan Esson, Emily Gallagher, Jesse Gregory, Umit Gurun, Sabrina Howell, David Keiser, Carolyn Kousky, Eli Liebman, Jeffrey Lin, Gregory Niehaus, Stephen Shore, Mohammad Soltani-Nejad, Justin Sydnor, Xuesong You, and seminar participants at the Federal Reserve Bank of Philadelphia, the University of Colorado Boulder, Temple University, the University of South Carolina, the University of Wisconsin-Madison, Florida State University, Bentley University, the University of Iowa, the Association for Public Policy Analysis and Management meeting, the Financial Management Association Annual Meeting, the ASSA-ARIA meeting, the Southern Finance Association meeting, the Western Risk and Insurance Association meeting, and the World Risk and Insurance Economics Congress for helpful comments and suggestions. The views, opinions, and/or findings contained herein are those of the authors and should not be construed as an official Government position, policy, or decision. The involvement of the Small Business Administration (SBA) is not an endorsement of the views, opinions, and/or findings of the authors. We thank the administrators at the SBA for their assistance in understanding the setting. All errors are our own.

[†]Fox School of Business, Temple University & Wharton Risk Management and Decision Processes Center (collier@temple.edu)

[‡]Tippie College of Business, University of Iowa (cameron-ellis@uiowa.edu)

1 Introduction

The costs of extreme weather events are large and growing. In the 1980s, 31 disasters in the U.S. exceeded a billion dollars in losses; in the 2010s, 128 billion-dollar disasters occurred (values in USD 2022, NOAA, 2022). Due to a combination of climate change and urbanization, this trend will continue. For example, flood losses in U.S. residential markets are projected to grow by as much as 60% over the next 30 years (First Street Foundation, 2021). Government intervention is often required to facilitate rebuilding after severe events, and in particular, many governments attempt to expand lending activity in affected areas. The most common policy tool is to offer or guarantee loans at subsidized interest rates to affected households or businesses.[1] Whether such policies are effective at spurring reinvestment depends crucially on the sensitivity of credit demand to interest rates in the wake of a natural disaster. However, credible estimates of credit demand elasticities are hard to come by and essentially non-existent in the case of natural disasters.

This paper estimates the interest rate elasticity of consumer demand for credit following extreme weather events in the United States. Our data are from the main U.S. program offering loans to households following natural disasters, the Federal Disaster Loan (FDL) Program. More than one million households applied to this program for a loan during our period of study, 2005 to 2018. These applicants span 70% of U.S. counties and 1,000 distinct disasters. Homeowners can borrow to repair damages to their property from a flood, hurricane, tornado, wildfire, etc. These damages are large: the median applicant in our data incurred $50,000 in uninsured damages. The median borrower takes a $24,000 loan with a 2.7% interest rate, 27 year maturity, and $148 monthly payment.

Using these data, we provide not only point estimates of the interest rate elasticity, but a nearly global demand curve for post-disaster credit. The FDL Program's offered interest rates update quarterly, and we use these rate changes to causally identify demand. We compare applicants who experienced a disaster in the days just before versus the days just after the rate adjustment to examine how changes in the interest rate, which is fixed for the life of the loan, affect whether approved applicants accept the loan. Our identifying assumption is that outside market conditions should be relatively stable during the few days on either side of the rate change. This strategy yields 24 natural experiments covering 20,000 individual borrowing decisions spread across 145 disasters. Each natural experiment allows us to identify a distinct point along the demand curve.

[1] Countries with disaster recovery loan programs include Australia (NSW, 2022), Canada (CMHC, 2022; Public Safety Canada, 2022), Japan (Japanese Finance Corporation, 2022), and the U.S. (FDL Program, 2022), among others. Several countries began providing recovery loans or expanded their programs during the COVID-19 pandemic including Germany (Kfw, 2020), Italy (MEF, 2020), the Netherlands (Government of the Netherlands, 2020), the U.K. (gov.uk, 2020), and the U.S. (SBA, 2022).

More generally, interest rates are one of the policy levers available to governments, making the elasticity of credit demand a parameter of interest. Despite its importance, identifying the effect of interest rates on consumer borrowing is challenging as interest rates co-move with borrower risk, loan contract terms, and economic conditions. Estimating credit demand over a broad range of interest rates is so difficult that only one paper has done it: Karlan and Zinman (2008) randomly assigned interest rates in a field experiment on microfinance loans in South Africa. For *long-term* lending contracts such as mortgages, assessments of credit demand effectively have been limited to local estimates (e.g., around the conforming loan limit, DeFusco and Paciorek, 2017).

We find a downward-sloping demand curve for post-disaster credit with an average semi-elasticity of -0.26: a 1 percentage point (pp) increase in the interest rate reduces take-up by 26%. For offered interest rates below the prevailing 30-year mortgage rate, the demand curve is effectively linear. Consumers respond to rate changes even when the offered rate is very low. For example, around 10% of approved applicants would accept a loan at 3 pp below, but not at 2 pp below, the prevailing mortgage market rates. When the offered rate is near the prevailing mortgage rate, the curve flattens somewhat, indicating that consumer take-up is more sensitive to the rate in this range. This increased sensitivity at higher rates likely reflects consumers' substitution toward outside options such as using savings or private credit. Regarding the *level* of credit demand, consumers are prone to reject rates at or below what the market would likely offer. Two-thirds of approved applicants would reject a recovery loan offered at the prevailing 30-year mortgage fixed rate; 40% would reject the loan at the prevailing 30-year Treasury rate.

While we would not necessarily expect credit demand elasticity to match across settings, the existing literature offers context for interpreting this rate sensitivity. Interest rate semi-elasticities are -0.03 for first mortgages (DeFusco and Paciorek, 2017), -0.1 for second mortgages and HELOCs, -0.4 for cash-out mortgage refinance loans (Bhutta and Keys, 2016), and around -1 for higher interest products like credit cards (Gross and Souleles, 2002) and microloans (Karlan and Zinman, 2019). Thus, our estimated semi-elasticity of -0.26 suggests that on average, consumers approach a recovery loan with a similar rate sensitivity as consumers considering a home-equity extraction loan.

Given the large damages that households incur, we would have expected credit demand to be more inelastic. First, the marginal utility of increasing consumption through borrowing is likely high at this moment as disaster damages can make portions of consumers' homes unusable (e.g., flooded living rooms, gaping roofs, burned-out kitchens) until they are repaired. Second, consumers respond to changes in the offered rate even when the rate is well below standard estimates of U.S. households' discount rates of around 4% (Gourinchas and Parker, 2002; Laibson et al., 2007). In standard models of credit demand, we would not expect consumers wanting a loan to reject it due to interest rate changes below their discount rate, yet we estimate a semi-elasticity of -0.16 on loans with interest rates below 4%. Finally, we would expect demand to be more inelastic because even if the offered interest rate exceeded the consumer's discount rate, the recovery loan is likely the lowest cost option. Using survey data of households affected by four recent major hurricanes, we show that households most frequently use savings to fund repairs. Similarly, Deryugina et al. (2018) find that households affected by Hurricane Katrina make early withdrawals on retirement savings. Yet, at the interest rates offered on recovery loans, taking the loan and keeping savings invested has a positive expected return. Moderate changes in the offered rate (e.g., increasing it from 2.7% to 3%) do not alter the conclusion that recovery loans are typically the lowest cost option. Overall, our findings suggest a reluctance among households to manage the shock through borrowing and that borrowing costs are central to this decision.

Toward understanding consumers' sensitivity to offered rates, we explore cross-sectional heterogeneity in the data. We find that demand is greater among consumers who would be deemed less creditworthy in private markets. Consumers with marginal credit scores, those with marginal debt-service-to-income (DTI) ratios, and consumers with high credit card utilization are all less sensitive to the offered interest rate. These differences in demand may result from pre-existing credit constraints or from differing attitudes toward borrowing (e.g., consumers with high DTIs may be more comfortable with debt).

We also examine potential wealth and cash-flow effects. Regarding possible wealth effects, we examine the size of the loss relative to the home's value as a measure of the wealth shock. We also examine household income. Neither of these appear to materially affect households' willingness to pay for credit, suggesting that the disaster-related reduction in lifetime wealth is not the key factor driving consumers' price sensitivity.

Regarding cash flows, we find that households who would have to commit a larger share of their discretionary income to servicing the loan have lower willingness to pay. These differences appear when we split the sample at the median where the new loan payment represents 5% of households' uncommitted monthly income. Our analyses do not clarify what drives these apparent cash flow effects, but potential explanations include differences in underlying risk (budget constraints are more likely to bind for some households), precautionary motives (preserving the ability to borrow in the future), or behavioral debt aversion. In sum, we find that post-disaster consumer credit demand seems to be attenuated by households' concerns regarding their monthly budget constraints and their access to substitutes.

As an alternative identification robustness check on our demand estimation, we exploit a discontinuity in the interest rate that is based on the applicant's credit score. Applicants with credit scores of at least 700 are more likely to be offered a loan with a "market" interest rate, which is approximately the prevailing average interest rate for 30-year, fixed-rate mortgages. Applicants with credit scores below 700 typically qualify for a "below-market" interest rate that is about half of the market rate. For example, the program's market and below-market interest rates were respectively 3.5% and 1.75% during Hurricane Harvey in 2017. Most approved applicants (87%) qualify for the below-market rate, which emphasizes that the program is largely a subsidized loan program. Using this credit score cut-off in a fuzzy regression discontinuity design, we estimate that households offered the market rate (typically about 2.5 pp above the below-market rate) are about 30 pp more likely to reject the recovery loan than those offered the below-market rate. This estimate is remarkably consistent with the demand curve using our time-based identification. For comparison with our main results, we also estimate a naïve specification using the full sample (instead of restricting the analysis to disasters occurring near a program rate change) and a Lasso estimation, which provides a more flexible treatment of model controls.

Finally, we assess the first-order welfare generated by the program. Consumer demand estimates are an essential ingredient in understanding the benefits of recovery loans; however, the key public policy metric is welfare, whether these benefits exceed program costs. We calculate consumer surplus from our estimated demand curve. We calculate costs from annual public records on the program's interest rate subsidies and administrative costs, also adding a 30% cost of raising funds through taxation.[2] Our estimates indicate that consumer surplus exceeds the programs' reported cost. We estimate that from 2005 to 2018, these recovery loans generated an average social surplus per borrower of $2,900, or $0.07 per dollar loaned.

Our findings add to the literature in several ways. First, we contribute to research on the management of households' exposure to weather and climate risks. Almost 43% of all U.S. homes, with a $6.6 trillion combined market value, are currently exposed to disaster risks (RealtyTrac, 2015). The literature examines the effects of increasing flood risk and sea level rise on coastal residential real estate markets (e.g., Bakkensen and Barrage, 2021; Bernstein et al., 2019; Keys and Mulder, 2020), the roles of the National Flood Insurance Program (e.g., Blickle et al., 2022; Sastry, 2022; Mulder, 2021) and disaster recovery grants (Fu and Gregory, 2019), and the effects of major events like Hurricanes Harvey and Katrina on households' finances (e.g., Gallagher and Hartley, 2017; Deryugina et al., 2018; Billings et al., 2022). Disaster recovery loans are a common policy tool intended to facilitate household rebuilding and reinvestment in the community after

[2] Our analysis focuses exclusively on first-order welfare: consumer surplus and program costs. It is possible that the program also generates second-order effects such as positive rebuilding externalities among neighbors or crowd-out of private sector lenders.

severe climate events (SBA, 2017, 2020). Spurring this reinvestment depends on loan take-up. Our findings suggest that the reach of recovery loans is limited by consumers' sensitivity to funding rebuilding through long-term commitments on their cash flows.

Second, we contribute to research on consumer credit demand. A series of papers estimate demand locally by exploiting exogenous interest rate variation for home loans (Bhutta and Keys, 2016; DeFusco and Paciorek, 2017), credit cards (Gross and Souleles, 2002; Ponce et al., 2017), auto loans (Argyle et al., 2020), and microcredit (Karlan and Zinman, 2008, 2019).[3] We estimate demand in a novel and economically important environment. By tracing a demand curve over a much larger range of rates than is typically possible, we find that post-disaster credit demand is non-linear: the elasticity is smallest at the lowest offered rates but increases around 1 pp below prevailing mortgage rates. The extent to which the shape of this curve extends to other credit markets is unclear, but it highlights the important, challenging task of incorporating point estimates from a range of interest rates when estimating credit demand elasticities.

Finally, we contribute to research on households' balance sheets and borrowing behavior (e.g., Mian et al., 2013; Aladangady, 2017; Di Maggio et al., 2017; DeFusco, 2018; Cloyne et al., 2019; Ganong and Noel, 2020; Guerrieri et al., 2020; Andersen and Leth-Petersen, 2020; Foote et al., 2020). A longstanding question is the extent to which consumers desire to smooth negative shocks over time through borrowing. At least some households manage unemployment shocks through credit, including mortgage refinancing (Hurst and Stafford, 2004) and sometimes payday borrowing (Bhutta et al., 2015; Keys et al., 2018). Yet, recent evidence shows that during unemployment spells, households may instead smooth credit card debt and bank overdrafts by adjusting consumption (Hundtofte et al., 2019). The shocks to households' balance sheets that we study are large relative to many of the events in the literature. In our setting, consumers' decisions to smooth shocks through borrowing depends crucially on the interest rate.

2 Setting and Data

This section describes the FDL program and our data, using material from FEMA (2019) and the program's Office of Disaster Assistance (2018).

[3] Additional examples include Alan and Loranth (2013), Attanasio et al. (2008), and Dobbie and Skiba (2013). Karlan and Zinman (2019) summarize estimates from the literature.

2.1 Federal Disaster Loan Program Overview

The FDL program began in 1953 and had made roughly $60 billion in recovery loans as of 2019. It is authorized to lend to households for the repair of uninsured damages to their primary residence, its contents (e.g., appliances, furniture), and their automobiles. The program is administered by the Small Business Administration (SBA) and also lends to businesses and non-profits, though the program predominantly lends to households. In 2017, households comprise 80% of applicants and 70% of the total loan volume. We limit our analysis to household lending.

Effectively all (98%) of household loans are associated with a presidential disaster declaration. For these declarations, the Federal Emergency Management Agency (FEMA) coordinates the local response, establishing temporary offices in affected neighborhoods. Households harmed by the disaster are encouraged to register with these FEMA offices. Households with incomes below a certain threshold (typically 125% of the federal poverty line) are referred to a FEMA grant program, which pays to repair or replace their lost property. FEMA refers households above the income threshold to the FDL program to apply for a loan. These households are then automatically contacted (via email, robocalls, and letters) by the loan program. Applying households who are declined for a loan are referred to the FEMA grant program for consideration.[4]

A household's eligibility depends on the issuance of a disaster declaration for its county, incurring a loss from the disaster, and some portion of the loss being uninsured. Figure 1 shows the geographic distribution of the program and illustrates its broad use across the contiguous U.S with an emphasis on the Gulf and South Atlantic coasts. The black areas in the figure denote ZIP codes that have at least one borrower in our data. We compare ZIP codes with applicants to other U.S. ZIP codes in more detail in Online Appendix G.

