

UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

Essays in Household Finance and Public Economics

Permalink

<https://escholarship.org/uc/item/0426q1w1>

ISBN

9798247938712

Author

Fitzpatrick, Michael

Publication Date

2026-06-02

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Essays in Household Finance and Public Economics

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Economics

by

Michael Fitzpatrick

Dissertation Committee:
Professor Matthew Freedman, Chair
Assistant Professor Marion Aouad
Associate Professor Jack Liebersohn

DEDICATION

This work is dedicated to my parents, without whom I would be nothing.

TABLE OF CONTENTS

	Page
LIST OF FIGURES	v
LIST OF TABLES	vii
ACKNOWLEDGMENTS	viii
VITA	ix
ABSTRACT OF THE DISSERTATION	x
1 Debt Aversion and Household Borrowing: Evidence from Mortgage Repayment	1
1.1 Introduction	1
1.2 Mortgage Valuation and the Cost of Prepayment	6
1.3 Data and Measurement	9
1.4 Empirical Patterns	11
1.4.1 Prevalence and Cost of Negative-Spread Prepayment	11
1.4.2 Assumptions	13
1.4.3 Insensitivity to Interest Rates	14
1.4.4 Sensitivity to Cash Flows	15
1.4.5 Prepayment Near Mortgage Completion	16
1.5 Mechanisms and Interpretation	17
1.5.1 Candidate Explanations	17
1.5.2 Evidence for Debt Aversion	21
1.5.3 Illustrative Bounds on the Prevalence and Intensity of Debt Aversion . . .	27
1.6 Conclusion	28
1.7 Figures and Tables	31
2 How do Economic Transfers affect Health? Evidence from the Alaska Permanent Fund Dividend	49
2.1 Introduction	49
2.2 Alaska Permanent Fund Background	54
2.3 Data	56
2.3.1 Summary Statistics	57
2.4 Methodology	59

2.5	Results	60
2.5.1	Drug and Alcohol Admissions	60
2.5.2	Heterogeneous effects on alcohol-related admissions	61
2.5.3	Placebo admissions and robustness checks	64
2.5.4	Liquidity Constrained Health Behavior	65
2.5.5	Responses around the Announcement Date	66
2.6	Discussion	67
2.7	Conclusion	69
3	Dope and Disorder: The effects of the reformulation of OxyContin on crime and homelessness	85
3.1	Introduction	85
3.2	Background	90
3.2.1	OxyContin and its reformulation	91
3.2.2	Conceptual framework	92
3.3	Data and methodology	93
3.3.1	Methodology	95
3.3.2	Estimating equation	96
3.4	Results	97
3.4.1	Crime Results	98
3.4.2	Homelessness Results	103
3.5	Conclusion	105
	Bibliography	117
	Appendix A Chapter 1 Appendix	128
	Appendix B Chapter 2 Appendix	140

LIST OF FIGURES

	Page
1.1 Mortgage prepayment and Treasury yields	31
1.2 Probability of prepayment by borrower characteristics	32
1.3 Interest rate spreads and net present value of prepayments	33
1.4 Distribution of excess payments	34
1.5 NPV of excess payments for negative interest rate spread mortgages	35
1.6 Negative rate spread statistics over time	36
1.7 Total prepayment amounts and real disposable income growth	37
1.8 Probability of prepayment around interest rate changes	38
1.9 Probability of prepaying by remaining principal balance	39
1.10 Placebo credit outcomes by remaining mortgage principal balance	40
1.11 Sophistication, attentiveness and inertia	41
 2.1 Total weekly hospitalizations relative to the dividend	 70
2.2 Total weekly drug-related hospitalizations relative to the dividend	71
2.3 Effects of APFD on drug-related emergency department visits	72
2.4 Effects of APFD admissions by dividend amount	73
2.5 Effects of APFD on placebo admissions	74
2.6 Heterogeneous effects of APFD on alcohol-related admissions: Male share of admissions	75
2.7 Heterogeneous effects of APFD on alcohol-related admissions: Age	76
2.8 Heterogeneous effects of APFD on alcohol-related admissions: Urban share of admissions	77
2.9 Heterogeneous effects of APFD on alcohol-related admissions: Income	78
2.10 Heterogeneous effects of APFD on alcohol-related admissions: Insurance type .	79
2.11 Effects of the APFD on Alcohol ED admissions: July 2020 and September 2022 only	80
2.12 Effects of APFD on elective surgery procedures	81
2.13 Effects of APFD announcement on drug-related emergency department visits . .	82
 3.1 Age adjusted death rates by exposure to OxyContin	 107
3.2 Effects of OxyContin reformulation on drug possession/sale arrests	108
3.3 Dynamic effects of OxyContin reformulation on crime rates	109
3.4 Age adjusted death rates by placebo exposure	110
3.5 Drug Possession/Sales Arrest – Alternative Measure	111

3.6	Placebo effects of OxyContin reformulation on crime rates	112
3.7	Effects of illicit opioid use on homelessness rates	113

LIST OF TABLES

	Page
1.1 Summary statistics	42
1.2 NPV losses under different assumptions	43
1.3 Prepayment heterogeneity	44
1.4 Predicting prepayment from credit behavior	45
1.5 Effects of interest rate spread on the probability of prepaying	46
1.6 ARMs vs FRMs	47
1.7 Effects of interest rate variation: Flexible rate vs “locked-in” loans	48
2.1 APFD Summary, 2016-2023	83
2.2 Summary Statistics	84
3.1 Acquisitive crimes	114
3.2 Non-acquisitive crimes	115
3.3 Effects of the reformulation on offenses known (Alternative Measure)	116

ACKNOWLEDGMENTS

I would like to thank the three members of my committee – Matthew Freedman, Marion Aouad, and Jack Liebersohn – for their guidance and supervision throughout the past five years. Your support has made an immeasurable impact on my work and life. I would also like to thank Jan Brueckner, David Neumark, and all other faculty in the Department of Economics at UC Irvine for their feedback and direction.

I am grateful for the administrative and financial support from the Department of Economics and School of Social Sciences at the University of California, Irvine.

I would like to thank my external coauthors Vikram Jambulapati and Nathan Blascak for their outstanding contributions to this dissertation. I want to additionally thank Nathan for taking a chance on me as an intern many years ago.

This dissertation would have never existed without my closest friends and colleagues. I'd like to thank Gamy De La Torre and Noah Kouchekinia for making the hours spent in the office each day more fun than the hours in between. Additional thanks to Kory Fenstemacher, Nick Fucci, Sydney Gordon, Tony Hu, Renee Lee, Russ Leonard, Kevin Peng, Michael Span, and countless others who were there to celebrate the highs and commiserate the lows.

To my family – Mom, Dad, Casey, Jake: thank you for your unwavering support and encouragement to pursue my dreams. I owe everything to you.

Lastly, I thank Katherine Rosencrance for her endless support, guidance, and love.

VITA

Michael Fitzpatrick

EDUCATION

Doctor of Philosophy in Economics	2026
University of California, Irvine	<i>Irvine, CA</i>
Masters of Arts in Economics	2023
University of California, Irvine	<i>Irvine, CA</i>
Bachelor of Business Administration in Economics	2018
Temple University	<i>Philadelphia, PA</i>

RESEARCH EXPERIENCE

Graduate Research Assistant	2022–2025
University of California, Irvine	<i>Irvine, California</i>

TEACHING EXPERIENCE

Teaching Assistant	2021–2026
University of California, Irvine	<i>Irvine, California</i>

ABSTRACT OF THE DISSERTATION

Essays in Household Finance and Public Economics

By

Michael Fitzpatrick

Doctor of Philosophy in Economics

University of California, Irvine, 2026

Professor Matthew Freedman, Chair

Chapter 1 studies partial mortgage prepayment, or curtailment. Many fixed-rate mortgage borrowers in the United States voluntarily pay down low-rate debt even when safe assets pay more. We study this behavior after the 2022 rate hikes, when safe rates rose above many existing mortgage rates and extra principal payments became financially costly. More than one in five borrowers whose mortgage rate was below the safe rate prepaid in a typical month, generating \$4.6 billion in present-value losses from 2022–2024. Prepayment is persistent and almost unresponsive to interest rate spreads, but it responds to cash flow shocks, including stimulus payments and increases in ARM payments. It is not concentrated among inattentive or unsophisticated households: prepayers have higher credit scores, are more likely to hold stocks and retirement accounts, and borrowers who refinanced when rates were low are just as likely to prepay. Prepayment nearly doubles towards the end of the mortgage, as the balance approaches zero. The findings point to debt aversion—a preference for reducing debt even when returns are negative—as a first-order force in household borrowing.

Chapter 2 studies the effect of a near-universally received, unconditional, cash transfers on health behaviors. We examine how receipt of the Alaska Permanent Fund Dividend (APFD) – a large, near universal, and annually recurring cash transfer to Alaska residents – affects the timing and quantity of health care consumption. We firstly determine if receipt of the

APFD heightens risky behaviors by examining the timing and quantity of drug- and alcohol-related hospital admissions. We then study if receipt of the APFD promotes the consumption of “health-improving” behaviors, specifically hip and knee replacements, cataract surgery and tonsillectomies, by relaxing pre-existing liquidity constraints. We do not find that the APFD affects the consumption of health-improving behaviors, as measured. However, there is evidence of “dangerous liquidity” for certain health behaviors. Alcohol-related emergency department admissions sharply increase by approximately 15.7 percent within one week of the APFD disbursement; we do not observe changes in amphetamine, opioid, or cocaine consumption. Further, we do not find that individuals increase their alcohol or drug consumption around the date the APFD is announced, only around the date of APFD disbursement, consistent with the importance of liquidity for certain risky behaviors.

Chapter 3 studies the effect of an exogenous opioid supply shock – the 2010 reformulation of OxyContin – on criminal activity and the incidence of homelessness. Prior literature has shown that this reformulation had unintended consequences, namely that it resulted in massive spikes in heroin and synthetic opioid use. Using geographic variation in the ex-ante “intensity” of OxyContin prior to the reformulation, I use a difference-in-differences framework to study how increases in illicit opioid use affects crime from 2006-2019. I find that highly exposed areas saw large increases in arrests for heroin possession, but no increases in arrests for the sale of heroin. Further, increases in illicit opioid use has largely no effect on both income-generating and non-income generating crimes. On the other hand, I show evidence that increases in heroin and fentanyl use are associated with an 8% increase in the incidence of homelessness. These results show that funds earmarked for drug-related crime prevention programs could be more wisely spent on alleviating homelessness.

Chapter 1

Debt Aversion and Household Borrowing: Evidence from Mortgage Repayment

1.1 Introduction

Evidence from lab experiments and surveys has shown that many households are averse to holding debt, even when borrowing is financially advantageous (Martínez-Marquina and Shi, 2024; Meissner and Altmann, 2022; Loewenstein and Prelec, 1992). Thaler (1990) argues that such preferences can explain systematic deviations from benchmark lifecycle and portfolio choice models: if households treat debt as psychologically costly, they will save too much, borrow too little, and hold portfolios that appear suboptimally safe. Recent work has found patterns consistent with debt aversion in student loan take-up (Eckel et al., 2007; Gopalan et al., 2021; Caetano et al., 2019) and investment decisions in low-income countries (Berkouwer and Dean, 2022). However, there is little systematic evidence on how prevalent debt aversion is among the broad population of households, or how much it costs. This paper provides such evidence, using one of the largest and most consequential components of household balance sheets in the United States: mortgage debt.

Measuring debt aversion in the wild is difficult because situations in which debt repayment is clearly unprofitable are rare. In most settings, interest rate movements change both the return to paying down debt and the return on other safe uses of cash in the same direction. The post-2022 mortgage market provides an unusually useful setting because it breaks that connection. Many households originated fixed-rate mortgages in a low-rate environment, but by 2022 safe outside yields had risen above the rate on those existing mortgages. In a frictionless benchmark, these households should stop making extra principal payments. This setting therefore allows us to study, at scale, whether households continue to repay debt even when the financial benefit to doing so has turned negative.

We show that more than one in five borrowers with a negative rate spread—meaning their mortgage rate is below the risk-free rate, after accounting for taxes—continue to make voluntary extra payments each month. The likelihood of prepaying a given fixed-rate mortgage did not decline even as risk-free yields rose, and the behavior is persistent within borrowers over time. We estimate that from January 2022 through November 2024, the present value of losses from negative-spread prepayment amounted to \$4.6 billion in the conforming fixed-rate single-family mortgage market. Even among high-income households, who are less likely to be liquidity constrained, 24% of mortgage borrowers prepay their mortgage.

We then ask which mechanisms best fit this behavior. Two patterns are especially important. First, prepayment responds strongly to changes in current cash flows, but not to the present value of prepayment: it rises after stimulus payments and predictable decreases in adjustable-rate mortgage payments. In contrast, prepayment is virtually unresponsive to interest rate changes within borrower and is unrelated to lagged interest rates or interest rates at the time of origination, making sluggish decision-making less compelling as the primary explanation. Second, several patterns are consistent with debt aversion: prepayment is higher near mortgage completion, more common among higher-credit-score borrowers, and correlated across debt products. Inattention, inertia, and lack of financial sophistication receive less support in the

data. Households who refinanced when rates were low—a proxy for financial attentiveness—are equally likely to prepay as others, and Survey of Consumer Finances data show that pre-payers are more likely to hold stocks and retirement accounts. Prepayment is similar for likely “auto-payers” as other households, suggesting that hassle costs do not explain prepayment.

Perhaps the most distinctive evidence comes from the end of the mortgage. The probability of prepayment nearly doubles as borrowers approach their final payment, reaching nearly 50%. This pattern is a sharp prediction of debt aversion with concave disutility: experimental evidence shows that the psychological benefit of eliminating debt entirely exceeds the benefit of a marginal reduction, so borrowers should be most willing to sacrifice returns when full repayment is within reach (Martínez-Marquina and Shi, 2024; Meissner and Altmann, 2022). The spike in prepayment near completion is difficult to reconcile with alternative explanations. It is not driven by retirement timing or lifecycle savings motives: payments on credit cards and other financial products do not change when mortgage prepayment rises, and controlling for borrower fixed effects and calendar time does not attenuate the result. Credit score incentives also cannot explain it, as credit scores do not discretely improve upon mortgage completion. Standard portfolio theory predicts no such change, since the return to prepaying does not change as the balance falls. The end-of-mortgage result does not imply that *all* negative-spread prepayment reflects debt aversion, but it provides direct evidence that debt aversion is operative for a substantial share of borrowers. Throughout, we use debt aversion to mean any nonpecuniary disutility from holding debt, so motives such as peace of mind or wanting to enter retirement debt-free are not alternatives to the mechanism we study but instances of it.

While prepayment is not related to the NPV of prepaying, it is related to short-term cash flows, similar to Ganong and Noel (2020). In aggregate data, prepayment spikes following the release of COVID stimulus checks and tracks the disposable income of high-income households. We also compare adjustable-rate (ARM) to fixed-rate (FRM) mortgages. In contrast to FRMs, interest rate increases reduce ARM borrowers’ monthly cash flows but have little effect on the

net present value of prepaying. Rate increases reduce ARM borrowers' propensity to prepay, suggesting that cash flows rather than present-value calculations drive changes in prepayment behavior. Therefore, we take the 22% monthly prepayment rate among negative-spread borrowers as a lower bound on the prevalence of a preference for debt aversion; some households might like to pay down debt, but their cash flows may not be high enough to do so. Using imputed income in the Fannie Mae data, we show that prepayment rises with income but flattens out high in the distribution: among borrowers high in the imputed-income distribution, about 24% prepay despite a negative spread. We interpret this as a conservative, liquidity-adjusted lower bound on the debt-averse share of mortgage borrowers. We also provide a lower bound on the intensity of debt aversion: the median negative-spread prepayer behaves as if the mortgage rate were about 90 basis points higher than its contractual rate, or roughly 30% higher relative to the contract rate in this group.

What do our results mean for economic theory? We highlight two implications of debt aversion, beyond the obvious modifications to models of portfolio choice. First, dispersion in debt aversion can generate sluggish aggregate responses to interest rate changes — not because households fail to observe rate movements (Sims, 2003; Reis, 2006), but because many are inframarginal with respect to prepayment. This mechanism can account for the muted pass-through of monetary policy to household balance sheets without requiring inattention. Second, if households are averse to borrowing per se, they may behave as if liquidity-constrained even when they have unused borrowing capacity (Thaler, 1990). Mortgages are the largest component of household debt, and roughly one in three mortgage borrowers prepays at least once per year, so these preference-driven distortions operate at a scale that matters for aggregate consumption and wealth accumulation.