The program describes its purpose as alleviating credit constraints: "Disaster loans are an important part of the recovery process because they provide eligible homeowners, renters and businesses with access to the funds they need to rebuild" (SBA, 2017). It recognizes that recovery loans may additionally generate positive externalities: "Disaster loans are a critical source of economic stimulation in communities hit by a disaster, spurring job retention and creation, revitalizing business health and stabilizing tax bases" (SBA, 2020). Results of recent research seem to align with this assessment that rebuilding may generate positive spillovers. For example, Fu and Gregory (2019) study an unrelated local rebuilding program following Hurricane Katrina, finding that it benefited the neighbors of participants. Successful rebuilding efforts may also reduce bur-

[4] We compare FDL borrowers to FEMA grant recipients in Online Appendix G. We confirm that grants are used by a different, less affluent population: grant applicants and recipients are concentrated in the bottom quartile of the income distribution. Only 24% of grant applicants are approved, and grants are small ($4,500 on average).

Figure 1: ZIP Codes with FDL Applicants, 2005 to 2018

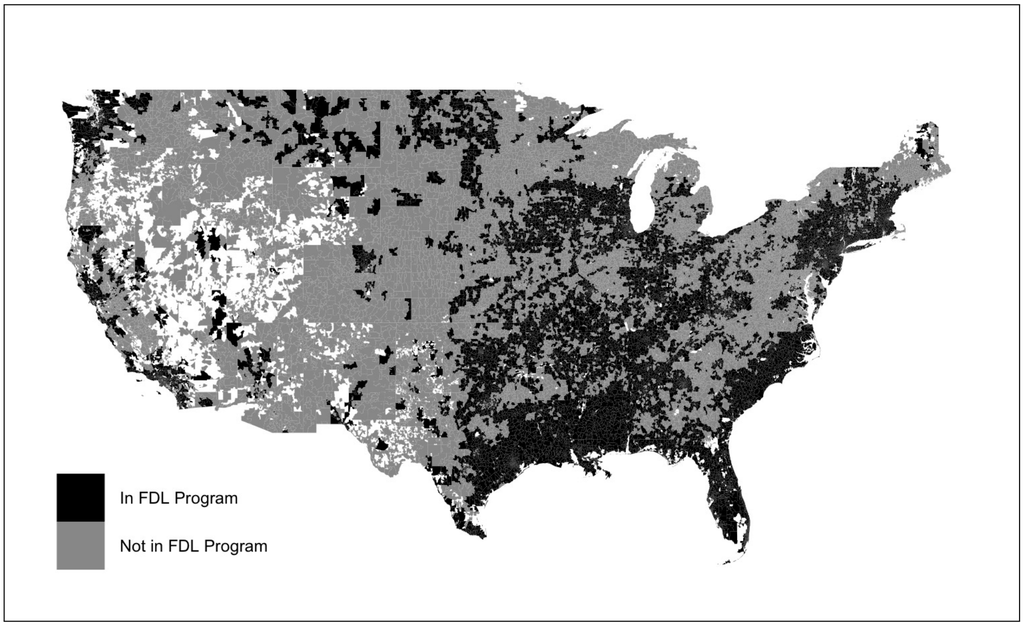

Note: Figure shows which ZIP codes had at least one applicant to the FDL program from 2005 to 2018.

dens on other public programs: Deryugina (2017) finds that spending on social safety nets such as publicly provided healthcare and unemployment insurance increases after major hurricanes.

2.2 Data, Lending Decisions, and Terms

Our data include all household FDL applications from 1 January 2005 to 31 May 2018 in the 50 U.S. states and the District of Columbia. During that time, more than one million households applied, and the program disbursed $12.5 billion in approved loans to 285,260 households. The data include a rich set of variables related to loan underwriting, the loss event, and the loan terms.[5] Table 1 provides summary statistics for households who completed an application, borrowers, and approved applicants who ultimately did not accept the loan. Here and throughout the paper, monetary values are in 2018 dollars (Federal Reserve, 2022b).

Lending decisions. The program is "a good faith lender and will only make a disaster loan if there is reasonable expectation that the loan can be repaid" (SBA, 2020). It collects information on an

[5] Our access to the program's data is through a data sharing agreement. The number of applications varies across years and is concentrated in large disasters (e.g., the 2005 and 2017 hurricane seasons). For example, a quarter of applications are from Hurricane Katrina, whereas 2014, a year with few large disasters, accounts for only 2% of all applications.

Table 1: Summary Statistics

	Mean	SD	Percentiles p10	p50	p90	Obs
			Completed Applications			
Income	71,338	75,102	21,207	54,341	128,153	791,637
Credit Score	632	107	497	625	789	872,127
DTI (%)	44	146	6	33	73	772,933
Loss Amount	89,332	117,487	9,978	46,951	221,816	596,555
			Borrowers			
Income	81,747	64,375	30,862	67,229	142,995	285,183
Credit Score	690	77	588	686	796	279,635
DTI (%)	33	96	7	31	54	281,623
Loss Amount	99,634	124,173	12,020	51,046	254,107	285,260
Home Equity	103,624	160,153	0	63,646	269,596	272,742
Insurance Claims	28,236	68,198	0	0	103,700	285,260
Loan Amount	46,050	60,302	8,931	24,175	121,390	285,260
Monthly Payment	238	246	57	148	529	285,259
Maturity (Years)	21	28	6	27	30	285,259
Interest Rate (%)	2.46	0.77	1.69	2.69	2.94	285,260
			Approved Applicants Who Cancel Loan			
Income	92,164	82,625	30,616	71,565	167,794	113,500
Credit Score	742	68	640	764	811	111,744
DTI (%)	27	43	4	24	47	111,922
Loss Amount	102,054	115,273	11,586	66,151	232,833	113,533
Home Equity	175,144	209,335	0	118,820	433,126	108,523
Insurance Claims	28,062	76,024	0	0	90,335	113,533
Loan Amount	60,598	59,908	7,683	39,992	143,482	113,533
Monthly Payment	348	324	94	261	706	113,533
Maturity (Years)	21	11	4	28	30	113,533
Interest Rate (%)	2.93	1.16	1.69	2.69	5.25	113,533

Note: Monetary values in $2018. Observations vary across variables because some applicants do not have credit reports or are declined early in the application process. "Income" is annual adjusted gross income. "Credit Score" is the FICO score of the primary applicant. "DTI" measures household's total monthly debt service payments (e.g., its mortgage) as a percent of monthly income. We report DTI for applicants with annual incomes of at least $1,000, which excludes about 0.1% of the sample. "Loss Amount" is the program's onsite assessment of property losses. "Home Equity" is the difference between the undamaged property value (determined during the onsite assessment) and the total balance on existing home debt (e.g., first and second mortgages, HELOCs) on the consumer's credit report at the time of application. Home equity is set to zero for non-homeowners. Income and home equity are winsorized at the 0.5% and 99.5% levels.

applicant's income from the IRS, outstanding debts from credit reports, insurance claims from its insurer, and property damages from an onsite loss inspection. Lending decisions largely depend on the interaction of the applicant's credit score and existing debt-service-to-income ratio (DTI).

While the rules vary over time, the program generally targets applicants with a credit score of at least 620 and an existing DTI below 40. Ultimately, 57% of households completing an application are declined by the program due to this underwriting process.

Table 1 describes the credit scores and DTIs of applicants and borrowers. The average credit score of FDL borrowers is 690, below that of Government-Sponsored Enterprise (GSE) borrowers, but around the national average. The average borrower has a DTI of 33, which is similar to GSE borrowers.[6]

Interest rates. Approved applicants are offered one of two interest rates, a "market rate" or "below-market rate." The applicant's rate depends on the disaster declaration date of the disaster that affected it. Thus, for a specific disaster the program offers a single market rate and a single below-market rate. Whether the applicant receives the market rate or below-market rate depends on an algorithm that accounts for the applicant's credit score, discretionary income, and wealth (Section 4.1). Applicants offered the below-market rate tend to have lower credit scores, higher DTI, and less wealth; however, because the assignment depends on several criteria, the distributions of characteristics for applicants offered the below-market rate overlaps with those offered the program's market rate.

Loan amounts. The program can lend up to $200,000 for damages to the residence and up to a combined total of $40,000 in damages to their contents and automobiles. While rarely used, households may be eligible to borrow additional funds for mitigation purposes or refinancing.

Collateral. The program does not make lending decisions based on borrower collateral. However, if the borrower has collateralizable assets, the program requires borrowers to secure their loans with collateral if the loan is above a certain amount (e.g., $25,000 as of 2018).[7]

Maturity. The structure of the recovery loan depends on slack in the household's budget. If necessary, the program will extend the loan up to 30 years to reduce the size of monthly payments. The program targets keeping the household's total monthly debt service, including the new loan, below 40% of its monthly income. The initial loan payment is due typically one year after the borrower accepts the loan.

The median borrower is approved for a loan of $24,176 with an annual interest rate of 2.7%, 27 year maturity, and amortized $148 monthly payment (Table 1). This loan payment represents a

[6] Specifically, the average U.S. FICO score was 689 in 2011 (the middle year of our data, Experian, 2020) and around 765 for the GSEs' mortgage borrowers (Fannie Mae, 2019b; Freddie Mac, 2019). The program's underwriting requirements are less stringent regarding both DTI and credit score than the GSEs Fannie Mae (2019a)

[7] Collier, Ellis, and Keys (2021) examine how the program's collateral requirements affect consumers' borrowing decisions. We discuss their findings in Section 3.2.

9% increase in the median borrower's monthly debt service.

Income. The median applicant has an annual income of $54,000 and lost $47,000 in the disaster (Table 1). Borrowers have higher incomes ($67,000). We additionally examine *relative* incomes in Online Appendix G. The median applicant has an income near the national median and that for its ZIP code and MSA. Applicants with above-median incomes are 30 pp more likely to be approved than below-median applicants in the same ZIP code. The difference in approval rates seems to result from underwriting based on credit scores and DTI, which correlate highly with income.[8]

Insurance coverage. On average, one quarter of a borrower's loss is insured, though the median borrower receives no insurance claims payments (Table 1). Gaps in insurance coverage are especially common for floods (e.g., about 70% of Hurricane Harvey-related flood damage was uninsured, Larsen, 2017). Many consumers, even those in very vulnerable locations, do not buy flood insurance (Walsh, 2017). Similarly, an insured household might have insufficient coverage: the National Flood Insurance Program has a maximum coverage limit of $250,000 on the home structure and does not tend to cover basements.

Approved applicants who cancel. About 28% of approved applicants do not accept the loan. They tend to have higher incomes, more home equity, higher credit scores, and lower existing DTI than borrowers (Table 1). The loan terms, a 2.7% interest rate and 28 year maturity, are similar to borrowers' terms at the median; however, the average interest rate is about 50 basis points higher for applicants who cancel.

2.3 Outside Funding Options

Why would a consumer fail to take a disaster loan after applying for it? Several sources of uncertainty could interfere with an applicant's ability to determine if they were willing to pay for the loan at the time of application. First, the applicant may be uncertain of the interest rate. The program provides a two-page, disaster-specific information sheet (Online Appendix B includes an example). The document reports the program's market rate and below-market rate, but includes insufficient details to determine which the applicant will receive. Second, the applicant may be unaware of the cost of repairs, which they will learn through the onsite loss inspection. Third, the applicant may be uncertain about the loan maturity, which depends on "borrower's ability to repay" with a maximum of 30 years. In sum, as applicants learn more about borrowing costs (e.g., the amount needed for repairs and the size of monthly payments), some approved applicants may

[8] Similarly, Billings et al. (2022) find that the allocation of federal disaster loans following Hurricane Harvey is regressive, and Begley et al. (2020) find that denial rates are higher in minority communities.

prefer to fund repairs through other means.

We use a survey to understand households' alternatives to a disaster loan. The survey includes 474 respondents who incurred damage to their home from one of four major hurricanes between 2017 and 2021.[9] Table 2 shows how households funded uninsured damages. The categories are not mutually exclusive and illustrate that consumers often use several sources to fund repairs. Half of the respondents reported drawing down on savings. This savings reduction aligns with Deryugina et al. (2018) who find that Hurricane Katrina significantly increased early withdrawals from individual retirement accounts.

Table 2: Funding Sources for Disaster Repairs

Funding source	%
Savings	51
Credit	
Credit cards	19
Formal loan from private bank or lender	8
Federal disaster loan	7
Transfers	
Family & friends	29
FEMA grant	19
Charity, nonprofit, or community group	10
Employer	8
Local government	5

Note: Table presents survey responses to the following question, "which of the following sources provided funds to help pay for the costs of repairing/rebuilding your home or for the costs of replacing items inside your home? (check all that apply)." Sample includes 474 respondents with home damage from four events: Hurricanes Harvey, Florence, Michael, and Ida.

Consumers using credit most frequently turned to credit cards, about a fifth of all respondents. Some credit card use may represent short-term, low-cost borrowing: del Valle et al. (2022) find that households affected by Hurricane Harvey open new credit cards at promotional rates and then pay off these balances before the promotion expires. The median credit card in their analysis has a $3,000 limit, suggesting an inability to fund large losses in this way. However, for other households, credit card usage may reflect borrowing at a high cost to fund repairs. For example, Morse (2011) and Dobridge (2018) find that disasters also increase payday loan borrowing.

[9] The studied events are Hurricane Harvey (n = 136), Hurricane Florence (n = 117), Hurricane Michael (n = 96), and Hurricane Ida (n = 125). The survey was primarily distributed through Qualtrics, which randomly sampled individuals in its internet panels who lived in disaster-affected areas. The survey was additionally distributed through (1) a geographically targeted Facebook ad campaign, (2) spots on local radio stations, and (3) community group outreach. Only individuals affected by the hurricane and who are the primary decision-maker in their household are included in the table. Participants were entered in a lottery to win gift cards valued at $20-30. See You and Kousky (2022) for additional details regarding data collection and results. We thank these authors for use of the data.

Regarding long-term loans, 8% of respondents funded repairs using a loan from a private lender. Similarly, 7% of respondents used an FDL. Respondents with more damages were more likely to use an FDL or private loan, indicating some composition effects. The results may also reflect consumers' inability to garner a loan and/or their reluctance to fund repairs through long-term borrowing.

The second most common category was transfers from family and friends, governments, non-profits, or employers. Twenty-nine percent of respondents received assistance from family and friends. Nineteen percent of respondents received a FEMA grant. These grants are typically small, averaging $4,500 during our period of study (Online Appendix G). Ten percent or fewer of respondents received assistance from a charitable organization or an employer or a local government.

In summary, savings, family and friends, and private credit (credit card or loan) appear to be the most frequent alternatives to a federal disaster loan. We generally would expect an FDL to provide cost advantages over these alternatives. Regarding savings, FDL applicants are typically offered a recovery loan at a sufficiently low interest rate that accepting the loan and keeping savings invested could generate a positive expected return at little risk. Similarly, a family member would likely be better off helping the household make recovery loan payments than providing a lump sum upfront; however, families may view a one-time intra-family gift differently from servicing a loan. Finally, because FDLs are subsidized typically, we would not expect a private loan to offer an interest rate that is competitive with the federal loan.[10]

3 Demand Estimation

We estimate the willingness to pay for recovery loans among approved applicants. Our analyses examine the extensive margin of credit demand, how interest rates affect whether an approved applicant accepts or rejects the offered loan. Figure 2, Panel A plots the two rates offered by the program over time against the rate for a 30-year fixed-rate mortgage (Federal Reserve, 2022a). The program's market interest rate is meant to reflect the prevailing interest rate; however, the program only adjusts the rate quarterly. We leverage these rate adjustments to piece together a demand curve.

We estimate demand for loans relative to the 30-year fixed-rate mortgage. That is, if the mort-

[10] Consistent with this expectation, survey respondents indicated that they found making payments on an FDL less burdensome than servicing a private loan or credit card debt (You and Kousky, 2022). Another alternative is to sell the un-repaired home. The media reports some instances of this, especially after Hurricane Harvey; however, the price discount on damaged homes appears steep (e.g., Putzier, 2019). As a result, households deciding to move would seem to benefit financially from borrowing to repair the home before selling it.

gage rate is 4.5%, the program's market rate is 5%, and the below-market rate is 2.5%; then we define the relative market rate as 0.5% and the relative below-market rate as -2%. Examining demand in relative rates is useful for considering potential substitutes for recovery loans. In particular, some households may be able to fund repairs through mortgage refinancing, something we examine explicitly in Section 3.3, and so may be especially prone to view the offered FDL rate relative to the prevailing mortgage rate.[11]

Several previous studies use discontinuity-based exogenous variation to estimate a full demand curve. Our estimation methodology is closest to Cohen et al. (2016) who estimate demand for Uber rides using discrete jumps in Uber's surge pricing algorithm. Another notable example is Finkelstein et al. (2019) who estimate health insurance demand using discontinuities in subsidy levels on Massachusetts' pre-ACA health insurance exchange.

Figure 2: Interest Rates Over Time

Panel A: Full Time Series

Panel B: One Window

Note: This figure plots the two interest rates offered by the FDL program over time and the average private market interest rate for a 30-year fixed mortgage. Panel A shows the full time series from 2005 through May 2018. The bolded, black vertical lines show the quarterly rate updates used in the estimation. Panel B shows an illustrative window, which includes two-weeks before and two weeks after Sept. 26, 2011, a date when the FDL program adjusted its rates.

[11] Besides interest rates, several additional considerations may influence a households' decision to use an FDL versus a private mortgage. FDLs lack closing costs and have less stringent collateral requirements, while private mortgages likely have fewer compliance requirements (e.g., loss verification and submitting receipts) and offer tax deductions for interest payments.

3.1 Time Identification

The program's quarterly interest rate adjustment provides a source of identification. Within a short window on either side of the quarterly change, unobserved conditions affecting credit demand should be stable while the program rate changes discretely. We bin the data into $windows$, two weeks before and two weeks after each observed rate change that has at least 20 approved applicants on each side. The bolded, black vertical lines show the rate updates used in the estimation. For the below-market rate, $windows$ span a 13 year period (2006 to 2018). For the market rate, the included $windows$ span a five year period (2006 and 2010) because of the restriction that each side of the window must include at least 20 approved applicants. Panel B of Figure 2 shows an example $window$ around the rate change on September 26, 2011. For example, consider an applicant who qualifies for the below-market rate. That household would receive a rate of 2.52% if it was affected by an event that was declared a disaster on September 25, but would receive a rate of 2.08% if it experienced an event that was declared a disaster on September 27, regardless of when the household applied or was approved.

We can use this setting to identify the causal effect of interest rates on loan take-up under two conditions: (1) the factors affecting credit demand are generally comparable during the two weeks before and after an FDL rate change and (2) the disaster declaration timing (pre vs. post rate change) is orthogonal to other, *unobserved* factors affecting credit demand. Toward assessing the latter condition, we test whether timing is correlated with *observed* factors that may affect credit demand by checking the balance of covariates around a rate change. This test is mainly concerned with whether political considerations related to this program appear to influence the timing of disaster declarations. Let the "lower-rate side" refer to the section of a $window$ in which the interest rate is lower. For example, in Panel B of Figure 2 the latter half of the $window$ is the lower-rate side as the quarterly adjustment reduced interest rates; however, in cases when the rate increases, the latter half of a window is the "higher-rate side." In Table 3, we aggregate our covariates by disaster and then compare the means for disasters that occur on the lower-rate side of a rate change to those on the higher-rate side. In addition to the political considerations mentioned above, we also examine covariates that may relate to systematic differences in applicants on the lower- versus higher-rate sides of the windows. For example, if more households apply when interest rates are low, their applications might receive less scrutiny, affecting who is approved. Moreover, anticipating less scrutiny, low-credit-quality applicants might be more likely to apply when the rate is low.[12] Table 3 shows that none of the covariates differ statistically, indicating

[12] While such systematic effects appear plausible, administrators indicate that program rules and operations are effectively held constant within a window, reducing the likelihood of some potential differences. For example, loan

Table 3: Disaster-Level Summary Statistics for All Rate Change $Windows$

	All		Lower Rate Side		Higher Rate Side		Diff. Means	
	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.	Diff.	t
Total Applicants per Disaster	403.26	931.34	399.59	1015	406.78	850.06	-7.19	-0.05
Approved Applicants per Disaster	141.06	351.50	134.77	351.42	147.08	353.87	-12.31	-0.21
Number of Loan Officers	46.30	51.20	46.30	53.49	46.30	49.28	-0.00	-0.00
30 Year Fixed Mortgage Rate	4.99	1.15	4.91	1.13	5.06	1.16	-0.16	-0.83
Monthly Income ($000s)	5.48	1.69	5.66	1.80	5.31	1.57	0.35	1.24
Credit Score	693.45	48.69	694.61	48.98	692.22	48.80	2.39	0.26
Loss Amount (2018 $000s)	87.73	95.44	96.75	101.92	79.08	88.61	17.67	1.11
Family Size	2.47	0.51	2.46	0.53	2.48	0.50	-0.02	-0.25
Monthly Fixed Debt	1573	676.12	1584	656.69	1563	698.59	20.98	0.19
Percent Renters	7.85	17.83	5.81	14.69	9.80	20.30	-3.99	-1.36
Home Value ($000s)	5.48	1.69	5.66	1.80	5.31	1.57	0.35	1.24
Home Equity ($000s)	87.73	95.44	96.75	101.92	79.08	88.61	17.67	1.11
N	145		71		74		-	

Note: Table presents summary statistics for disasters within $windows$ that are used in our demand estimation (i.e., 2 weeks on either side of a rate change with at least 20 approved applicants on each side). The first columns summarize all disasters that are in a $window$. The second set of columns includes only disasters on the lower side of a rate change. The remaining disasters, those on the higher side of a rate change, are represented in the third set of columns. The final columns provide two-way t-tests comparing the differences between the two sides of the windows.

that the $windows$ are balanced on relevant observed factors. Such aggregate comparisons may obscure differences, and we additionally include a large set of covariates that may affect demand as controls in our estimations.

We model the effect of a rate change on a household's acceptance probability linearly and control for observable measures of credit demand. Our identification strategy treats applications within a $window$ as randomly assigned based on whether the disaster affecting them occurred before versus after the rate change. Formally, we estimate the following:

$$P(Accept_{i,t}) = \alpha + f(rate_{i,t};\theta) + X_{i,t}\beta + \varepsilon_{i,t} \tag{1}$$

$$f(rate_{i,t};\theta) = \sum_{j=1}^{J} \theta_j 1\{window_{j,t}\} * rate_{i,t} + \theta_0 1\{i \notin J\} * rate_{i,t}$$

$$\begin{aligned} X = \{&\text{30-year Fixed Mort. Rate}, \text{30-year Fixed Mort. Rate}^2, \text{30-year Fixed Mort. Rate}^3, \\ &\text{Time}, \text{Time}^2, \text{Time}^3, \text{Credit Score}, \log(\text{Income}), 1\{\text{Income} \leq 0\}, \text{Loss Amount}, \\ &\text{Monthly Fixed Debt}, \text{Home Value}, \text{Home Equity}, \text{Renter}, \\ &\text{Loan Officers per Applicant}, \text{State}, \text{Year}\} \end{aligned}$$

officers work remotely from a national office or handful of regional offices. The same loan officer pool typically handles applications on both sides of a window. If the program expands the pool of loan officers to increase capacity, the new officers also would tend to handle applications on both sides of a window.

where our unit of observation is approved applicant i who suffered a disaster that was declared at time t. $Accept_{i,t}$ is a binary indicator for whether the approved applicant accepts or rejects the loan offered with relative rate $rate_{i,t}$. $window_{j,t}$ is a binary indicator for t being within two weeks of rate change j; J is the number of observed windows with at least 20 approved applicants on each side of the rate change. Similar to the demand estimation in Cohen et al. (2016), coefficient θ_0 captures the decisions of consumers who are in none of the windows. θ_0 is not well identified, and thus is not used in our demand curve estimation, but its inclusion increases the efficiency of the estimates for the control variables and so allows us to more accurately estimate $\theta_{1:J}$. Model controls include the raw level of the 30-year fixed-rate mortgage, credit scores, income, loss amounts, monthly fixed debt, the value of the home, whether or not the applicant rents, the total home equity, the number of loan officers per applicant for each disaster, and whether the loan is collateralized.[13] Additionally, we include state and year fixed effects and time (measured in days) trends. We model both time trends and mortgage rates using linear, quadratic, and cubic terms, which control for potential multi-window trends in the time series that could affect acceptance rates. The $\theta_{1:J}$ give local estimates of $\frac{\partial P(Accept)}{\partial rate}$ at many different price points based on exogenous price variation.[14]

We additionally estimate an average demand elasticity. While the global demand curve estimated from Equation (1) is our focus, the average demand elasticity is useful for comparison of our results with previous studies. For this average elasticity measure, we combine all of the time-based discontinuities. We limit our data to only those applicants within 2 weeks of a rate change and then use an indicator for whether the applicant was on the lower-rate side as an instrument for the rate they are offered. Being on the lower side of a rate change strongly impacts the interest rate the applicant is offered but, under the same identification assumptions as Equation (1), should have no other effect on an applicant's decision to accept or decline the loan. We then use the instrumented interest rate, via 2SLS, to estimate an average elasticity. Formally, we estimate

$$\begin{aligned} log(Rate_{i,t}) &= \alpha_0 + \alpha_1 1\{\text{Lower Rate Side}_{i,t}\} + \alpha_2 1\{\text{Below Market Rate}_{i,t}\} + X_{i,t}\gamma + \nu_{i,t} \\ P(Accept_{i,t}) &= \beta_0 + \beta_1 \widehat{log(Rate_{i,t})} + \beta_2 1\{\text{Below Market Rate}_{i,t}\} + X_{i,t}\theta + \varepsilon_{i,t} \end{aligned} \tag{2}$$

[13] We log income to reduce the effects of outliers. A few approved applicants have non-positive income (e.g., a self-employed applicant might have recently started her/his business). For non-positive incomes, we recode the data such that $\log(\text{Income}) = 0$ and include the indicator $\{\text{Income} \leq 0\}$. Variables in $X_{i,t}$ are measured at the time of the application, not the disaster declaration.

[14] The estimating equation incorporates a slight deviation from the typical regression discontinuity design, due to the structure of the data. While the typical design would include time as a running variable and the rate change date as the forcing variable, disaster declarations dates are too "lumpy" to do so. Because an applicant's interest rate is determined by its disaster declaration date, a $window$ may include hundreds or even thousands of applicants, but these applicants represent only a few distinct disaster declaration dates.

Table 4: Disaster-Type Summary Statistics for All Rate Change $Windows$

Disaster Type:	Total (N)	Total (%)	Low-Rate Side (%)	High-Rate Side (%)
Earthquake	1	0.7	0	1.3
Fire	15	10.3	12.3	9.1
Hurricane	6	4.1	5.5	3.9
Other Storm / Flooding	79	54.5	53.4	55.8
Tornado	43	29.7	27.4	29.9
Tsunami	1	0.7	1.4	0

Note: Table presents summary statistics disasters that occur within two weeks of a rate change. The full list of these disasters is in Online Appendix C.

where $1\{\text{Lower Rate Side}_{i,t}\}$, our instrument, is a binary indicator for applicant i being on the low side of a rate change within a $window$; $X_{i,t}$ is the same vector of control variables and fixed effects as in Equation (1); $1\{\text{Below Market Rate}_{i,t}\}$ is a binary indicator for if the applicant is offered the below market rate; and β_1 then gives the average, quasi-elasticity estimate.[15] This approach of stacking our time windows to estimate an average effect is similar to two previous studies that stack credit score discontinuities. Credit scores discretely affect credit card limits for Agarwal et al. (2018) and interest rates and maturities on auto loans for Argyle et al. (2020).

Most of the disasters in the *windows* are small: the median disaster has 19 applicants.[16] Table 4 presents summary statistics on the disasters in the *windows*. These include 145 disasters. Six are hurricanes; the most common events are un-named storms/flooding (76) and tornadoes (43). The disaster types are similarly distributed between the low and high rate sides. For example, 53% of the disasters are un-named storms/flooding on the low-rate side versus 56% on the high-rate side. One advantage of using these events is that they may be less prone to a broader set of economic spillovers that can accompany an extremely large disaster. For example, events like Hurricanes Katrina and Harvey can affect migration and local economies, though the effects on earnings and employment appear to be small and temporary (Deryugina et al., 2018; Billings et al., 2022).[17] Thus, the size and local nature of the disasters that we use for identification suggests an interpretation that consumers' credit demand is driven by their property damages.

[15] We use the gross interest rate, not the relative rate, in these regressions. The relative rates can be negative and thus their logged values are negatively infinite. The average semi-elasticity in Equation (1), which uses relative rates, is very close to the semi-elasticity estimated following Equation (2).

[16] For reference, Hurricane Harvey, which did not occur in the *windows*, had nearly 100,000 applicants.

[17] Katrina caused persistent, large out-migration but small reductions in household earnings, which disappeared within a year (Deryugina et al., 2018). In contrast, Harvey is associated with small, in-migration and no clear effects on earnings and employment (Billings et al., 2022). Billings et al. (2022) attribute the differences between these large disasters to unusual features of New Orleans during Katrina.