Our findings are robust to a range of assumptions about interest rates and investment alternatives. We measure prepayment using publicly available Fannie Mae loan performance data and replicate key results using a credit registry. The central finding—that prepayment rates

do not respond to changes in interest rate spreads—does not depend on specific assumptions about tax rates or outside investment options, since it is identified from within-borrower variation over time. The *level* of losses does depend on assumptions about households’ marginal tax rates, which determine the after-tax mortgage rate and the after-tax return on alternative investments. Using NBER TAXSIM and data from the Survey of Consumer Finances, we estimate that the share of borrowers with negative rate spreads ranges from 39% to 69%, and aggregate losses from \$2.2 billion to over \$8.4 billion, depending on assumptions about the best available alternative.

Our findings contribute to several literatures. In the same spirit as Genesove and Mayer (2001), who provide the first large-scale evidence of loss aversion in the housing market, we provide the first such evidence for debt aversion. While debt aversion has been documented in lab experiments and student loan markets (Gopalan et al., 2021; Eckel et al., 2007; Prelec and Loewenstein, 1998), ours is the highest-stakes and most widespread setting in which it has been studied.¹ Our results also speak to the “co-holding puzzle”—that many households simultaneously hold high-interest debt and low-interest assets (Gross and Souleles, 2002). Existing explanations emphasize the lack of fungibility between debt and credit (Telyukova and Wright, 2008) or insurance against credit limit volatility (Fulford, 2015), but these do not explain why households would pay down *low*-interest debt while earning *higher*-interest savings. Finally, unlike some documented financial “mistakes”—such as failure to refinance (Keys et al., 2016) or failure to shop for mortgages (Woodward and Hall, 2012)—mortgage prepayment

¹Several earlier papers have documented mortgage curtailment (prepayment without refinancing) as a descriptive phenomenon. Amromin et al. (2007) estimate \$1.5 billion in losses from prepayment in the 1990s, when the losses arose primarily from the tax arbitrage between mortgage interest deductions and tax-deferred retirement accounts—a channel reduced by the Tax Cuts and Jobs Act of 2017 (McClelland et al., 2022). Adelman et al. (2010) link prepayment to saving propensity and liquidity risk. McCollum et al. (2015) document continued curtailment using loan-level records for 2001–2011, and Xu (2023) provides additional descriptive evidence using Fannie Mae data. Choi (2022) reviews personal finance advice on curtailment, noting that such advice may have been reasonable when mortgage rates were high but is harder to justify in the current environment. None of these studies quantify curtailment in the post-2022 regime in which risk-free yields exceed fixed-rate mortgage coupons, making curtailment a strictly dominated strategy. Relative to this prior work, our contribution is three-fold: we identify a setting where prepayment is unambiguously costly, we connect prepayment to debt aversion as a mechanism rather than treating it as a descriptive puzzle, and we exploit credit registry data to link mortgage curtailment to broader household financial behavior.

is an *active* choice requiring deliberate action, making it more informative about household preferences and harder to address through nudges alone (Bhutta et al., 2020).

The remainder of the paper is organized as follows. Section 1.2 derives the NPV of mortgage prepayment and establishes the frictionless benchmark prediction that no household should prepay when the risk-free rate exceeds its mortgage rate. Section 1.3 describes our data sources. Section 1.4 documents four empirical patterns: the prevalence and cost of negative-spread prepayment, its insensitivity to interest rate changes, its sensitivity to cash flow shocks, and the spike in prepayment near mortgage completion. Section 1.5 develops three candidate explanations—attention costs, debt aversion, and cash flow constraints—and presents evidence that debt aversion best accounts for the data. Section 1.6 concludes.

1.2 Mortgage Valuation and the Cost of Prepayment

A conventional fixed-rate mortgage with a monthly interest rate r_m and a term of N months can be viewed as a stream of N fixed payments. Let c denote the equal monthly payment. The principal value P of this mortgage is then the present value of all future payments:

$$P = \sum_{t=1}^N \frac{c}{(1 + r_m)^t}. \quad (1.1)$$

Households, however, have the option to *curtail* the loan—that is, make a payment larger than c in a given month. Unlike a standard extra payment on credit card debt, prepayment via curtailment reduces only the final month’s payment (rather than all future payments). As we will show, this feature has important implications for the net present value (NPV) of the decision to curtail.

Consider a prepayment of amount Δ made when N payments remain. The principal balance of the loan falls by Δ , thereby reducing the terminal payment by an amount γ that can be derived from the mortgage amortization formula. Specifically, letting P denote the remaining

principal balance,

$$P - \Delta = \sum_{t=1}^{N-1} \frac{c}{(1+r_m)^t} + \frac{c-\gamma}{(1+r_m)^N}. \quad (1.2)$$

Solving for γ yields

$$\gamma = \Delta(1+r_m)^N.$$

Thus, a partial prepayment now lowers the terminal payment by γ , but the earlier monthly payments c remain unchanged.

To evaluate the financial cost or benefit of such a prepayment, denote by r the household's opportunity cost of capital (e.g., the after-tax yield on an alternative investment). The net present value of a curtailment of size Δ is then

$$\text{NPV} = \Delta \frac{(1+r_m)^N}{(1+r)^N} - \Delta. \quad (1.3)$$

Whenever $r_m < r$, this expression is negative because the present value of the reduced final payment, discounted at rate r , is below the return that could be earned at rate r . In other words, if $r_m - r < 0$, a household strictly prefers to invest Δ at the higher outside rate r rather than to pay down its mortgage. Conversely, if $r_m > r$, prepaying (or refinancing) produces a positive NPV because the mortgage's effective rate of return exceeds r . Prepaying will always have zero NPV for an adjustable-rate mortgage having $r \equiv r_m$.

Naturally, r_m may exceed r if, for example, long-term market rates have fallen below the contractual mortgage rate since origination. In that case, prepayment has a positive NPV, but refinancing may be even more beneficial than prepayment. Our analysis focuses on the case $r_m < r$, in which case both refinancing and prepayment are unambiguously suboptimal from a purely financial perspective. We thus abstract from the option-like nature of refinancing.²

Prepayment is also a relatively illiquid investment compared with risk-free assets, since ex-

²See, for instance, Agarwal et al. (2013) for a derivation of optimal refinancing rules and Keys et al. (2016) for empirical evidence on refinancing decisions.

tracting equity from a house requires fees, a higher interest rate, and time costs. One caveat is that borrowers can *recast* their mortgage—re-amortize the remaining balance to obtain a lower monthly payment. However, industry evidence suggests that recasting is extremely rare due to stringent lender requirements (Lerner, 2021; Brown, 2023).

A frictionless benchmark

The NPV formula above implicitly defines a frictionless portfolio problem. A household receives a liquid cash flow F and allocates it between (i) an outside asset earning an after-tax, risk-free return r , and (ii) an extra principal payment on a fixed-rate mortgage earning an after-tax return r_m . Let $x \in [0, 1]$ be the share of F allocated to prepayment. Expected wealth next period is:

$$W(x) = xF(1 + r_m) + (1 - x)F(1 + r). \quad (1.4)$$

Maximizing $W(x)$ implies $x = 1$ when $r_m > r$ and $x = 0$ otherwise. The benchmark therefore predicts (i) no dominated prepayment, and (ii) an aggregate prepayment share that jumps discontinuously when the spread $r_m - r$ crosses zero.

The frictionless benchmark predicts that no household should prepay its mortgage when a higher after-tax return is available at the risk-free rate. The next sections show that this prediction is violated. We first describe our data (Section 1.3), then document four empirical patterns that characterize prepayment behavior (Section 1.4): its prevalence and cost, its insensitivity to interest rate changes, its sensitivity to cash flow shocks, and an increase near mortgage completion. Section 1.5 then develops three candidate explanations—attention costs, debt aversion, and cash flow constraints—and assesses which best fits the evidence.

Importantly, we think of this benchmark as the prediction of a highly-simplified model, not as a normative welfare criterion. In a model where households receive utility from reducing debt, negative-NPV prepayment is not necessarily a “mistake”. Rather, the foregone return measures

how strong that preference must be to rationalize the observed choice.

1.3 Data and Measurement

Mortgage data. Our main source of data is the Fannie Mae Single-Family Loan Performance dataset, which includes monthly loan-level performance history for fully amortizing, single-family fixed-rate conforming mortgages with terms of 30 years or less. We include both purchase and refinance loans originated on or after January 1, 1999, exclude loans that were ever delinquent, and limit the sample to January 2011–November 2024. We take a 1% random sample for computational convenience.

The dataset includes variables measuring the principal balance, interest rate, maturity at origination, and a recalculated maturity that accounts for past prepayment. We use the mortgage amortization formula to calculate payments in excess of the required monthly payment. Because principal balances are rounded to the nearest \$1,000, the formula sometimes predicts prepayment due to rounding error; we therefore ignore prepayment amounts less than \$20, which exceeds the maximum measurement error. We verify the calculation by confirming that changes in remaining maturity match expected values and by recovering known prepayment in a simulated dataset. Because we identify prepayment from changes in the outstanding principal balance rather than from total payment amounts, escrow payments for property taxes and homeowner’s insurance—which do not reduce principal—cannot be misclassified as prepayment in our data.