3.2 Estimation Results

We find an average semi-elasticity -0.26, a 1 pp increase in interest rates reduces take-up by 26%, using Equation (2).[18] This estimate is in the range of Bhutta and Keys (2016) who examine home equity extraction and find an extensive-margin, semi-elasticity for cash-out refinancing of -0.4 and for junior liens of -0.1. Household demand appears *less* elastic for disaster recovery loans than for higher interest credit products such as credit cards and microloans, which have estimated semi-elasticities from around -0.8 to -1.1 (e.g., Gross and Souleles, 2002; Karlan and Zinman, 2019). But demand for recovery loans is more elastic than that of first mortgages, which is near zero (DeFusco and Paciorek, 2017).

The *local* demand estimates, $\frac{\partial P(Accept)}{\partial rate}$ for each window in Equation (1), are presented in the first column of Table 5. The second column shows 95% confidence intervals estimated via percentiles from a traditional bootstrap and are *not* clustered.[19] The third column shows the identifying rate variation (relative to the 30-year fixed mortgage rate) with an average change of 23 basis points. Our slope estimates are reasonably stable across different relative interest rates. With a few exceptions, raising the interest rate by 1 pp would lower the acceptance rate by around 9-18 pp, with an average of 14.6 pp. Across the 24 different windows (6 for the market rate and 18 for the below-market rate), all of the point estimates are significantly below zero – an increase in price reduces demand in every estimation window.[20]

We convert these local demand estimates into a global demand curve along observed rates (Figure 3).[21] Following the conventions of demand curves, we place quantities (the likelihood that an applicant accepts the loan) on the horizontal axis and prices (the relative interest rate) on the vertical axis. Table 6 provides semi-elasticities by segments of the interest rate weighted by the

[18] We derive the semi-elasticity by first noting that the *elasticity* $e = (\Delta q/q)/(\Delta p/p)$ where q is the acceptance probability and p is the interest rate. Within windows, $q = 0.69$ and average $p = 0.028$. In Equation (2), we estimate the effect of a change in the logged interest rate, $log(\widehat{Rate}_{i,t})$, on the acceptance probability to be $\beta_1 = -0.50$. Let $\partial q/(\partial p/p) = \beta_1$. Together, $e = (\partial q/\partial p)(p/q) = -0.50/0.68 = -0.73$. We derive the *semi-elasticity* by re-arranging terms in the elasticity equation and solving for how a 1 pp change in prices affects take-up: $\Delta q/q = e \times \Delta p/p = -0.72 \times 0.01/0.028 = -0.26$.

[19] We bootstrap the estimation 100 times and then define our 95% confidence interval as the 2.5th and 97.5th percentiles of the bootstrap distribution. While we cluster by disaster in our other regressions, we do not do so here because each window includes a small number of disasters (5 on average). Clustering standard errors when the number of clusters is small can lead to problems (Cameron et al., 2008).

[20] As noted in Section 3.1, the market rate $windows$ occur during a shorter period (2006 to 2010) than the below-market rate $windows$. We examine whether the interest-rate responses for below-market rate $windows$ differed during this period, but find they are similar: an average $\partial P(Accept)/\partial rate = -0.118$ for the below-market rate over the full time series vs. -0.112 for 2006 to 2010. While this comparison does not guarantee a similar pattern for the market rate, it reduces concerns that interest-rate responses during this period are not comparable to the rest of the time series.

[21] Specifically, we construct a 5-degree polynomial approximation from the regression results in Table 5 via weighted OLS. The regression weights are the inverse of the estimates' standard errors, such that more precisely estimated windows carry more weight in the approximation. We then integrate the polynomial approximation and set the constant C in the integral so that the estimated demand model matches the observed mean rate and mean acceptance. Finally,

Table 5: Results for Time Identification

(1) $\frac{\partial P(Accept)}{\partial rate}$	(2) 95% Confidence Interval	(3) Rate Variation	(4) N for Lower Rate	(5) N for Higher Rate	(6) Date Cutpoint ($Window$)
-0.096	[-0.102, -0.091]	-3.905 – -3.843	736	3342	2006-06-30
-0.135	[-0.145, -0.126]	-3.933 – -3.683	391	58	2008-07-10
-0.102	[-0.111, -0.090]	-3.663 – -3.475	340	509	2006-04-01
-0.087	[-0.099, -0.075]	-3.495 – -3.433	64	477	2007-10-10
-0.104	[-0.114, -0.091]	-3.373 – -3.310	46	702	2006-10-04
-0.139	[-0.154, -0.124]	-3.163 – -3.100	34	361	2008-03-30
-0.123	[-0.139, -0.101]	-3.233 – -2.983	34	91	2009-07-01
-0.146	[-0.155, -0.137]	-2.518 – -2.455	1871	131	2010-04-02
-0.180	[-0.196, -0.149]	-2.560 – -2.247	74	20	2011-03-28
-0.121	[-0.136, -0.106]	-2.322 – -2.197	181	335	2016-01-02
-0.156	[-0.179, -0.121]	-2.242 – -2.055	65	86	2013-06-26
-0.097	[-0.114, -0.079]	-2.130 – -2.005	50	287	2017-07-03
-0.109	[-0.130, -0.084]	-1.997 – -1.935	33	681	2016-06-27
-0.149	[-0.178, -0.125]	-2.012 – -1.887	129	29	2015-04-02
-0.092	[-0.110, -0.078]	-2.070 – -1.820	93	667	2010-09-30
-0.141	[-0.150, -0.132]	-1.920 – -1.857	3717	36	2016-09-30
-0.130	[-0.141, -0.114]	-1.972 – -1.785	280	675	2012-07-02
-0.086	[-0.105, -0.070]	-2.027 – -1.590	31	1505	2011-09-26
-0.471	[-0.554, -0.385]	-0.870 – -0.620	37	210	2008-07-10
-0.167	[-0.227, -0.097]	-0.905 – -0.530	85	327	2006-06-30
-0.205	[-0.273, -0.104]	-0.920 – -0.420	95	88	2007-07-02
-0.178	[-0.274, -0.069]	0.025 – 0.525	139	263	2009-04-01
-0.138	[-0.330, 0.002]	0.170 – 0.420	839	52	2010-04-02
-0.180	[-0.237, -0.120]	0.680 – 1.180	32	125	2010-09-30

Note: This table presents the results of the Equation (1). 95% confidence intervals are calculated via percentiles from a bootstrap. Rate Variation represents the variation in rates on either side of the Date Cutpoint ($window$). Price changes are observed once the next disaster (following the quarterly change) occurs. The Date Cutpoint is the middle date of the two disasters where a rate change occurs. Disaster within 15 days on either side of this date are included within the $window$. There are 98 observed price changes, but only $windows$ with at least 20 people on each side are included. Columns 4 and 5 show the number of approved applicants within each window on either side of the rate change. This regression has a total N (which includes approved applicants who are not in a window) of 396,451, an F-Stat of 921, and an Adjusted R^2 of 0.20.

number of applicants in each window. Each segment represents approximately 5 (of the 24) local

we invert the approximated function to get the inverse demand function:

$$\theta_{1:J} = \frac{\partial P(Accept)}{\partial rate} = f(rate;\beta) = \beta_0 + \beta_1 rate + \beta_2 rate^2 + \beta_3 rate^3 + \ldots$$

$$P(Accept) = F(rate;\beta, C) = \beta_0 rate + \frac{\beta_1}{2} rate^2 + \frac{\beta_2}{3} rate^3 + \frac{\beta_3}{4} rate^4 + \ldots + C$$

$$C^* = \arg\min_C |P(\overline{Accept}) - F(\overline{rate};\beta, C)|$$

$$rate = F^{-1}(P(Accept);\beta, C^*)$$

Because our polynomial approximation of the derivative is strictly below zero, the integral is monotonically decreasing and thus (numerically) invertible.

estimates and 100 basis points of the interest rate.

Figure 3: Demand Curve

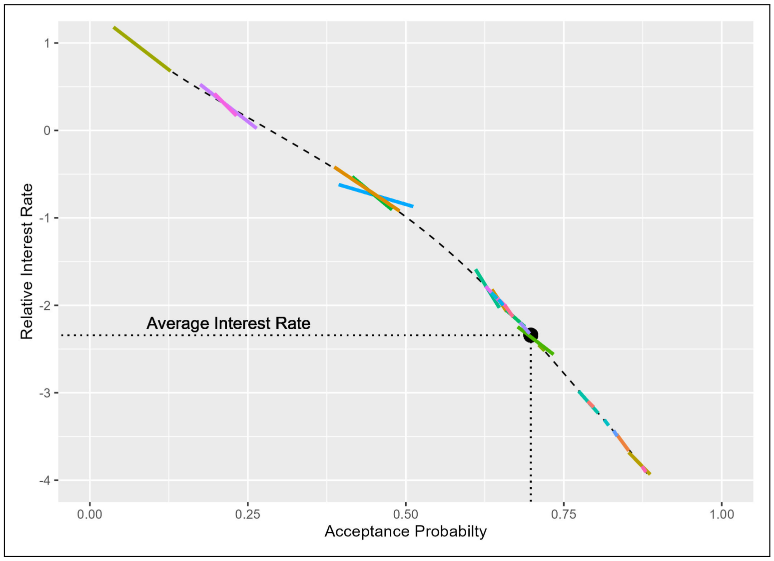

Note: This figure presents our local demand curves (each row of Table 5; solid, colored lines) as well as our estimated global demand curve (dashed line) for the SBA loans based on the interest rate, relative to the rate for a 30-year fixed-rate mortgage, holding all other factors that impact demand at their mean.

The demand curve is slightly concave overall, indicating that consumers are less sensitive to rates when they are low than when they are high. The slope is approximately linear from relative rates of -4 to -1 where the semi-elasticity is -0.18 (Table 6). For example, the figure shows that 80% of households would accept a loan at 3 pp below the prevailing mortgage market rate, but only 67% would accept a loan at 2 pp below the prevailing rate. No direct substitutes in private consumer credit markets exist for this range of interest rates. Thus, households who reject these loans when rates increase may be forgoing repairs or financing them through a less direct substitute such as drawing down savings or borrowing from family and friends.

The curve flattens somewhat as the relative rate approaches zero. For relative rates of -1 to 0, we estimate a semi-elasticity of -0.68. Households rejecting loans in this range may similarly be drawing down savings, delaying repairs, etc., but at these interest rates, some households may also be able to access more direct substitutes such as mortgage refinancing at a similar price.

Finally, the curve steepens slightly for the highest offered rates, for relative rates of 0 to 1.2 ($\partial P(Accept)/\partial rate$ = -0.17 vs. -0.27 for rates between -1 and 0, Table 6). We still find the largest semi-elasticity of -1.02 in this rate range. The combination of both the slope and semi-elasticity increasing is due to low take-up in this range: while most consumer reject the loan at these interest rates, those who accept the loan are relatively less sensitive to the price. Consumers willing to pay for a loan in this range of rates may have differing attitudes toward debt or fewer outside options; we explore heterogeneity in demand in the next section. Regarding the level of demand, the curve shows that *most* approved applicants (66%) would reject a recovery loan at the 30-year mortgage rate, typically one of the lowest interest rates available to households in private credit markets.

Table 6: Semi-Elasticity by Segments of Interest Rates

Relative Interest Rate Range	P(Accept)	$\frac{\partial P(Accept)}{\partial rate}$	Semi-Elasticity
[-3.9, -2.9]	0.83	-0.11	-0.13
(-2.9,-2.0]	0.71	-0.13	-0.19
(-2.0,-1.0]	0.58	-0.12	-0.21
(-1.0, 0.0]	0.40	-0.27	-0.68
(0.0, 1.2]	0.17	-0.17	-1.02

Note: This table shows the semi-elasticity of demand across different segments of interest rates. The acceptance probability (Column 2) and the change in acceptance probability due to the relative interest rate (Column 3) are generated from the main estimating framework (Equation 1) and reflect the average consumer by holding at their mean all modeled factors that impact demand besides the interest rate.

In addition to the decision of whether to accept the loan, consumers might adjust on the intensive margin, changing their loan amounts in response to interest rates. For example, Collier, Ellis, and Keys (2021) find intensive-margin responses using data from this program to study attitudes toward collateral. They examine the credit decisions of borrowers who incur damages around the size of the collateral threshold (e.g., set at $25,000 in 2018), finding that these borrowers reduce their loan amounts to avoid posting collateral. We additionally examined the effects of interest rates on *loan amounts* using our windows-based estimation strategy (Equations 1 and 2). We found an imprecise null result: the effect of interest rates on loan amounts was not statistically distinguishable from zero, but the model's standard errors are large. Thus, while our main analyses reveal consumers' strong extensive-margin response to the offered interest rate, we cannot rule out that consumers may also adjust on the intensive margin such that our main results may underestimate the total effect of interest rates on applicants' credit decisions.[22]

[22] We also examined whether approval likelihoods differ on each side of the windows, using Equation 1. The estimation includes the credit quality controls (e.g., credit score), but some model controls (e.g., loss amounts) are typically unavailable for declined applicants. An effect of interest rates on approval might be generated by two processes: 1) selection on unobservables – applicants who apply when interest rates are lower may differ in ways that are unobserved by the econometrician but that affect approval and/or 2) loan officer scrutiny may vary with the interest rate (poten-

We conduct several robustness tests, described in detail in Section 4. First, we estimate demand locally using an alternative identification strategy. The demand curve incorporates the interest rate responses from both below-market-rate and market-rate recipients. The key assumption is that after including model controls, the below-market-rate and market-rate recipients respond to the same interest rate variation similarly. To examine this assumption, we exploit a discontinuity in qualification for the below-market rate versus the market rate, which depends notably on whether the applicant's credit score is below 700. Using a fuzzy regression discontinuity design, we estimate a local demand curve for comparison with our main results. Second, while our windows-based estimation is tightly identified, it omits most applicants because it only includes disasters occurring within two weeks of the quarterly rate adjustment. We conduct a naïve estimation using the entire sample of approved applicants. Third, we repeat our main, windows-based estimating framework (Equation 1), incorporating controls more flexibly using a Lasso estimation, which includes all continuous control variables up to the fifth power and their logs and allows the algorithm to select the most predictive ones. This analysis is motivated by the possibility that the functional forms of the controls may be mis-specified, affecting our demand estimates. In each of the three tests, the estimated demand curve is largely similar to our main estimates.