We study adjustable-rate mortgages using the Freddie Mac Single-Family Non-Standard Dataset, which has a similar structure but includes ARMs.

Tax rates. We use the 2010–2019 Surveys of Consumer Finances (SCF) and the NBER TAXSIM program to calculate marginal tax rates. We subset SCF respondents to those with fixed-rate mortgages, calculate their federal and state marginal tax rates, and match these rates to loans

in the Fannie Mae data by year, origination principal, and state.³ For our baseline analysis, we assume that households with a tax-deferred retirement plan available face an effective marginal tax rate of zero on investment returns. If households were contributing to the maximum annual limit to such accounts, then their marginal tax rate would be higher; however, households typically do not do so: as of 2024, only about 14% of households contribute the maximum amount to their tax-deferred retirement accounts each year (Vanguard, 2024). We then calculate the after-tax mortgage interest rate (r_m in Equation 1.3) and after-tax maturity-matched Treasury yield (r in Equation 1.3). We also present results under alternative tax assumptions.

Discount rates. Our baseline results use maturity-matched Treasury rates as the risk-free rate: the 10-year Treasury note yield for mortgages with five or more years remaining, the 5-year note for two to five years, the 2-year note for one to two years, and the 1-year Treasury bill for less than one year. We also show results using alternative discount rates.

Consumer credit records. Other credit variables come from the University of California Consumer Credit Panel (UC-CCP), which consists of loan-level credit records for a 2% sample of the U.S. population, observed quarterly from 2004 to 2023. We restrict our sample to mortgage borrowers with information on actual and scheduled payment amounts, so we do not need the amortization formula to identify prepayment. The UC-CCP also contain location information as well as borrowers' credit scores and balances on other financial products each quarter, allowing us to link mortgage prepayment to broader financial behavior.⁴

The main limitation of the UC-CCP is that it does not include interest rates, so we cannot calculate the NPV of prepayment in these data. A secondary limitation is the inability to distinguish true prepayment (which reduces remaining maturity) from “clerical” early payments (e.g., making next month’s payment before its due date). The Fannie Mae data allow us to

³Since the mortgage data do not include the state of residence, we randomly assign mortgages to states proportional to state population. See Appendix A for details.

⁴Similar to Amromin et al. (2020) and Adelman et al. (2010), we use credit card ownership and balances to predict mortgage prepayment, but we rely on administrative data rather than self-reported values.

remove clerical prepayment (see Appendix A), but this is not possible in the UC–CCP so the credit-registry series indicates a somewhat higher level of prepayment. Figure A.2 illustrates the discrepancy: the benchmark Fannie Mae prepayment flag and the UC–CCP flag differ, but when we add back clerical prepayers in the Fannie Mae data, the two series converge.

1.4 Empirical Patterns

The frictionless benchmark in Section 1.2 predicts that no household should prepay when $r_m < r$, and that aggregate prepayment should respond to changes in the interest rate spread. We now examine the data along four dimensions that bear on this prediction and help distinguish between candidate explanations. First, we document the prevalence and cost of negative-spread prepayment. Second, we test whether prepayment responds to interest rate changes. Third, we ask whether prepayment responds instead to cash flow shocks. Fourth, we examine prepayment near mortgage completion. Section 1.5 then develops candidate explanations formally and assesses which best fits the evidence.

1.4.1 Prevalence and Cost of Negative-Spread Prepayment

We begin by documenting the scale of prepayment when it is financially costly. Panel A of Figure 1.3 shows the distribution of after-tax interest rate spreads ($r_m - r$) in the Fannie Mae data. During 2011–2021 (blue), the mass is almost entirely positive, as outstanding mortgage rates were well above risk-free yields. After 2022 (red), more than half of observations are negative. As an illustrative exercise, Panel B shows the hypothetical NPV of a \$1 prepayment for a mortgage with an interest rate of 4% and 100 months remaining. Holding r_m and Δ fixed, we vary r from 0.01% to 8%, corresponding to spreads of about -4 to $+4$ percentage points. Mechanically, for negative interest rate spreads, the net present value is negative. The slope of the line of best fit is approximately 0.08, indicating that for every percentage-point decrease in the spread, households lose nearly 10 cents per dollar prepaid—or about 10% of their excess

payment.

Despite this shift in the interest rate environment, prepayment behavior barely changed. Figure 1.1 plots the distribution of outstanding mortgage rates against the 10-year Treasury yield (Panel A) and total prepayment volume by spread sign (Panel B). The volume of negative-spread prepayment (red bars) grows rapidly after early 2022, but positive-spread prepayment does not decline commensurately—total prepayment volume is roughly stable. Among borrowers with a negative spread, about 22% prepay in any given month.

The upper panels of Figure 1.4 show histograms of excess payment amounts for 2011–2021 and 2022–2024, right-truncated at \$600. The distribution has a very long right tail, so the average payment (\$556) is substantially larger than the median. Most prepayments are small, with pronounced bunching at round numbers—especially \$100 and \$200—which reflects deliberate payment choices rather than data rounding, since the excess amount is calculated relative to the expected amortization payment. Aggregate losses are therefore driven disproportionately by the minority making large payments. Panels C and D scale prepayment by the loan’s monthly payment, right-truncated at 1 (capturing about 90–95% of payments). The pattern is similar in both time periods: the mean and median prepayment as a share of the monthly payment are about 23% and 14%, respectively. Panels E and F show the NPV of these payments. From 2011–2021, nearly all excess payments carried a positive NPV, since most mortgages had a positive rate spread.⁵ From 2022–2024, the distribution shifts inward, with over half of observed payments carrying a negative NPV.

Table 1.1 provides summary statistics from the Fannie Mae dataset at the loan-month level, comparing prepayers and non-prepayers. Panel A covers the full sample period (January 2011–November 2024); Panel B restricts to January 2022–November 2024; Panel C further restricts to negative-spread mortgages in the later period. Across the full sample, prepayers and non-prepayers look remarkably similar on most dimensions: origination balances of

⁵Of course, during 2011–2021, the NPV from refinancing could have been even higher.

roughly \$220,000, remaining balances near \$180,000, and comparable marginal tax rates on both income and mortgage interest (about 12% each). Prepayers have slightly lower mortgage interest rates (3.93% vs. 4.10%) and shorter origination terms, on average. Naturally, prepaying households have shorter effective remaining terms: 219 months vs. 259 months across all years, and 228 vs. 265 months in the 2022–2024 subsample. Among negative-spread loans in Panel C, the differences are larger: prepayers have origination balances of about \$268,000 and remaining balances of \$222,000, compared to \$281,000 and \$246,000 for non-prepayers. The key observable distinction throughout is that prepayers tend to have higher credit scores, a pattern we return to in Section 1.5.

How large are the aggregate losses? Figure 1.5 shows the time series. There was almost no negative-NPV prepayment prior to February 2022. By the fourth quarter of 2023, total losses reached about \$250 million per month, implying an average cost of roughly \$115 per prepaying borrower. The cumulative total from January 2022 to November 2024 is approximately \$4.6 billion for GSE-backed mortgages, with an average NPV of nearly $-\$87$ per payment.

Another way to gauge the cost is in terms of forgone return. Figure 1.6 shows the average rate spread for negative-spread prepayers, weighted by dollar volume: this increased in absolute value from about -0.2 to nearly -1.25 percentage points over the period. On average, these borrowers forgo nearly a full percentage point of return.

1.4.2 Assumptions

Our calculations rely on assumptions about tax benefits of mortgage interest and the appropriate opportunity cost of capital. Table 1.2 shows how different assumptions affect the results.

The first row reports our benchmark: investment returns are taxed at the marginal income tax rate, except for households with an eligible tax-deferred retirement account (for whom the effective tax rate is zero); mortgage interest is tax-deductible for itemizers; and the maturity-

matched Treasury rate is the opportunity cost. Under these assumptions, 52% of mortgages carry a negative spread from January 2022–November 2024, and total losses are \$4.6 billion.

The availability of tax-free retirement accounts matters substantially. The second row assumes all investment returns are taxed at τ , removing the tax-free channel. Because the average marginal tax rate for prepayers is about 12%, this reduces the effective 10-year Treasury return from roughly 4% to about 3.5%. Since this yield is close to the median mortgage rate, even a half-percentage-point shift substantially changes the share of mortgages with a negative spread: it falls from 52% to 39%, and total losses fall by more than half, to \$2.2 billion. In the third row, we ignore taxes entirely, which raises losses to \$6.3 billion.