It is interesting from the demand curve that 8% of approved applicants would reject the loan at the lowest observed relative rate of -4. One possibility is that these consumers may be willing to pay for a loan at a lower interest rate than what is observed in the data. Alternatively, these consumers might reject the loan regardless of the offered rate. For example, consumers may misperceive the full cost of repairs when applying, but learn more about these costs during the application process. If such non-interest-rate motivations for rejecting the loan are *orthogonal* to interest rates, they would not affect the slope of the demand curve; however, they would affect the level of demand, influencing the estimated elasticities. Omitting these never-takers from the elasticity estimates seems desirable from a policy perspective: if we assume that 8% of all consumers would reject the loan regardless of the rate, the average semi-elasticity is -0.24.[23] Instead, non-interest-rate factors affecting take-up could *correlate* with the interest rate, potentially influencing the shape of the demand curve. For example, certain populations such as low-credit-quality borrowers might be prone to misperceptions of repair costs and more likely to apply in certain ranges of the interest rate. We have taken three steps to mitigate the effects of this possibility. We include a

tially due to variation in case loads), affecting approval rates. We do not find a clear effect of interest rates on approval probability: across the windows, the coefficient on the interest rate is typically insignificant and its sign flips irregularly. It is possible that the relationship between rates and approval likelihood is most relevant in a subset of windows (e.g., those with the lowest interest rates) and our tests are under-powered for detecting an effect for this subset.

[23] We adjust the elasticity by netting out 8% from the average acceptance rate q. Since $q = 0.68$, the net acceptance rate $q' = 0.68/0.92 = 0.74$, resulting in an elasticity of $e' = -0.50/0.74 = -0.67$ (vs. the unadjusted $e = -0.73$) and semi-elasticity of -0.24 (vs. an unadjusted of -0.26).

rich set of controls for observable characteristics of applicants including credit quality in the main and Lasso specifications. We report semi-elasticities by segment of the interest rate (Table 6).[24] Finally, we examine heterogeneity in demand along a number of dimensions in the next section.

3.3 Heterogeneity in Demand

We next examine how demand varies across subsets of the sample, considering factors related to credit quality, wealth effects, and liquidity. These dimensions are not mutually exclusive and our analysis of them is descriptive, not causal. For example, household liquidity and wealth are likely both related to household income, among other factors.

We present the average elasticity estimates by group in Table 7. The first column shows the average elasticity (as reported in Section 3.2). The first and second rows show the coefficient estimate and standard errors. The third row shows the elasticity implied by the coefficient and average acceptance rate. The p-value reports whether the difference in elasticities between the compared populations is statistically significant. We focus our discussion on the interest-rate elasticity (instead of the semi-elasticity) as a unit-free measure of demand because the average offered interest rate sometimes varies within a sample split (e.g., the offered rate differs by about 35 bp for high vs. low income consumers); however, the table also includes semi-elasticities for comparison. Figure 4 plots the demand curves corresponding to Table 7 for a subset of the sample splits.

The first set of columns compares applicants based on measures of credit quality, splitting the sample at the median by DTI, credit score, and credit card utilization.[25] We find strong evidence that demand varies with credit quality. The median approved applicant has a DTI of 29%. Households with lower DTI have elastic demand (-1.2), a 1% increase in the interest rate reduces take-up by 1.2%, while those with higher DTI have an elasticity of -0.5. Panel A of Figure 4 shows a level shift in demand: at any given interest rate, the higher DTI consumers are more likely to accept the loan. The figure also shows that the difference in slopes between low and high DTI households

[24] For example, supposing these problems are worst at the lowest offered rate, we can compare the elasticity estimates at the lowest offered rates to those at slightly higher rates. Table 6 shows that the estimated semi-elasticity in the lowest range of rates (-4 to -3) matches those in the next two bands of the interest rate (-3 to -2 and -2 to -1). Similarly, we find the same semi-elasticity in the lowest interest rate range (-0.13) if we omit the window with the lowest observed interest rates in the data, reducing concerns that unobserved selection may be problematic especially for our estimates at the lowest observed rate.

[25] The median splits are based on values for applicants in the windows sample. We observe the consumer's credit report at the time of application. We restrict our utilization measure to cards that were already open on the disaster declaration date. High card utilization is typically a proxy for liquidity/credit constraints (e.g., Agarwal et al., 2010; Gross and Souleles, 2002), but del Valle et al. (2022) also find increased utilization among (better off) consumers who open new cards at promotional rates and pass spending through them following Hurricane Harvey. Our utilization measure may still reflect disaster-related purchases; however, restricting our assessment to existing cards seems to better connect the analysis to research examining utilization as a proxy for constraints.

Table 7: Average Demand Elasticity By Group

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Global	Debt Service to Income		Credit Score		Credit Card Utilization	
		Low	High	Low	High	Low	High
Stacked RD Coefficient	-0.50	-0.70	-0.42	-0.48	-0.71	-0.61	-0.43
	(0.05)	(0.10)	(0.07)	(0.22)	(0.07)	(0.09)	(0.16)
Implied Elasticity	-0.73	-1.21	-0.54	-0.58	-1.28	-1.04	-0.55
Split (Median)	-	0.29		704		0.27	
Difference	-	0.68		-0.69		0.50	
(p-value)	-	(0.00)		(0.02)		(0.05)	
Implied Semi-Elasticity	-0.26	-0.40	-0.21	-0.24	-0.40	-0.35	-0.21
Average Acceptance Rate	0.68	0.57	0.79	0.83	0.56	0.58	0.78
Average Interest Rate	2.81	3.04	2.57	2.39	3.16	3.01	2.60
	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
	Enough Extractable Equity?		Loss Size Relative to Home Value		Monthly Income		Loan Payment Relative to Disc. Income	
	No	Yes	Low	High	Low	High	Low	High
Stacked RD Coefficient	-0.54	-0.62	-0.43	-0.52	-0.63	-0.51	-0.32	-0.56
	(0.09)	(0.13)	(0.24)	(0.46)	(0.09)	(0.12)	(0.15)	(0.32)
Implied Elasticity	-0.75	-0.99	-0.62	-0.76	-0.92	-0.74	-0.42	-0.96
Split (Median)	-		0.41		4998		0.05	
Difference	-0.24		-0.14		0.18		-0.55	
(p-value)	(0.32)		(0.85)		(0.41)		(0.34)	
Implied Semi-Elasticity	-0.27	-0.34	-0.21	-0.31	-0.35	-0.25	-0.15	-0.34
Average Acceptance Rate	0.72	0.63	0.69	0.68	0.68	0.68	0.77	0.58
Average Interest Rate	2.76	2.91	2.97	2.47	2.64	2.98	2.75	2.87

Note: This table shows the results of the average price elasticity estimate, following Equation 2. The first column shows the aggregate average elasticity. The remaining columns divide the data by DTI, credit score, card utilization on credit cards that were open before the disaster, a binary indicator for the availability of home equity to cover the loan amount, the loan size as a percentage of the applicant's initial home price, monthly income, and the monthly loan payment as a percentage of the applicant's discretionary income. The first row shows our estimate of the slope coefficient. Standard errors, in parentheses, are clustered at the disaster by rating category level. The third row shows the elasticity implied by the coefficient and average acceptance rate. The fourth row shows the value at which the sample is split. The fifth row shows the difference in implied elasticities for each split. The sixth row shows the p-value for the equality of implied elasticities. The remaining rows respectively show the implied semi-elasticity, average take-up rate, and average offered interest rate for each subgroup.

grows as the FDL rate approaches the 30-year mortgage rate. In particular, high DTI households appear completely inelastic for relative rates in the range of 0 to 1.

We find a similar result regarding credit scores: less creditworthy households are both more likely to accept the loan and are less sensitive to the rate. The median approved applicant has a credit score of 704. Applicants with higher credit scores have an elasticity of -1.3; those with lower scores have an elasticity of -0.6. We are unable to estimate a full demand curve for below-median credit score households due to sample size.[26] For the parts of the curve that we can estimate, the demand of lower credit score households is shifted to the right of households with higher credit scores in a similar fashion to DTI (Panel B of Figure 4). Regarding utilization, we find a qualitatively similar result: consumers with low card utilization have unit-elastic demand while high-utilization consumers have an elasticity of -0.55. Like the low-credit-score sample, we cannot estimate the full demand curve for the high-utilization households (Panel C of Figure 4) but again find a level shift showing that higher utilization applicants are more likely to accept the loan.

Second, we examine collateral constraints by considering whether a household has sufficient home equity to finance the disaster loss. We measure home equity by subtracting the total current balance on the applicant's home loans (first mortgage, second mortgage, home equity lines, etc.) from the home's pre-loss value. This home value is from the program's onsite inspection, which determines what the value of the home would have been had it not been damaged. We categorize applicants as having "enough extractable equity" if this home equity measure exceeds the disaster loss amount. We also include homeowners without a mortgage as having *sufficient* extractable equity and categorize applicants who are not homeowners as having *insufficient* extractable equity. Several studies examine home equity extraction and conclude that households generally use extracted equity for real outlays, including consumption expenditures and making home improvements (e.g., Mian and Sufi, 2011; Bhutta and Keys, 2016; DeFusco, 2018).

We do not find a statistical difference in average demand based on sufficient extractable home equity (Columns 8 and 9 of Table 7); however, this average effect masks differences across the demand curve, shown in Panel D of Figure 4. When the FDL interest rate is below the 30-year mortgage rate, the price sensitivity for applicants with enough extractable equity is highly similar to those without. However, once the FDL rate crosses the private mortgage rate, the curves diverge such that applicants with enough home equity to finance repairs are more price sensitive. For example, when a recovery loan is priced 50 bp above the mortgage rate, 9% of households with enough equity to cover the loss accept the loan versus 43% of households *without* enough equity. These results suggest that applicants with sufficient home equity substitute toward private home loans such as mortgage refinancing when FDL rates are high.

Next, we examine demand by relative loss size and income, but do not find meaningful effects. To proxy wealth effects of the loss, we measure loss size as the amount of damages relative to the home's value. We do not find an effect of relative loss size on the interest rate elasticity (Columns 10 and 11). Income is a policy-relevant metric and potentially connected to all of the mechanisms

[26] The market rate windows do not have at least 20 approved applicants with below-median credit scores on each side.

considered here; however, we do not find a difference in the elasticity when splitting by median income (Columns 12 and 13). One potential explanation for this null is that income combines competing mechanisms as lower income households may have lower credit quality but also be less liquid.

We also examine the impact of monthly cash flows on demand. Our measure of cash-flow liquidity is the ratio of the monthly recovery loan payment to the household's monthly discretionary income (income minus its existing debt payments). At the median, monthly recovery loan payments represent 5% of a household's discretionary income. Less liquid households are less likely to accept the loans, a difference of almost 20 pp in the average acceptance rate. This is consistent with borrowers targeting monthly payments. The difference appears most prominently in Panel E. Regarding slopes, the demand of the less liquid households (Column 15) is unit elastic, versus -0.4 for more liquid households. The estimated difference in elasticities is large but not statistically significant due to substantial variation among less liquid households.

In sum, credit quality and relative monthly payment size appear important in explaining heterogeneity in demand. This heterogeneity is useful for assessing price sensitivity but also targeting, understanding what populations are likely to fund rebuilding with long-term recovery loans. These population differences may reflect a combination of market-based factors and preferences. For example, high-credit-quality households likely have access to lower cost private credit; however, they remain price sensitive at very low interest rates, which do not appear to have a direct substitute in private credit markets. These consumers' behavior may reflect a greater aversion to additional debt and/or higher psychological or hassle costs of borrowing.[27] Andersen et al. (2020) find that psychological/hassle costs reduce the likelihood that busier households (e.g., better off) refinance. Interestingly, we find little effect of income on demand for recovery loans, potentially pointing toward behavioral or preference-based explanations such as debt aversion. Our results regarding demand heterogeneity in monthly payment size connect with findings in other credit markets: auto loan borrowers (Argyle et al., 2020) and microfinance borrowers (Karlan and Zinman, 2008) engage in payment targeting, and monthly liquidity constraints appear central to explaining mortgage default (e.g., Ganong and Noel, 2020). Recent research on how households manage liquidity highlight a combination of financial frictions and behavioral factors that may contribute to the demand heterogeneity that we observe (e.g., Baugh et al., 2021; Ganong and Noel, 2019; Olafsson and Pagel, 2018).

[27] While applicants have already incurred some hassle costs by applying, accepting the loan might create additional hassles such as maintaining documentation of repair costs and servicing an additional loan.

Figure 4: Demand Heterogeneity

Panel A: DTI

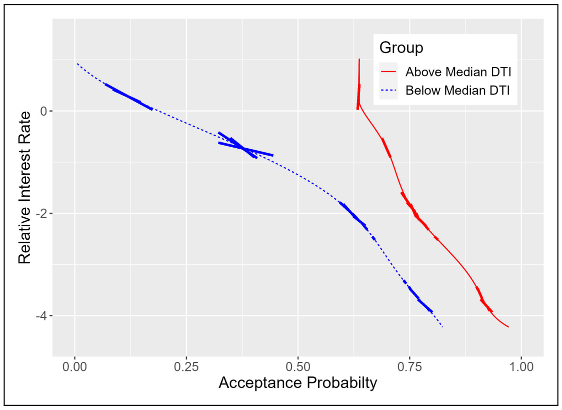

Panel B: Credit Score

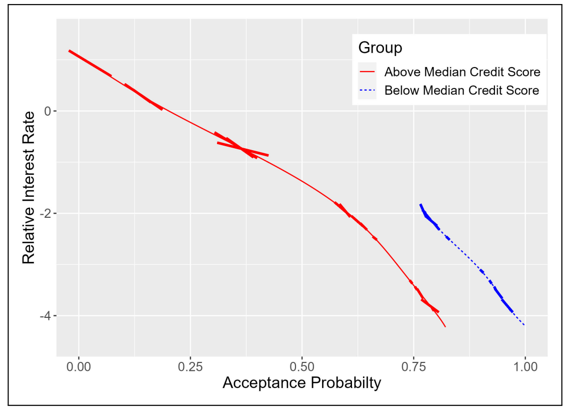

Panel C: Credit Card Utilization

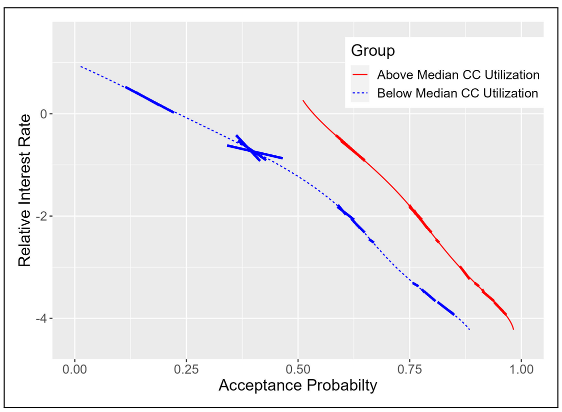

Panel D: Extractable Equity

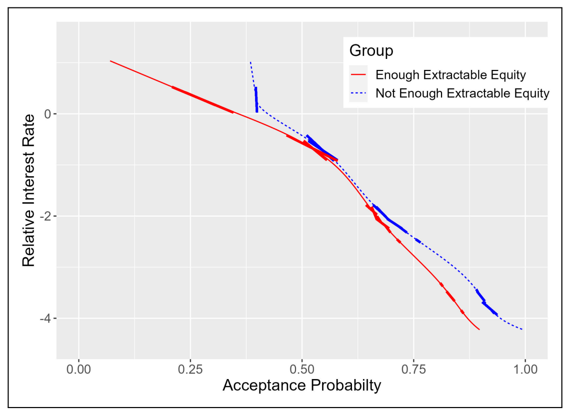

Panel E: Payment / Disc. Income

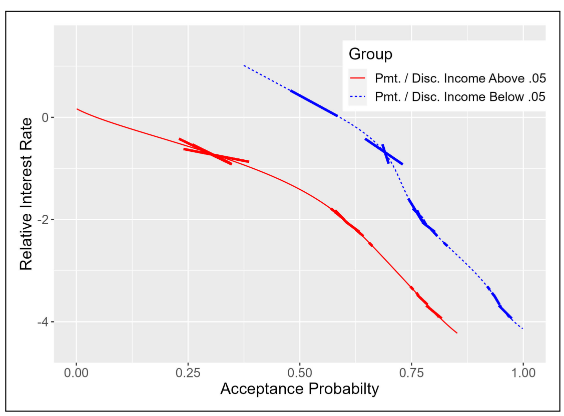

Note: Figures compare demand based on DTI, credit score, card utilization, whether the household has sufficient home equity to cover the disaster loss, and the share of discretionary income taken up by the loan payment. Estimations follow Equation (1) and interact rate changes with a binary indicator for the subgroups in each panel (e.g., in Panel D an indicator for extractable equity). Windows that do not contain at least 20 applicants on each side of the window are removed (e.g., this removes all of the market rate windows for the below-median credit score group).