Finally, although the maturity-matched Treasury rate is a reasonable benchmark, mortgage payments should arguably be discounted at a somewhat higher rate due to refinancing and default risk. When we use the yield on maturity-matched mortgage-backed securities instead, the negative-spread share rises to 69% and the average loss per payment nearly doubles to \$120. These estimates suggest that our benchmark may be a lower bound on the true magnitude of losses. Under all assumptions, the aggregate cost is substantial.

1.4.3 Insensitivity to Interest Rates

The frictionless benchmark predicts a strong link between interest rate movements and prepayment. Figure 1.1 already suggests little aggregate correlation. We now test this formally.

We estimate the linear probability model:

$$Prepay_{it} = \beta \text{Spread}_{it} + \Omega X_{it} + \varepsilon_{it}, \quad (1.5)$$

where Spread_{it} is the difference between the mortgage rate and a comparable risk-free rate. Table 1.5 reports results. In the cross-section (Column 1), each percentage-point higher spread is associated with a 1 pp *lower* probability of prepayment—the wrong sign. Adding loan fixed

effects (Column 2), so that time-series variation in r identifies the effect, the coefficient falls to about 0.3 pp and is economically negligible. Column 3 confirms the cross-sectional pattern: households with lower mortgage rates prepay more.

Two robustness exercises reinforce the null. First, we add annual lags of the spread up to four years (Table A.1); all lagged coefficients are precisely estimated zeros. Second, a binscatter (Figure A.4) reveals a flat relationship across the entire range of spreads, in both the full sample and the 2022–2024 subsample. There is no evidence of a threshold or nonlinearity that might indicate switching costs or inaction bands.

1.4.4 Sensitivity to Cash Flows

If prepayment does not respond to interest rates, what explains its time-series variation? We examine whether disposable cash flow—rather than the present value of future payments—drives extra mortgage payments.

Income shocks. Figure 1.7 plots total monthly prepayment alongside real disposable income growth for high-income (top 10%) and middle-income (middle 40%) households, using data from Blanchet et al. (2022).⁶ Among the top 10%, prepayment closely tracks income gains and shows abrupt jumps following COVID-19 relief payments (dashed vertical lines).

ARM vs. FRM comparison. Adjustable-rate mortgages provide a more direct test. When market rates rise, ARM borrowers face both a higher discount rate *and* a higher monthly payment. FRM borrowers face only the former. If households respond to current cash flow rather than NPV, ARM prepayment should decline more.

We estimate:

$$\text{Prepay}_{it} = \beta \cdot \text{ARM}_i \times \text{LIBOR}_t + \delta_i + \gamma_t + \varepsilon_{it}, \quad (1.6)$$

⁶Data are downloaded from <https://realtimeinequality.org/>.

using a sample of FRMs and ARMs past their introductory fixed-rate period. Table 1.6 reports the results. Column 1 confirms that a 1 pp LIBOR increase raises ARM rates by about 0.644 pp, corresponding to a roughly \$52 increase in monthly payments. Column 2 shows that this reduces ARM prepayment by 1.9 pp, while FRM borrowers—whose payments are unchanged—show no effect. Columns 3–5 decompose by payment size: small-dollar prepayments (\$20–\$500) decline most, consistent with a liquidity mechanism.

Comparing ARMs still in their introductory period (“locked in”) with those already floating reinforces this pattern. Table 1.7 shows that after the lock-in period, LIBOR fluctuations pass through fully to payments and generate a stronger negative effect on prepayment.

We confirm the timing using a local projection around ARM rate resets:

$$\text{Prepay}_{i,t+k} - \text{Prepay}_{i,t-1} = \beta_k(\text{Rate}_{i,t} - \text{Rate}_{i,t-1}) + \gamma X_{it} + \varepsilon_{it}, \quad k = -2, 0, 1, 2. \quad (1.7)$$

Figure 1.8 plots the estimated β_k coefficients. A 1 pp rate increase reduces prepayment probabilities by 0.5–1.1 pp over the following two months, relative to a baseline of 18%.

1.4.5 Prepayment Near Mortgage Completion

We next examine how prepayment varies with the remaining mortgage balance. Figure 1.9 plots the probability of prepaying against remaining principal using UC–CCP data.⁷ For large balances, prepayment occurs at a roughly constant rate of about 26%. As the balance falls below roughly \$5,000, the probability rises sharply: at \$500 remaining, it reaches nearly 50%. Even between \$500 and \$2,000, the probability at the lower balance is 58% higher.⁸

Is this spike driven by other factors that change near mortgage completion, such as shifts in consumption opportunities or credit behavior? Figure 1.10 plots four placebo outcomes—

⁷The baseline probability is higher in this dataset than in Fannie Mae because we cannot remove “clerical” prepayers. See Section 1.3 and Appendix A.

⁸Fu (1998) finds a similar pattern using data from the 1960s–1990s.

originating a new auto loan, opening a new credit card, becoming delinquent on any trade line, and revolving credit utilization—against remaining mortgage balance. None show a similar spike. We formalize this in a within-borrower regression:

$$y_{it} = \sum_{j=1}^7 \beta_j \mathbb{1}(\text{bal}_{it} \in B_j) + \delta_i + \gamma_t + \varepsilon_{it}, \quad (1.8)$$

where $\text{bal} \in [0, 1]$ is the share of remaining principal and the B_j are bins from 0–1% through 40–50%, with balances above 50% as the omitted category. Table A.2 confirms that, even controlling for borrower and time fixed effects, prepayment rises near completion while the placebo outcomes do not. Because Equation (8) includes borrower fixed effects, this comparison is within-borrower, across time: it compares the same borrower to herself as the remaining balance declines.

1.5 Mechanisms and Interpretation

The patterns in Section 1.4 present a puzzle. A substantial share of borrowers make prepayments that are financially costly (Section 1.4.1), do not adjust when the cost increases (Section 1.4.3), do respond to cash flow shocks (Section 1.4.4), and are far more likely to prepay when full repayment is within reach (Section 1.4.5). We now develop three candidate explanations—attention costs, debt aversion, and cash flow constraints—and assess which can account for these facts.

1.5.1 Candidate Explanations

We consider three mechanisms, not mutually exclusive, that can rationalize prepayment even when it is not NPV-maximizing.

Attention and search costs. Households may face costs of monitoring interest rates or finding alternative uses of funds. A tractable representation is a fixed adjustment cost $\kappa > 0$: the

household prepays only if the perceived gain exceeds that cost,

$$(r_m - r)F \geq \kappa. \quad (1.9)$$

This creates an inaction band: small rate changes leave behavior unchanged, while sufficiently large changes trigger adjustment (Andersen et al., 2020). A closely related microfoundation is rational inattention, in which households choose information precision at a cost (Reis, 2006; Berger et al., 2024). The key predictions are: (i) eventual sensitivity to sufficiently large rate changes, (ii) response to lagged as well as current rates, and (iii) more attentive or sophisticated borrowers should be *more* responsive to rate changes. Additionally, when a household receives an increase in available cash, it should be *less* likely to prepay under this theory, since the benefits of finding higher-return uses of funds increase.

The insensitivity of prepayment to interest rate changes mirrors a related finding in the mortgage origination literature: Anenberg et al. (2023) show that household borrowing amounts also respond little to changes in interest rates. Our setting is particularly well-suited to study this insensitivity. Changes in bond yields or CD rates alter the opportunity cost of paying down a fixed-rate mortgage without changing the mortgage’s own characteristics, allowing us to isolate household sensitivity to economic incentives while holding mortgage terms constant. Since mortgages constitute the largest component of household debt and excess repayments affect roughly one-third of mortgages each year, understanding this insensitivity is important for both household finance and monetary policy transmission.

Debt aversion. Attention-based stories imply that households *would* change behavior if they knew about outside opportunities. Debt aversion offers a different explanation, rooted in preferences. Experimental evidence shows that consumers overvalue debt reduction per se, even at the cost of higher-yielding opportunities (Loewenstein and Prelec, 1992; Prelec and Loewenstein, 1998; Meissner, 2016; Duffy and Orland, 2020; Meissner and Albrecht, 2022; Martínez-

Marquina and Shi, 2024). Beyond the lab, evidence has come from field experiments (Field, 2009; Caetano et al., 2019) and observational studies of student loans (Gopalan et al., 2021; Eckel et al., 2007), where results have been mixed. No previous research has documented the costs of debt aversion in a setting as widespread as ours.