4 Robustness and Extensions

This section includes two supporting analyses. The first is a robustness test in which we use a discontinuity in applicants' credit scores to estimate credit demand for comparison with our estimates derived in Section 3. Additionally, we conduct a welfare analysis, which compares consumer demand to the cost of supplying recovery loans.

4.1 Robustness: Demand Estimation and Credit Score Discontinuity

We exploit a separate form of identification, qualification for the below-market rate, as a robustness check on our demand estimates. In particular, we are concerned with whether, after including model controls, the below-market-rate and market-rate recipients respond similarly to the same interest rate variation and can thus be combined into a single demand curve. Applicants are offered the below-market interest rate when the program concludes they cannot access credit elsewhere. Credit score is a key criterion: the applicant is deemed as having limited access to private credit if the primary wage earner on the application has a credit score below 700.[28] Panel A of Figure 5 shows the discontinuity in offered interest rate at a credit score of 700. Around 95% of applicants with a credit score of 699 qualify for the subsidized, below-market rate. Increasing the applicant's credit score by one point triples the likelihood that the applicant is instead offered the market rate. Households at the cutoff should differ only in the average prices they are offered. Thus we can exploit this discontinuity to assess local demand among market-rate and below-market-rate recipients who are, otherwise, effectively identical. If the local demand curve for this marginal household closely matches both the sections of the global demand curve that are estimated using variation in the below-market rate as well as the sections that are estimated using variation in the market rate, then it provides support that these sections can be combined in an aggregate demand curve.[29]

We estimate a fuzzy regression discontinuity design. We follow standard practice and use a two-stage least squares approach (Hahn et al., 2001; Lee and Lemieux, 2010). In the first stage, the threshold is used as an instrument for the interest rate to which the household is assigned.

[28] The credit elsewhere test also considers discretionary income and net worth: (1) The applicant's discretionary monthly income is greater than 1/3 of the disaster loan payment structured using the market interest rate and a 15 year loan term, and (2) The applicant's adjusted net worth is over four times the uncompensated disaster loss. The adjusted net worth is calculated as total assets minus total liabilities minus $100,000. Applicants who meet at least two of the three requirements are deemed as having access to credit elsewhere and offered the program's market rate.

[29] Consumers could plausibly manipulate their credit scores, for instance by applying for new credit cards, in the 60 days between the disaster and the deadline for loan application. To check for this, we run the McCrary (2008) sorting test and fail to find evidence of sorting (Figure E1 in Appendix E).

Figure 5: Alternative Demand Estimations

Panel A: Fuzzy Discontinuity First Stage

Panel B: Demand Curve: Fuzzy Discontinuity

Panel C: Demand Curve: Full Sample

Panel D: Demand Curve: LASSO Covariates

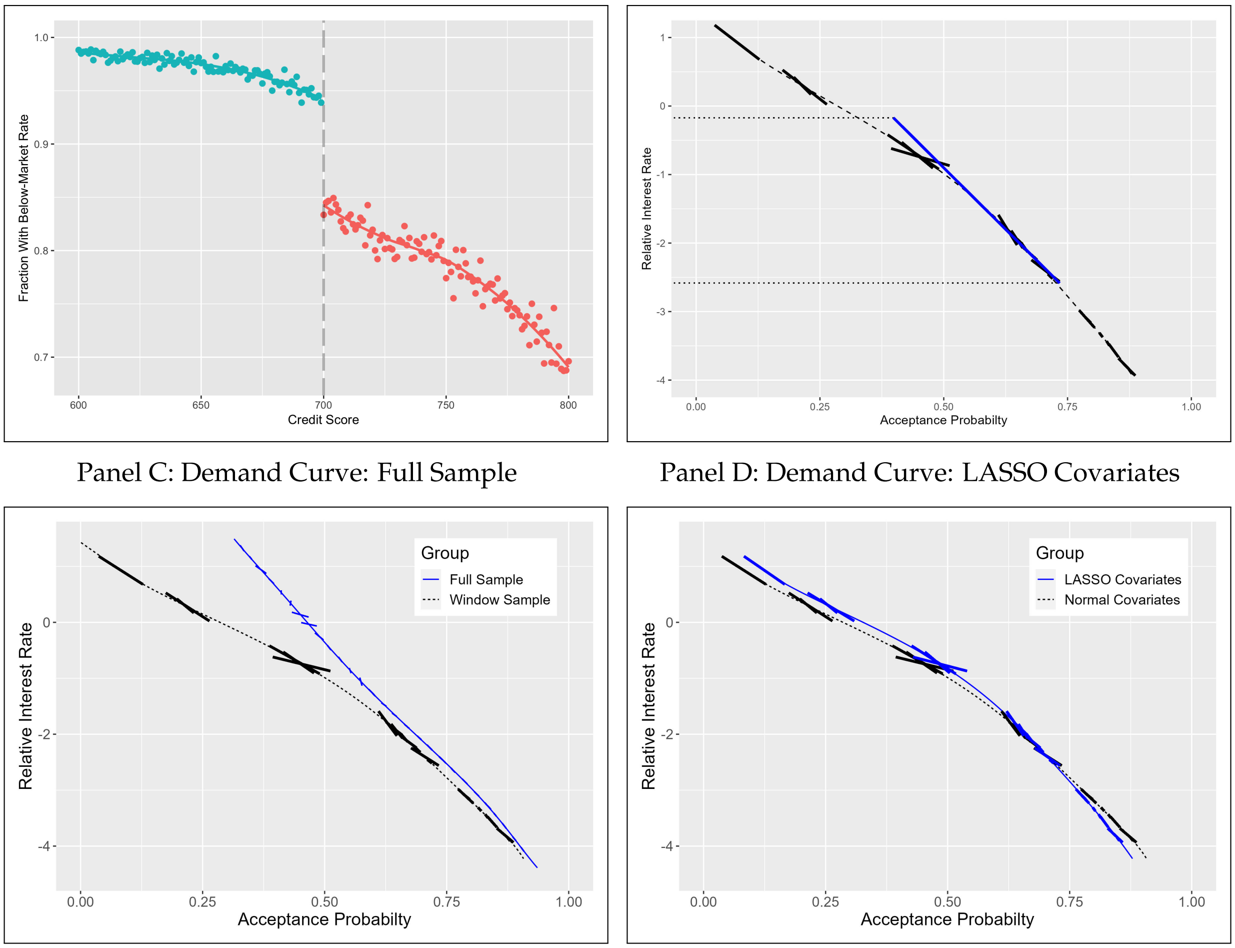

Note: Panel A shows the fraction of approved applicants who are offered the program's below-market interest rate by credit score. Each point represents a credit score. One criteria for the interest rate determination is that the applicant has a credit score below 700. Panel B presents the local demand curve from the fuzzy discontinuity (bandwidth of 4; solid, blue line) overlaid on our time-based, local demand curves. Panel C presents a naïve demand curve estimation using the full sample and defining "windows" based on 54 pairs of adjacent rate bins. Panel D presents a demand curve estimated using the window sample, but with additional covariates selected by Lasso estimation.

The second stage estimates how the household's (extensive margin) demand for the loan changes when offered the market vs. below-market rate. Formally, for household i, we estimate:

$$D_i = \alpha_1 + \tau_1 1\{CreditScore_i \geq 700\} + f(CreditScore_i; \theta_1, 1\{CreditScore_i \geq 700\}) + \varepsilon_i$$
$$P(Accept_i) = \alpha_2 + \tau_2 \widehat{D_i} + f(CreditScore_i; \theta_2, 1\{CreditScore_i \geq 700\}) + u_i \tag{3}$$

Where D_i is a binary indicator for the household receiving the market interest rate (our treatment).

$\widehat{D_i}$ is the predicted value from the first stage discontinuity; $f()$ is a local linear approximation parameterized by θ weighted by a triangular kernel; and τ_2 is our estimated local average treatment effect.

Table 8: Fuzzy Regression Discontinuity Models for Loan Acceptance

Dependent Variable: Loan Acceptance Running Variable: Credit Score			
	(1)	(2)	(3)
I(Market Rate)	-0.312***	-0.335***	-0.298***
	(0.045)	(0.044)	(0.054)
Bandwidth	3	4	5
Number Observations LATE	9,232	12,334	15,268

Note: *p<0.1; **p<0.05; ***p<0.01

This table shows the results of a fuzzy regression discontinuity for loan acceptance based on being offered the "market" interest rate using a fuzzy cutoff of 700 in the credit score, following Equation 3. The columns are differentiated by bandwidth. The kernel for the local linear estimation is triangular.

Table 8 shows the results of Equation (3). The columns represent different bandwidths, all weighted with a triangular kernel. For example, the results shown in Column (1) use a bandwidth of 3 and so compare approved applicants with credit scores of 697 to 699 to those with credit scores of 700 to 702 weighting the scores closer to 699.5 more. The results show applicants' sensitivity to the interest rate: on average, increasing the offered interest rate by 2.5 pp (the market rate is typically double the below-market rate) decreases the likelihood that approved applicants accept the loan by about 30 pp.

Panel B of Figure 5 plots the estimated slope from the middle column, which has the smallest standard error, of Table 8 onto the global demand curve. Despite using a completely different form of identification, the fuzzy RD estimate closely mirrors the previously estimated demand curve when spanning from the average relative market rate to the average relative below-market rate.

For comparison with the main results, we estimate the demand curve using two additional strategies. The first leverages the full sample of approved applicants, instead of including only applicants who experienced a disaster within two weeks of a quarterly rate update. In this case, we round the relative interest rate to the nearest 0.1% (i.e., we create 10 bp bins) and examine how consumer take-up varies from one interest rate to the next-closest observed rate. Otherwise, we use the same estimating framework described in the main results (Section 3.1, e.g., the same model controls as in Equation 1). While this approach increases statistical power by using a larger sample, we view it as a naïve estimate as it compares consumer take-up, not within a narrow window of a few weeks, but instead across months or even years. Since the FDL rate is fixed each quarter, the mortgage rate is the key source of variation in the relative rate for this naïve specification. As mortgage rates co-move with a variety of factors (e.g., economic cycles, real estate prices, savings rates), potential omitted variables abound that may affect credit demand. Panel C of Figure 5 shows the estimated demand curves for the full sample and the windows sample. The curves are highly similar for relative rates below -1, the range of relative rates most frequently offered in the program. The curves separate somewhat for higher interest rates: the full-sample curve is effectively linear instead of flattening like the windows-based curve. It is possible that this divergence is due to sample composition (e.g., some of the largest disasters are omitted from the windows sample due to their timing); however, given the identification concerns for the full-sample estimates, the divergence likely reflects estimation bias that leads the naïve specification to underestimate the price elasticity at higher interest rates.

The second approach uses the time-based sub-sample from our main analyses in a Lasso estimation, which includes more flexible model controls (see Online Appendix D). Lasso estimation involves a trade-off: added flexibility can improve model predictions but can be prone to over-fitting. Panel D of Figure 5 plots the demand curves using our primary model ("normal covariates") and the Lasso estimation. The results are nearly identical – the semi-elasticity from the Lasso estimation is -0.249 (vs. -0.259 for the primary model) – indicating that a more flexible treatment of the modeled controls does not meaningfully affect estimates of the demand curve.

4.2 Extension: Estimating Welfare of Recovery Loans

We use the demand curve estimated in Section 3 to measure the first-order welfare for recovery loans. Welfare analyses offer additional context for consumers' willingness to pay by comparing it to the cost of providing recovery loans. While some consumers have a large willingness to pay, many do not so the costs may outweigh the consumer surplus. We calculate costs from annual public records on the program's interest rate subsidies and administrative costs and impose a 30% cost of raising funds through taxation. Additional details and the full results are in Appendix A.

We find the loans generate first-order, positive social surplus – the consumer surplus derived from the estimated demand curve exceeds the programs' reported cost. We estimate that from

Figure 6: Social Surplus

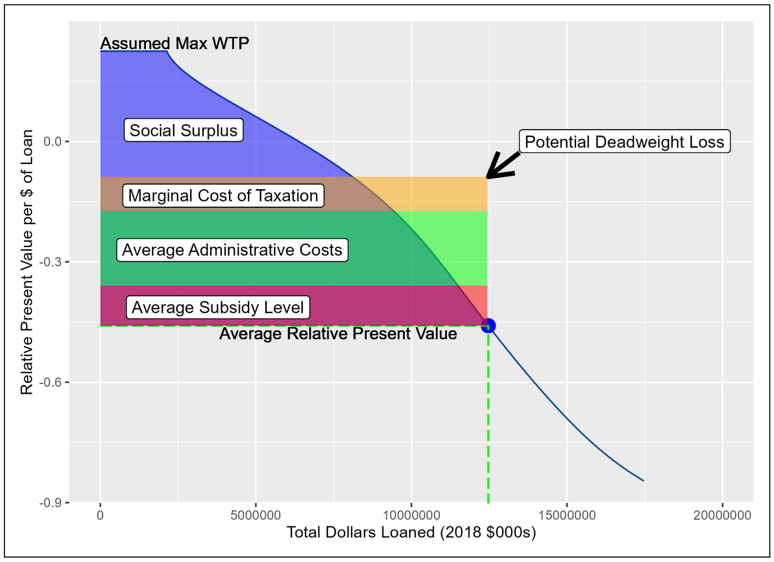

Note: This figure shows the estimated first-order social surplus, subsidy amount, and direct deadweight loss for the program from Jan. 2005 - May 2018. Aside from the marginal cost of taxation, this figure does not account for any externalities, positive or negative, from the program.

2005 to 2018, these recovery loans generated an average social surplus per borrower of $2,900, or $0.07 per dollar loaned. Figure 6 illustrates the magnitudes of consumer and producer surplus and the components of program costs.[30] While the program creates first-order welfare gains, it operates at a financial loss: the program would have needed to charge an interest rate of 4.5% to cover the costs of existing borrowers, which is 1.9 pp above its average interest rate but still 0.4 pp below the concurrent 30-year fixed mortgage rate. Our demand estimates indicate that only 62% of existing borrowers would be willing to pay the break-even rate. This welfare estimate only considers the consumer surplus of borrowers and loan costs; potential second-order effects such as positive, local spillovers from rebuilding (Fu and Gregory, 2019) may add to total welfare and contribute to the program's decision to keep interest rates low.

We also consider how welfare is allocated across sub-populations (e.g., based on creditworthiness). We compare welfare for below-median versus above-median DTI because demand differs so notably between these groups (Section 3.3). This back-of-the-envelope welfare calculation relies on additional assumptions regarding how costs are allocated between groups. For consumer surplus, we use the dis-aggregated demand curves for the low DTI and high DTI households.

[30] The flat cost curve depicted in the figure is equivalent to assuming no selection effects. Adverse or advantageous selection would not affect the magnitude of social surplus for existing borrowers, but would affect estimates of the deadweight loss and counterfactual interest rate analyses.

For producer surplus, we assign administrative costs based on the share of dollars loaned to each group and assign subsidy costs based on the share of loan dollars charged-off for each group. We find that recovery loans increase welfare for both groups but that around 85% of the social surplus accrues to the high DTI (i.e., low credit quality) group. We estimate a welfare gain of $0.17 per dollar loaned for high DTI households and $0.03 for low DTI households.

5 Conclusion

We provide new evidence on credit demand. Using rich administrative data on a large lending program, we estimate a household demand curve for credit after severe climate events. We find that a 1 pp increase in the offered interest rate lowers loan acceptance by 26% on average. Demand is lowest among high-credit-quality households and those for whom the new loan would represent a larger share of their discretionary income.

Our results offer insights on households' willingness to absorb large weather shocks through borrowing. Other adverse events similarly can generate sizable, unfunded shocks to households' balance sheets: unanticipated medical procedures, extended unemployment spells, short-term disability, household liability judgments, etc. Evaluating whether households' credit demand varies in response to these other events is beyond the scope of this paper, though our detailed administrative data provide variation on many dimensions, including household characteristics, time, and geography in our setting. A key implication of our findings pertains to managing increasing climate risks. Without interest rate subsidies, take-up of disaster recovery loans is likely to be low, potentially reducing or delaying recovery (e.g., Fu and Gregory, 2019) and increasing pressure on other public programs (Deryugina, 2017).

References

Agarwal, S., S. Chomsisengphet, and C. Liu (2010). The importance of adverse selection in the credit card market: Evidence from randomized trials of credit card solicitations. Journal of Money, Credit and Banking 42(4), 743–754.

Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2018). Do banks pass through credit expansions to consumers who want to borrow? The Quarterly Journal of Economics 133(1), 129–190.

Aladangady, A. (2017). Housing wealth and consumption: Evidence from geographically-linked microdata. American Economic Review 107(11), 3415–46.

Alan, S. and G. Loranth (2013). Subprime consumer credit demand: Evidence from a lender's pricing experiment. The Review of Financial Studies 26(9), 2353–2374.

Andersen, H. Y. and S. Leth-Petersen (2020, 03). Housing Wealth or Collateral: How Home Value Shocks

Drive Home Equity Extraction and Spending. Journal of the European Economic Association 19(1), 403–440.

Andersen, S., J. Y. Campbell, K. M. Nielsen, and T. Ramadorai (2020). Sources of inaction in household finance: Evidence from the danish mortgage market. American Economic Review 110(10), 3184–3230.

Argyle, B., T. Nadauld, and C. Palmer (2020). Monthly payment targeting and the demand for maturity. Review of Financial Studies. https://doi.org/10.1093/rfs/hhaa004.

Attanasio, O. P., P. Koujianou Goldberg, and E. Kyriazidou (2008). Credit constraints in the market for consumer durables: Evidence from micro data on car loans. International Economic Review 49(2), 401–436.

Bai, J. and S. Ng (2008). Forecasting economic time series using targeted predictors. Journal of Econometrics 146(2), 304–317.

Bakkensen, L. A. and L. Barrage (2021). Going underwater? flood risk belief heterogeneity and coastal home price dynamics. The Review of Financial Studies.

Baugh, B., I. Ben-David, H. Park, and J. A. Parker (2021). Asymmetric consumption smoothing. American Economic Review 111(1), 192–230.

Begley, T. A., U. Gurun, A. K. Purnanandam, and D. Weagley (2020). Disaster lending: 'fair' prices, but 'unfair' access. https://ssrn.com/abstract=3145298.

Belloni, A., D. Chen, V. Chernozhukov, and C. Hansen (2012). Sparse models and methods for optimal instruments with an application to eminent domain. Econometrica 80(6), 2369–2429.

Belloni, A., V. Chernozhukov, and C. Hansen (2014a). High-dimensional methods and inference on structural and treatment effects. The Journal of Economic Perspectives 28(2), 29–50.

Belloni, A., V. Chernozhukov, and C. Hansen (2014b). Inference on treatment effects after selection among high-dimensional controls. The Review of Economic Studies 81(2), 608–650.

Belloni, A., V. Chernozhukov, C. Hansen, and D. Kozbur (2016). Inference in high-dimensional panel models with an application to gun control. Journal of Business & Economic Statistics 34(4), 590–605.

Bernstein, A., M. T. Gustafson, and R. Lewis (2019). Disaster on the horizon: The price effect of sea level rise. Journal of financial economics 134(2), 253–272.

Bhutta, N. and B. J. Keys (2016). Interest rates and equity extraction during the housing boom. American Economic Review 106(7), 1742–74.

Bhutta, N., P. M. Skiba, and J. Tobacman (2015). Payday loan choices and consequences. Journal of Money, Credit and Banking 47(2-3), 223–260.

Billings, S. B., E. Gallagher, and L. Ricketts (2022). Let the rich be flooded: The distribution of financial aid and distress after hurricane harvey. Journal of Financial Economics 146(2), 797–819.

Blickle, K. S., J. A. Santos, et al. (2022). Unintended consequences of" mandatory" flood insurance. Federal Reserve Bank of New York.

Bureau, U. C. (2022). American community survey 1-year estimates. Accessed via R.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. The Review of Economics and Statistics 90(3), 414–427.

Caner, M. (2009). Lasso-type gmm estimator. Econometric Theory 25(1), 270–290.

Carson, J. M., C. M. Ellis, R. E. Hoyt, and K. Ostaszewski (2020). Sunk costs and screening: Two-part tariffs in life insurance. Journal of Risk and Insurance 87(3), 689–718.

Census Bureau (2018). Understanding and using American Community Survey Data: What all data users need to know.

Chernozhukov, V., C. Hansen, and M. Spindler (2015). Post-selection and post-regularization inference in linear models with many controls and instruments. The American Economic Review 105(5), 486–490.

Cloyne, J., K. Huber, E. Ilzetzki, and H. Kleven (2019). The effect of house prices on household borrowing:

a new approach. American Economic Review 109(6), 2104–36.

CMHC (2022). Residential Rehabilitation Assistance Program. Canadian Mortgage and Housing Corporation, *https://www.cmhc-schl.gc.ca/en/professionals/project-funding-and-mortgage-financing/funding-programs/all-funding-programs/residential-rehabilitation-assistance-program*.

Cohen, P., R. Hahn, J. Hall, S. Levitt, and R. Metcalfe (2016, September). Using big data to estimate consumer surplus: The case of uber. Working Paper 22627, National Bureau of Economic Research.

Collier, B. L., C. M. Ellis, and B. J. Keys (2021). The cost of consumer collateral: Evidence from bunching. NBER Working Paper 29527, https://www.nber.org/papers/w29527.

DeFusco, A. A. (2018). Homeowner borrowing and housing collateral: New evidence from expiring price controls. The Journal of Finance 73(2), 523–573.

DeFusco, A. A. and A. Paciorek (2017). The interest rate elasticity of mortgage demand: Evidence from bunching at the conforming loan limit. American Economic Journal: Economic Policy 9(1), 210–40.

del Valle, A., T. C. Scharlemann, and S. H. Shore (2022). Household financial decision-making after natural disasters: Evidence from Hurricane Harvey.

Deryugina, T. (2017). The fiscal cost of hurricanes: Disaster aid versus social insurance. American Economic Journal: Economic Policy 9(3), 168–98.

Deryugina, T., L. Kawano, and S. Levitt (2018). The economic impact of Hurricane Katrina on its victims: Evidence from individual tax returns. American Economic Journal: Applied Economics 10(2), 202–33.

Di Maggio, M., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao (2017, November). Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. American Economic Review 107(11), 3550–88.

Dobbie, W. and P. M. Skiba (2013). Information asymmetries in consumer credit markets: Evidence from payday lending. American Economic Journal: Applied Economics 5(4), 256–82.

Dobridge, C. L. (2018). High-cost credit and consumption smoothing. Journal of Money, Credit and Banking 50(2-3), 407–433.

Einav, L., A. Finkelstein, and M. R. Cullen (2010). Estimating welfare in insurance markets using variation in prices. The Quarterly Journal of Economics 125(3), 877–921.

Experian (2020). 2019 consumer credit review. Prepared by Matt Tatham, *www.experian.com/blogs/ask-experian/consumer-credit-review/*.

Fannie Mae (2019a). Selling guide: Fannie mae single family. https://www.fanniemae.com/content/guide/selling/b3/6/02.html.

Fannie Mae (2019b). Single-family fixed rate mortgage dataset. https://loanperformancedata.fanniemae.com/lppub/index.html.

FDL Program (2022). Disaster loan assistance. Federal Disaster Loan Program, *https://disasterloanassistance.sba.gov/ela/s/*.

Federal Reserve (2022a). 30-year fixed rate mortgage average in the United States. Federal Reserve Bank of St. Louis, https://fred.stlouisfed.org/series/MORTGAGE30US.

Federal Reserve (2022b). Consumer price index (cpi) data. Retrieved from Federal Reserve Economic Data (FRED) database via R.

Federal Reserve (2022c). Market yield on u.s. treasury securities at 30-year constant maturity, quoted on an investment basis. Federal Reserve Bank of St. Louis, https://fred.stlouisfed.org/series/3DGS30.

FEMA (2019). Individual disaster assistance. Federal Emergency Management Agency.

Finkelstein, A., N. Hendren, and M. Shepard (2019). Subsidizing health insurance for low-income adults: Evidence from massachusetts. American Economic Review 109(4), 1530–67.

Finkelstein, A. and R. McKnight (2008). What did medicare do? the initial impact of medicare on mortality and out of pocket medical spending. Journal of Public Economics 92(7), 1644–1668.

First Street Foundation (2021). The cost of climate: America's growing flood risk. https://assets.firststreet.org/uploads/2021/02/The_Cost_of_Climate_FSF20210219-1.pdf.

Foote, C. L., L. Loewenstein, and P. S. Willen (2020). Cross-Sectional Patterns of Mortgage Debt during the Housing Boom: Evidence and Implications. The Review of Economic Studies 88(1), 229–259.

Freddie Mac (2019). Single family loan-level dataset. http://www.freddiemac.com/research/datasets/sf_loanlevel_dataset.page.

Fu, C. and J. Gregory (2019). Estimation of an equilibrium model with externalities: Post-disaster neighborhood rebuilding. Econometrica 87(2), 387–421.

Gallagher, J. and D. Hartley (2017). Household finance after a natural disaster: The case of hurricane katrina. American Economic Journal: Economic Policy 9(3), 199–228.

Ganong, P. and P. Noel (2019). Consumer spending during unemployment: Positive and normative implications. American economic review 109(7), 2383–2424.

Ganong, P. and P. Noel (2020). Liquidity versus wealth in household debt obligations: Evidence from housing policy in the great recession. American Economic Review 110(10), 3100–3138.

GAO (2016). Credit reform: Current method to estimate credit subsidy costs is more appropriate for budget estimates than a fair value approach. U.S. Government Accountability Office, GAO-16-41.

Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the life cycle. Econometrica 70(1), 47–89.

Government of the Netherlands (2020). Support and recovery package substantially expanded. *https://www.government.nl/latest/news/2021/01/21/support-and-recovery-package-substantially-expanded*.

gov.uk (2020). Recovery loan scheme. *https://www.gov.uk/guidance/recovery-loan-scheme*.

Gross, D. B. and N. S. Souleles (2002). Do liquidity constraints and interest rates matter for consumer behavior? evidence from credit card data. The Quarterly Journal of Economics 117(1), 149–185.

Guerrieri, V., G. Lorenzoni, and M. Prato (2020). Slow household deleveraging. Journal of the European Economic Association 18(6), 2755–2775.

Hahn, J., P. Todd, and W. Van der Klaauw (2001). Identification and estimation of treatment effects with a regression-discontinuity design. Econometrica 69(1), 201–209.

Hendren, N. and B. Sprung-Keyser (2020, 03). A Unified Welfare Analysis of Government Policies*. The Quarterly Journal of Economics. qjaa006.

Hundtofte, S., A. Olafsson, and M. Pagel (2019). Credit smoothing. Technical report, National Bureau of Economic Research.

Hurst, E. and F. Stafford (2004). Home is where the equity is: Mortgage refinancing and household consumption. Journal of Money, Credit and Banking, 985–1014.

Japanese Finance Corporation (2022). Special purpose loans. *https://www.jfc.go.jp/n/english/sme/loans.html*.

Karlan, D. and J. Zinman (2019). Long-run price elasticities of demand for credit: evidence from a countrywide field experiment in mexico. The Review of Economic Studies 86(4), 1704–1746.

Karlan, D. S. and J. Zinman (2008). Credit elasticities in less-developed economies: Implications for microfinance. American Economic Review 98(3), 1040–68.

Keys, B. J. and P. Mulder (2020). Neglected no more: Housing markets, mortgage lending, and sea level rise. National Bureau of Economic Research.

Keys, B. J., J. Tobacman, and J. Wang (2018). Rainy day credit? unsecured credit and local employment shocks.

Kfw (2020). Kfw coronavirus aid. *https://www.kfw.de/About-KfW/Reporting-Portal/KfW-coronavirus-aid/*.

Laibson, D., A. Repetto, and J. Tobacman (2007). Estimating discount functions with consumption choices over the lifecycle.