We model a psychological cost of carrying mortgage debt D . Let x be the share of cash flow F devoted to prepayment. The household maximizes next-period wealth net of this cost:

$$V(x) = W(x) - u(D - xF), \quad (1.10)$$

where $W(x)$ is defined in Equation (1.4) and the cost function satisfies $u' > 0$, $u'' < 0$. The first-order condition for an interior solution is

$$r_m - r = -u'(D - x^*F), \quad (1.11)$$

so prepayment can occur with a negative return spread whenever the spread is compensated by the marginal utility of reducing debt. If F is small, the first-order condition is unlikely to hold exactly, and households are more likely to choose a boundary solution at $x^* = 0$ or $x^* = 1$ (depending on which side of Equation 1.11 is greater). An additively separable debt-utility term of this sort arises naturally from the mental-accounting framework of Thaler (1990) and Prelec and Loewenstein (1998).

The concavity of u implies that the marginal benefit of debt reduction is higher when the balance is low. Martínez-Marquina and Shi (2024) provide experimental evidence for exactly this pattern: households are more willing to sacrifice returns to pay off debt when balances are smaller. This generates a distinctive prediction: prepayment should spike near mortgage completion, when the psychological benefit of full elimination is largest. Standard portfolio theory predicts no such change, since the financial return to prepaying does not vary with the remaining balance.

Two additional predictions follow. First, if debt aversion is a stable preference, households that prepay their mortgage should also exhibit similar behavior on other debts. Second, although debt aversion does not directly predict sluggish responses, it can generate low *aggregate* sensitivity to rate changes if many households are inframarginal—that is, if heterogeneity in debt aversion means few households are near the margin between prepayment and investment.

Cash flow constraints. In mortgage modification settings, Ganong and Noel (2020) show that changes in NPV have little effect on behavior, whereas changes in short-run cash flow have large effects.⁹ By analogy, households that ignore NPV may still respond to cash flow shocks. The prediction is straightforward: prepayment should track liquidity shocks regardless of NPV. An important implication is that the observed prepayment rate is a *lower bound* on the prevalence of debt aversion, since cash-constrained households may be unable to act on their preference even if they hold it.

Mortgage insurance One potential benefit of prepayment is earlier cancellation of mortgage insurance for high loan-to-value (LTV) borrowers. If mortgage insurance were an important driver of the patterns we document, prepayment should rise discretely around the LTV threshold at which that insurance can be removed – namely, when the LTV ratio is greater or equal to 80. To test this, we plot the probability of prepaying against LTV at origination in Figure A.6. Origination LTV does not vary over time in the data, so we restrict the sample to include only observations from the first 24 months of each loan’s life. This approach allows us to isolate the time in the amortization schedule for which households paying mortgage insurance would be most incentivized to prepay. We do not see such a discontinuity around 80 percent LTV, suggesting that mortgage insurance is not a first-order explanation for the behavior in our data. Further, removing loans with an $LTV \geq 80$ does not qualitatively change any key results: the probability of having a negative rate spread in Jan 2022-Nov 2024 is .56 (vs .52 for all loans);

⁹Cookson et al. (2022) provide related evidence that household responses to cash flows cannot be explained by standard forward-looking models.

the probability of prepaying conditional on a negative rate spread is .24 (vs .22 for all loans); and the average NPV loss of such prepayments in that period is -\$88.12 (vs. \$86.80). In sum, while mortgage insurance may induce prepayment for some households, it is not a key driver of the results nor of the behavior in general.

1.5.2 Evidence for Debt Aversion

We now bring together evidence from Section 1.4 and additional cross-sectional tests to assess which explanation best fits the data.

Stable household-level preferences. A natural starting point is whether prepayment reflects stable individual characteristics or transitory shocks. Table 1.4, Panel A, decomposes the variation in prepayment using UC–CCP data. Individual fixed effects alone explain about 25% of the variation ($R^2 = 0.255$); adding time fixed effects raises this to 0.376. Prepayment is highly persistent: conditional on prepaying in month t , the probability of prepaying in $t + 1$ exceeds 80%. Conditional on *not* prepaying, the probability in $t + 1$ is below 5%.

Within-borrower variation reinforces the pattern. Adding credit covariates (Column 3), higher credit scores are associated with more prepayment: a 100-point increase raises the probability by 4.8 pp. Credit card ownership and larger credit card balances predict *less* prepayment. Including individual fixed effects (Column 4), these coefficients strengthen: the credit score coefficient rises to 0.114, and the credit card indicator increases in magnitude from -5.2 to -8.7 . When a household’s credit score improves, it becomes more likely to prepay; when it acquires a credit card, it becomes less likely. These within-borrower patterns point to a stable relationship between aversion to debt across products.

Cross-product consistency. If debt aversion is a preference rather than a product-specific mistake, it should manifest across debt types. Panel B of Table 1.4 tests this using auto loans. Households that prepay their mortgage are 8.5 pp more likely to also prepay an auto loan in

the same quarter, a roughly 50% increase relative to the 17% base rate. These households also choose shorter auto loan terms—about 2–3 months shorter on both average and maximum term length. The cross-product correlation holds conditional on age, credit score, gender, and state fixed effects, indicating that it does not simply reflect higher income or better financial position.

Prepayment near completion and concave disutility. Section 1.4.5 documented that the probability of prepaying rises from about 26% to nearly 50% as the remaining balance approaches zero (Figure 1.9), while four placebo credit outcomes show no such spike (Figure 1.10). This pattern is consistent with the prediction of concave $u(D)$: the psychological benefit of eliminating debt entirely may exceed the benefit of a marginal reduction (Martínez-Marquina and Shi, 2024; Meissner and Albrecht, 2022).

Standard portfolio theory predicts no change near completion, since the financial return to prepaying does not vary with the balance. Several alternative explanations receive less support in the data. Retirement timing is also unlikely to explain the completion spike. Panel B of Figure 1.2 shows prepayment probabilities by borrower age, calculated from UC–CCP data. The profile follows an inverse-U shape: the likelihood of prepayment rises steadily from the early 30s through the mid-50s and then declines through the mid-70s. This hump-shaped pattern closely mirrors the lifecycle profile of household income, which peaks in the late 40s to mid-50s before declining into retirement. Crucially, there is no discontinuity at age 65, when the probability of retirement increases sharply. If mortgage prepayment were driven by a desire to eliminate debt at retirement, we would expect a discrete jump in prepayment at or around this age; no such jump exists. Rather, the age profile is consistent with prepayment tracking available cash flow over the lifecycle: households prepay more when income is high and less when it falls, regardless of proximity to retirement.

The rise in repayment near completion hardly changes after adding borrower and time fixed

effects Table A.2), which absorb any age-related or cohort-related confounders. Credit score incentives also cannot explain the pattern, since credit scores do not discretely improve upon mortgage completion. The spike near completion (as seen in Figure 1.9) is also unlikely to be a byproduct of moving or refinancing. A household can move at any point in the amortization schedule, and refinancing incentives are weakest when the remaining balance is small. More importantly, we replicate this finding in the Fannie Mae data where we do not count the final month of a loan as curtailment, so the result is not mechanically generated by loans disappearing from the data. The rise occurs before payoff, and the same part of the balance distribution does not show a corresponding change in new borrowing, delinquency, or credit utilization.

Financial literacy. If negative-NPV prepayment reflected financial ignorance, we would expect it to be concentrated among less sophisticated borrowers. The evidence is difficult to reconcile with financial illiteracy or lack of access to higher-return alternatives as primary explanations, however. Table 1.3 breaks out prepayment behavior by borrower characteristics. Borrowers with above-median credit scores are *more* likely to have a negative spread (in part because they secured lower rates at origination) *and* are more likely to prepay conditional on having one. They also make larger extra payments. Figure 1.2, Panel A, shows the gradient directly: prepayment rises from about 10% per month for borrowers in the 600s to above 20% for those near 800.

First-time home buyers are *less* likely to engage in negative-spread prepayment, not more. Perhaps most strikingly, borrowers who originated in 2020–2021—a group enriched for attentive households who refinanced when rates were low—are no less likely to prepay a negative-spread mortgage (Table 1.3). These borrowers demonstrated sophistication in one domain (refinancing) while exhibiting the same “costly” behavior in another (prepayment), consistent with a deliberate preference rather than a financial mistake.

Stock market participation. If negative-NPV prepayment reflected a lack of access to higher-return alternatives, we would expect prepayers to be *less* likely to participate in equities. Data from the 2022 Survey of Consumer Finances show the opposite. We classify a household as a prepayer if it reports being “ahead” on its mortgage payment schedule relative to “on track” or “behind.”¹⁰ Households that prepay their mortgage are 27.5% more likely to hold any stock (defined as a positive total market value of directly and indirectly held equities) than non-prepayers. Unconditional on stock ownership, prepayers’ equity holdings are about 36% larger than those of non-prepayers (excluding the top 5% of the stock value distribution, which is dominated by extreme right-tail observations). Even conditional on owning stocks, prepayers hold about 6% more in equities. Prepayers are also 11% more likely to hold an IRA or Keogh retirement account. These patterns indicate that households who engage in costly mortgage prepayment are not simply unaware of or unable to access higher-return investments—they actively participate in equity markets while simultaneously choosing to pay down low-interest debt.

Inattention and inertia. Inattention predicts that sufficiently large rate changes, or sufficiently attentive borrowers, should eventually respond. Neither pattern appears in the data.

We construct two proxies for attentiveness. First, we identify 2020–2021 refinancers, who demonstrated active rate-monitoring. Second, we classify borrowers by the residuals from an auxiliary regression of mortgage rates on credit scores, LTV, and origination-quarter fixed effects:

$$\text{Rate}_i = \alpha + \beta_1 \text{CreditScore}_i + \beta_2 \text{LTV}_i + \gamma_{q(i)} + \varepsilon_i, \quad (1.12)$$

with the lowest-residual quartile representing those who secured especially favorable rates. For both groups, we re-estimate the spread-sensitivity regressions from Section 1.4.3. Figure 1.11, Panel A, shows that even among the most attentive borrowers, the estimated sensitivity of

¹⁰This categorical variable is available in the SCF and captures households that have made payments in excess of the amortization schedule.

prepayment to interest rate spreads is economically negligible—close to zero in all subsamples.

Inertia is also an unlikely explanation. If borrowers began prepaying when it was financially beneficial and simply failed to stop, we would expect those who started in a low-rate environment to show higher continuation rates than those who started in a high-rate environment. Figure 1.11, Panel B, shows the opposite: borrowers who first prepaid after March 2022 (when rates were high and prepayment was costly) converge to the same ~90% continuation probability within two months as those who began in a low-rate environment. Persistence is high but unrelated to the interest rate environment at origination.

Finally, dropping borrowers who make exactly the same payment each month—a group for whom autopay or habit could explain persistence—does not change the results. Table A.3 replicates the rate-sensitivity regressions from Table 1.5 excluding these borrowers; the coefficients are nearly identical.

One additional question is whether some borrowers may follow biweekly or 13th-month payment strategies—making half a payment every two weeks or one extra full payment per year—that mechanically generate “excess” payments without reflecting a deliberate monthly decision. In the distribution of payments as a fraction of the expected monthly payment (Figure 1.4, Panels C–D), such strategies would appear as a mass at approximately 1.5 times the expected payment (corresponding to an extra half-payment in a given month). We do observe some bunching at this ratio, but it is quantitatively small relative to the overall prepayment volume. The vast majority of excess payments do not correspond to round multiples of the scheduled payment, indicating that biweekly plans account for at most a minor share of the behavior we document.

More broadly, the high month-to-month continuation probability (above 80%) might suggest that prepayment is set-and-forget rather than an active monthly decision. However, the fact that prepayment responds to contemporaneous cash flow shocks—rising after stimulus pay-

ments and falling when ARM payments increase—indicates that households do revisit this decision, at least in response to material changes in disposable income. As we have shown, dropping borrowers who make exactly the same dollar payment each month, the group most likely on autopay, also does not change the results.

A related interpretation is commitment demand: households may prepay to tie up cash in home equity rather than leave it available for spending. The evidence on commitment demand is mixed; Sadoff and Samek (2019) emphasize that explicit commitment is often limited in real-world settings, while also showing that information and experience can increase demand for commitment. This motive may therefore explain part of the behavior we document, especially repeated extra payments that look like a saving rule. But it does not fit the totality of the evidence as well as debt aversion. A mortgage already commits the household to regular payments, and paying off the final balance removes that commitment rather than strengthening it. More importantly, commitment predicts a value from locking away wealth, not a sharp rise in repayment when the balance approaches zero. It also does not naturally explain why prepayers have higher credit scores, hold stocks and retirement accounts, and repay other debts similarly. We therefore view commitment as a possible contributing motive, but not the primary explanation for our findings.

Several commonly cited motives for prepayment—a desire to be debt-free before retirement, “peace of mind,” or eliminating the monthly payment—are sometimes presented as alternatives to debt aversion. We argue that they are better understood as *instances* of it. Debt aversion, as modeled in Equation (1.10), encompasses any disutility from holding debt not captured by its financial cost. A household that sacrifices investment returns to eliminate its mortgage before retirement is exhibiting exactly this preference. The question is not whether such households exist—they clearly do—but whether the preference is large enough to generate the patterns we document, and whether it is better described as a stable preference than as a mistake. The evidence in this section suggests that a stable preference-based account fits the data better than

a simple mistake-based account.

1.5.3 Illustrative Bounds on the Prevalence and Intensity of Debt Aversion

We use the stylized framework to interpret what the data imply about the prevalence and intensity of debt aversion. Section 1.4.4 showed that prepayment responds strongly to cash flow shocks: aggregate prepayment tracks income growth for high-income households and declines for ARM borrowers whose monthly payments increase. This has an important implication for interpreting the 22% prepayment rate documented in Section 1.4.1. If some debt-averse households lack the liquidity to act on their preference, the observed prepayment rate understates the true prevalence of debt aversion.

We gauge how loose this lower bound may be using imputed income in the Fannie Mae data. We proxy for income by dividing the scheduled mortgage payment by origination DTI.¹¹ Figure A.5 shows that both the probability and amount of prepayment rise with imputed income. The extensive-margin relationship is clearly concave. Among households high in the imputed-income distribution, the probability of prepaying despite a negative spread is about 24%, only modestly above the unconditional 22% rate. This suggests that liquidity constraints matter, but that the raw lower bound is not far below the true prevalence on the extensive margin. We therefore interpret 24% as a conservative, liquidity-adjusted lower bound on the share of mortgage borrowers who are debt-averse.

The results near mortgage completion point in the same direction. Figure 1.9 shows that prepayment approaches 50% when the remaining balance becomes very small. When the utility gain from eliminating debt is especially large, willingness to prepay becomes much more common. Taken together, the evidence suggests that the debt-averse share is above the raw

¹¹Because DTI is based on required payments on all debt, not just the mortgage, this measure is noisy and likely understates true income. We therefore use it to rank households rather than to measure income precisely.

22% monthly lower bound and plausibly closer to one-quarter in ordinary months.

We can also provide a lower bound on the intensity of debt aversion among debt-averse households. The first-order condition in Equation (1.11) implies that any household prepaying at a negative spread reveals a marginal debt-aversion premium of at least $|u'(D - x^*F)| \geq |r_m - r|$: the household is willing to forgo at least $|r_m - r|$ in annual return per dollar of debt reduced. Under our benchmark assumptions, the median negative spread among prepayers is about 90 basis points. Put differently, the median prepayer behaves as if the mortgage rate were about 90 basis points higher than its contractual rate, or roughly 30% higher relative to the contract rate in this group. This is a lower bound, since inframarginal households may have considerably stronger preferences.

1.6 Conclusion

Consider a mortgage borrower who has extra money each month. She could invest in a 401(k) or other retirement plan; or she could put the extra cash towards her mortgage, helping pay it off faster. Personal finance writer Dave Ramsey’s best-selling *Total Money Makeover* counsels the latter: one of his seven “financial baby steps” is to *pay off your home early*. For Ramsey, ridding oneself of debt is not merely a sound investment—it is a moral imperative.¹² This view is at odds with standard portfolio theory, which holds that no special consideration should be given to whether cash is spent on a new investment or used to pay off debt, so long as the risk-adjusted return is maximized.¹³ Our findings suggest that Dave Ramsey’s audience is large—and that they are willing to pay for the privilege of following his advice.

Between one-fifth and one-quarter of U.S. conforming mortgage borrowers voluntarily prepay each month, and this share does not respond to changes in interest rates. Monetary policy

¹²On the *Dave Ramsey Show*, he states that “the borrower is slave to the lender, voluntarily signing up to give our lives away. Well, this show is about breaking free.” See: <https://www.youtube.com/watch?v=3qZ9CMbtDQA&t=7m30s>.