Larsen, T. (2017). Hurricane harvey: Identifying the insurance gap. CoreLogic, corelogic.com/intelligence/hurricane-harvey-identifying-the-insurance-gap.

Lee, D. S. and T. Lemieux (2010). Regression discontinuity designs in economics. Journal of Economic Literature 48(2), 281–355.

McCrary, J. (2008). Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test. Journal of Econometrics 142(2), 698–714.

MEF (2020). Liquidity support for households, businesses and local authorities. Ministero dell'Economia e delle Finance (MEF), *https://www.mef.gov.it/en/covid-19/Liquidity-support-for-households-businesses-and-local-authorities/*.

Mian, A., K. Rao, and A. Sufi (2013). Household balance sheets, consumption, and the economic slump. The Quarterly Journal of Economics 128(4), 1687–1726.

Mian, A. and A. Sufi (2011). House prices, home equity-based borrowing, and the us household leverage crisis. American Economic Review 101(5), 2132–56.

Morse, A. (2011). Payday lenders: Heroes or villains? Journal of Financial Economics 102(1), 28–44.

Mulder, P. (2021). Mismeasuring risk: The welfare effects of flood risk information. https://faculty.wharton.upenn.edu/wp-content/uploads/2017/07/MismeasuringRisk_Mulder2021.pdf.

NOAA (2022). U.S. billion-dollar weather and climate disasters. U.S. National Oceanic and Atmospheric Administration, DOI: 10.25921/stkw-7w73.

NSW (2022). Disaster recovery loans. New South Wales Government, *https://www.raa.nsw.gov.au/disaster-assistance/disaster-recovery-loans*.

Office of Disaster Assistance (2018). Standard operating procedure, Disaster Assistance Program. Small Business Administration, 31 May 2018.

Olafsson, A. and M. Pagel (2018). The liquid hand-to-mouth: Evidence from personal finance management software. The Review of Financial Studies 31(11), 4398–4446.

Ponce, A., E. Seira, and G. Zamarripa (2017). Borrowing on the wrong credit card? evidence from mexico. American Economic Review 107(4), 1335–61.

Poterba, J. M. (1996). Government intervention in the markets for education and health care: how and why? In Individual and social responsibility: Child care, education, medical care, and long-term care in America, pp. 277–308. University of Chicago Press.

Public Safety Canada (2022). Disaster assistance programs. *https://www.publicsafety.gc.ca/cnt/mrgnc-mngmnt/rcvr-dsstrs/dsstr-ssstnc-prgrms/dsstr-ssstnc-prgrms-ll-en.aspx*.

Putzier, K. (2019). A natural disaster wipes out your home. Then the buyers come calling. *Wall Street Journal*, https://www.wsj.com/articles/the-new-storm-chasers-real-estate-disaster-investors-11564498767.

RealtyTrac (2015). 2015 U.S. natural disaster housing risk report.

Sastry, P. (2022). Who bears flood risk? evidence from mortgage markets in florida. .

SBA (2017). SBA Disaster Loan Program: Frequently asked questions. Small Business Administration, sba.gov/sites/default/files/articles/sba-disaster-loans-faq.pdf.

SBA (2020). Sba disaster guide. Small Business Administration, sba.gov/sites/default/files/disaster_assistance_transcript_1.pdf, Accessed 6 May 2020.

SBA (2022). Covid-19 economic injury disaster loan. Small Business Administration, *https://www.sba.gov/funding-programs/loans/covid-19-relief-options/eidl*.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society. Series B (Methodological) 58(1), 267–288.

Walsh, M. W. (2017). Homeowners (and taxpayers) face billions in losses from Harvey flooding. New York Times, 28 August 2017, https://www.nytimes.com/2017/08/28/business/dealbook/flood-insurance-harvey.html.

Ware, M. (2017). Storm watch: Making sure SBA's disaster loan program is prepared.

You, X. and C. Kousky (2022). Funding household disaster recovery: Evidence on the role of insurance.

Wharton Risk Center Working Paper.

Appendix A Welfare

A.1 Consumer Surplus

To estimate aggregate consumer surplus, we convert the acceptance rate into a quantity loaned and interest rates into prices.

$$\begin{aligned} \text{Social Surplus} &= CS + PS \\ &= \int_0^{q^*} \left[D^{-1}(q) - p^*\right] dq + PS \end{aligned} \tag{A1}$$

where consumer surplus CS depends on the relationship between quantity loaned q and loan prices p estimated in the inverse demand function $D^{-1}(\cdot)$. Values q^* and p^* respectively represent the total amount the program loaned and the average price on those loans. Regarding quantity, we assume that loan decisions are only made on the extensive margin, consumers accept or decline the loan, and that the intensive margin, the size of the loan, is determined by the size of the uninsured portion of the loss.[31] Using this assumption, we calculate the quantity loaned by multiplying the total number of approved applicants per year, the acceptance percentage, and the average loan amount.

To translate relative interest rates into relative prices, we define the price p as the present value, to the borrower, of repayments per dollar of a 30-year, amortized loan at the FDL program's offered rate relative to the amortized amount at the private mortgage rate. Specifically,

$$p = \sum_{t=1}^{360} \beta^t \left(A_{FDL} - A_{Mort}\right)$$

where A is the amortized monthly payment per dollar loaned and β is the assumed discount rate. For example, consider an applicant offered a recovery loan at a 3% interest rate while the mortgage rate was 5%. A $100,000 loan amortized over 30 years would require monthly payments of $422 for the FDL loan and $537 for the mortgage rate. An annual discount rate of 0.98 implies a present value of these repayments of $113,543 for the FDL loan and $144,485 for the mortgage rate. These present values imply prices (per dollar loaned) of $1.14 and $1.44. We thus define our relative price as the difference, which here equals -$0.30.

Unlike the quantity transformation above, the transformation of relative rates to relative prices is not linear and thus we re-run the OLS estimation in Equation (1) using relative prices (amortized loan amounts) instead of relative rates. The inverse demand curve $D^{-1}(\cdot)$ in Equation (A1) represents this estimation using relative prices. To match the present value calculations used by the government, which we use to calculate producer surplus in the next section, we define the discount rate for each applicant based on the average 30-year treasury rate at the time of application (Federal Reserve, 2022c). Over our data's timeframe, the average treasury rate is 3.7%, which implies an annual discount rate of 0.963 and a monthly $\beta = 0.9968$. We do not observe relative prices above a certain amount as we do not observe interest rates higher than 1.4 pp above 30 year mortgage rates. Therefore, we assume that the maximum willingness to pay is the highest price

[31] As described in Section 3.2, we do not find any statistically significant evidence of intensive margin effects in our main analyses; however, the estimates are too noisy to claim a null effect.

per dollar loaned that we observe. This assumption is conservative: it may be that a portion of borrowers would pay higher prices and, if so, consumer surplus would be larger.

The shaded blue in Figure 6 shows the result of these transformations on the demand curve. The point on the curve identifies the total amount loaned q^* ($12.47 billion in 2018 dollars) and the average relative price p^* (-$0.46 per dollar loaned in 2018 dollars). The shaded area represents the consumer surplus from the program and totals $5.44 billion.

$$\text{Social Surplus} = \$5.44B + PS$$

A.2 Producer Surplus

Calculating social surplus requires an estimate on how much the program costs. Due to low offered interest rates, delinquencies, and administrative costs, the program, on average, loses money and thus producer surplus is negative. To estimate the cost of subsidized interest rates and delinquencies, the program calculates a subsidy rate for its loans following the Federal Credit Reform Act of 1990 (FCRA). The FCRA subsidy calculation is approximately the present value of the expected loss on issued loans.[32] The program reports administrative costs in addition to the subsidy amount.[33] Additionally, we assume that the program's subsidy is paid by taxpayers and the additional marginal cost of raising these public funds is 30% (following Poterba, 1996; Finkelstein and McKnight, 2008), which accounts for the deadweight loss due to taxation.

Table A1 shows the costs for the program, separated by fiscal year and source, over our time frame. The final column, cost per dollar loaned, shows the scale economies of disaster lending. The three years with the largest lending (2005, 2006, and 2017) also feature some of the lowest average costs. These scale economies appear at least partially due to fixed costs: for years with less costly disasters (e.g., 2014) the costs per dollar loaned exceed 1. The final row of the table sums the costs and implies the total producer surplus from Jan. 2005 to May 2018 is -$4.63 billion.

$$\begin{aligned}\text{Social Surplus} &= \$5.44B + PS \\ &= \$5.44B - \$4.63B \\ &= \$0.82B\end{aligned}$$

These calculations result in a first-order social surplus over the same time frame of $820 million, which equates to $2,700 per borrower or $0.07 per dollar loaned. These numbers imply a marginal value of public funds (MVPF) dedicated to the program of 1.53, which is similar to other government programs for adults (Hendren and Sprung-Keyser, 2020).[34]

[32] The subsidy rate is, based on current treasury rates with similar durations, the present value PV of expected cash outflows minus the PV of expected cash inflows divided by the former ($SubsidyRate = (PV_{Out} - PV_{In})/PV_{Out}$). The FCRA measure incorporates delinquencies and federal borrowing rates but not administrative costs (GAO, 2016).

[33] Cost breakdowns are found in the Agency Financial Reports and are available at https://www.sba.gov/document/report--agency-financial-report.

[34] The MVPF compares consumer surplus to the government's cost of offering a program. The standard cases considered by Hendren and Sprung-Keyser (2020) examine how fiscal externalities affect MVPF (e.g., how expanding scholarships affects future tax revenues). In these standard cases, consumer surplus is already given (e.g., the dollar value of offered scholarships), and the focus is on causal estimates of second-order government costs. In contrast, the challenge in our setting is estimating consumer surplus for loans, and we take government costs from federal reports as given. Following this literature, we do not include the marginal cost of raising public funds in the MVPF calculation.

Table A1: Costs of FDL Loan Program

(1) Fiscal Year	(2) Total Loaned	(3) Admin Costs	(4) Subsidy Costs	(5) Marginal Cost of Taxation	(6) Total Costs	(7) Cost per $ Loaned
2005	1102.11	75.75	55.56	39.39	170.70	0.15
2006	5651.89	773.94	483.31	377.17	1634.42	0.29
2007	237.56	125.36	120.86	73.87	320.08	1.35
2008	347.53	96.61	145.85	72.74	315.19	0.91
2009	316.21	121.15	134.23	76.61	331.99	1.05
2010	243.01	79.21	81.02	48.07	208.31	0.86
2011	362.07	99.88	102.59	60.74	263.21	0.73
2012	211.38	87.68	30.92	35.58	154.18	0.73
2013	841.24	130.16	42.13	51.69	223.97	0.27
2014	96.78	99.46	28.56	38.40	166.42	1.72
2015	186.98	105.88	-51.43	16.33	70.78	0.38
2016	749.99	149.24	-76.81	21.73	94.15	0.13
2017	1135.53	162.31	4.16	49.94	216.41	0.19
2018	943.32	203.94	149.58	106.06	459.58	0.49
Total	12425.62	2310.56	1250.52	1068.32	4629.40	0.37

Note: This table shows the estimated costs, in millions of 2018 dollars, for the program separated by fiscal year and source. All numbers are estimated from total program costs to represent the costs of household loans only by weighting based on the relative dollars loaned to households vs. businesses in each fiscal year (1 Oct. to 30 Sept.). The costs for 2005 and 2018 are adjusted to match our data timeline based on average monthly loan rates. The split between admin and subsidy costs for 2005 and 2006 are estimated based on the reported total program cost for those years and the relative split from other years. Listed subsidy costs include re-estimates of subsidies from prior years. The marginal cost of taxation is estimated at 30% of the combined admin and subsidy costs.

Figure 6 plots a graphical depiction of the cost of the program onto our demand curve. We equally weight the average cost for each dollar loaned, which is equivalent to assuming no selection effects (Einav et al., 2010). In contrast to Section 3, which focuses on the share of approved applicants who accept the loan, the discussion of consumer surplus is on borrowers. For existing borrowers, the program would need to charge an average interest rate of 4.5% to fully cover its costs, which is 1.9 pp above its average interest rate but still 0.44 pp below the concurrent 30-year fixed mortgage rate. About 51% of borrowers would be willing to pay the prevailing mortgage rate, and another 12% of current borrowers would be willing to pay above the program's average cost but below the prevailing mortgage rate. The remaining 38% of borrowers would not accept the loans if they were offered at the program's average cost.

The figure shows the potential deadweight loss generated by the program subsidy, again assuming no selection effects. Of the -$4.6B in producer surplus, -$0.72B is deadweight loss. This amounts to -$0.06 per dollar loaned. If adverse (advantageous) selection affects the program such that households with higher (lower) willingness to pay have higher delinquency, the deadweight loss would be smaller (larger).[35] Our results from Section 3.3 suggest that less creditworthy applicants have a higher willingness to pay and so imply that we are likely over-estimating the first-order deadweight loss from the program. Selection effects would also influence the welfare associated with counterfactual interest rates. For example under adverse selection, raising rates would increase the program's average cost by changing the composition of borrowers and so if the program intended to charge a break-even interest rate, that rate would need to be higher than 4.5% (the break-even rate for existing borrowers) since raising rates would cause high quality

[35] Note that about two thirds of program costs are associated with administration and the marginal cost of taxation to fund administration, which would tend to be less affected by any selection.

borrowers to disproportionately select out of the program.

Our first-order welfare analyses do not consider several potential second-order effects. As discussed in Section 2, second-order effects may include *positive* externalities such as benefits to neighbors generated from repairing damaged homes (Fu and Gregory, 2019) and possibly a reduction in households' use of other public safety-net programs (Deryugina, 2017). It is also possible that the *existence* of recovery loans creates first-order, *ex ante* benefits not considered here. For example, the consumption and savings allocations of households in disaster-prone areas might be influenced by the ability to take a low-interest loan if disaster strikes.

The most prominent potential *negative* externality is crowd-out: the potential that in the absence of the FDL program, the private market would have lent to these households. We believe this is not the case for most consumers given our estimates of credit demand. These households often have limited collateral, low credit scores, and/or high DTI and so many may be unable to find private lenders willing to offer rates that they would pay. Thus, to the extent that the program crowds out private lenders, it appears to do so by offering low interest loans to the (small) set of households who might otherwise have turned to personal finance loans, credit cards, or other high cost credit products.

Our first-order welfare results also raise questions regarding why the private market has not addressed household demand. As Table A1 illustrates, the provision of recovery loans may benefit from substantial economies of scale. Additional documentation seems to support this market frictions interpretation. At the behest of Congress in 2008, the SBA developed a program for guaranteeing private loans for post-disaster rebuilding. However, this program remains unimplemented due to a lack of interest from the private sector. The Inspector General of the SBA explained that private lenders' objections included a lack of expertise and the high administrative costs required to offer disaster recovery loans (Ware, 2017).

We extend this approach to provide a back-of-the-envelope comparison of how welfare is allocated between low and high DTI applicants in Online Appendix F.