¹³This holds the liquidity and risk of different options the same.

tightening in 2022 raised risk-free yields above the rate on roughly half of outstanding mortgages, yet borrowers did not reduce prepayment, resulting in \$4.6 billion in losses from GSE mortgages alone—a pure transfer from households to mortgage investors. Prepayment is persistent within borrowers, correlated across debt products, and spikes as the remaining balance approaches zero. Households respond to cash flow shocks but not to changes in the present value of repayment. Taken together, these patterns point to debt aversion—a preference for reducing debt that goes beyond its financial cost—as an important driver of repayment behavior.

Strikingly, negative-NPV prepayment is not concentrated among unsophisticated borrowers. Households with higher credit scores are *more* likely to prepay at a loss, and those who refinanced optimally when rates were low are equally prone to costly prepayment afterward. Survey data confirm that prepayers are more likely to hold stocks and retirement accounts. These patterns are difficult to reconcile with financial illiteracy as the primary explanation and instead point to a deliberate preference.

Our findings have several implications for economic theory. First, debt aversion can generate aggregate unresponsiveness to interest rate changes even without inattention or adjustment costs. If many households are inframarginal with respect to repayment, marginal changes in rates will be insufficient to alter portfolio allocations, producing a “sluggish” aggregate response of the sort often attributed to sticky information or rational inattention (Sims, 2003; Reis, 2006). Second, households that are averse to borrowing may act liquidity-constrained despite having unused borrowing capacity, reluctant to smooth consumption through additional debt (Thaler, 1990). This could help explain why household portfolios appear suboptimally “safe” from the perspective of traditional lifecycle theory (Cocco et al., 2005). Third, the scale of prepayment—roughly one-third of mortgages each year—means that portfolio distortions from debt aversion are not confined to a small set of unsophisticated households but represent a first-order feature of household balance sheets.

An important open question is why debt-averse households borrow in the first place. If carrying debt generates disutility, one might expect such households to avoid mortgages altogether or to borrow less at origination. One possibility is that debt aversion operates more strongly on existing balances than on the decision to take on new debt, consistent with the concave disutility function in our model: the psychological cost of holding a specific debt balance may be more salient than the *ex ante* prospect of future indebtedness. Another is that housing market constraints leave even debt-averse households with little alternative to mortgage financing. Studying how debt aversion shapes borrowing *entry* decisions is a natural direction for future work.


Our results also connect to work on tax salience (Chetty et al., 2009). Just as consumers respond to posted prices rather than after-tax prices, mortgage borrowers appear to respond to their monthly payment obligations rather than to the opportunity cost of prepayment. The lack of salience of investment returns has direct implications for how monetary policy transmits to household borrowing decisions: even large changes in the risk-free rate may fail to alter mortgage repayment behavior if households do not attend to the spread between their mortgage rate and available returns.

Finally, unlike most documented financial “mistakes,” prepayment is an *active* choice. Failure to refinance, failure to shop for a mortgage, and failure to rebalance a portfolio all involve inaction, and can in principle be addressed through default options, reminders, or other nudges. Prepayment requires the opposite: households must deliberately send extra money each month. Changing this behavior would require not a nudge but a more direct intervention—one that addresses the underlying preference for debt reduction rather than simply making the optimal choice easier to find.


1.7 Figures and Tables

Figure 1.1: Mortgage prepayment and Treasury yields

(a) Rates on outstanding mortgages and average 10-year Treasury yield




(b) Total excess payments by loans with above/below 10yr Treasury yields




Note: Panel (a) shows the 10th, 50th, and 90th percentile for outstanding mortgage interest rates in black/dashed lines and average 10-year Treasury yield in gray, by month. Panel (b) shows total volume of prepayment amounts each quarter. The blue represents the share of prepayment volume from loans with mortgage interest rates above the prevailing 10-year Treasury yield, and the red shows the share from loans with below the 10-year Treasury yield. Interest rate data come from Fannie Mae and Treasury data come from FRED.

Figure 1.2: Probability of prepayment by borrower characteristics

(a) Probability of prepayment by borrower credit score




(b) Probability of prepayment by borrower age




Note: Panel (a) shows the likelihood a household prepays its mortgage by credit score of borrower at origination. Credit scores are rounded to the nearest 10. Panel (b) shows the likelihood a household prepays its mortgage by age of borrower. Authors' calculations based on Fannie Mae Single-Family Loan Performance dataset and the UC-CCP.

Figure 1.3: Interest rate spreads and net present value of prepayments

(a) Distribution of interest rate spreads




(b) Hypothetical NPV of a \$1 prepayment




Note: Panel (a) shows the distribution of interest rate spreads for 2011-2021 (blue) and 2022-2024 (red). Panel (b) shows the hypothetical NPV of a \$1 prepayment with 100 months remaining against different values of r . Authors' calculations based on Fannie Mae Single-Family Loan Performance dataset.

Figure 1.4: Distribution of excess payments



Note: The top row shows the distribution of prepayment amounts in dollar values. The middle row shows the distribution of prepayment amounts but scaled by the monthly payment. The bottom row shows the NPV of these payments. The left column is 2011-2021, and the right column is 2022-2024. NPV is calculated using tax-adjusted mortgage interest rates and tax-adjusted maturity-matched Treasury yields.


Figure 1.5: NPV of excess payments for negative interest rate spread mortgages




Note: Panel (a) shows the total NPV losses by per month for loans with negative interest rate spreads. Panel (b) shows the average NPV of these payments. NPV is calculated using tax-adjusted mortgage interest rates and tax-adjusted maturity-matched Treasury yields

Figure 1.6: Negative rate spread statistics over time

(a) Share of negative rate spread households prepaying




(b) Average negative rate spread (weighted by prepayment amount)




Note: Panel (a) shows share of all households that have a negative interest rate and who prepay by month. Panel (b) shows the average interest rate spread for negative rate spread loans, weighted by the dollar value of prepayment.

Figure 1.7: Total prepayment amounts and real disposable income growth

(a) Top 10%




(b) Middle 40%




Note: Green and red lines are real disposable income growth for the top 10% and middle 40% of the income distribution, respectively. Blue line is total prepayment amounts per month in millions of dollars. Dashed lines represent months when COVID-19 stimulus checks were sent (April 2020, January 2021, March 2021). Income data come from Realtime Inequality.

Figure 1.8: Probability of prepayment around interest rate changes




Note: Figure plots estimates of β_k 's from the local projection in Equation (1.7). Specifically, we estimate $Prepayment_{i,t+k} - Prepayment_{i,t-1} = \beta_k(Rate_{i,t} - Rate_{i,t-1}) + \gamma X_{it} + \varepsilon_{it}$ for $k = -2, 0, 1, 2$. Authors' calculations based on Freddie Mac Single-Family Non-Standard dataset.

Figure 1.9: Probability of prepaying by remaining principal balance



Note: Figure shows the likelihood of a loan prepaying by remaining principal balance. Balances are rounded to the nearest \$500 and capped at \$300,000. Authors' calculations based on the University of California Consumer Credit Panel (UC-CCP).


Figure 1.10: Placebo credit outcomes by remaining mortgage principal balance




Note: All figures show mean outcome against remaining principal balance. Balances are rounded to the nearest \$500 and capped at \$300,000. Panel (a) plots the probability of originating a new auto loan. Panel (b) plots the probability of opening a new credit card. Panel (c) plots the probability of current delinquency on any trade line. Panel (d) plots total credit utilization. Authors' calculations based on the University of California Consumer Credit Panel (UC-CCP).

Figure 1.11: Sophistication, attentiveness and inertia

(a) Effects of interest rate spreads on prepayment by measures of attentiveness



(b) Conditional continuation probabilities by timing of first prepayment



Note: In Panel A, each dot represents the coefficient on *Spread* from Equation (1.5) using different samples. The blue dots correspond to samples by quartiles of residuals from Equation (1.12). The red dots correspond to the 2020-2021 purchase and refinance samples. Panel B shows the probability of prepaying at time $t + k$ month conditional on prepaying in the k previous months. The blue line represents loans that first prepaid in a high-rate environment (March 2022-November 2024). The red line represents loans that first prepaid in a low-rate environment (Jan 2011-February 2022). Authors' calculations based on the Fannie Mae Single-Family Loan Performance dataset.

Table 1.1: Summary statistics

Panel A: All loans (Jan 2011 -Nov 2024)						
	All		Curtailers		Non-Curtailers	
	Mean	SD	Mean	SD	Mean	SD
Origination balance	220,101	127,033	217,338	124,026	220,740	127,710
Current balance	178,191	130,929	178,730	123,144	178,067	132,665
Term (in months)	306	83	293	88	309	95
Months to legal maturity	260	95	243	97	264	95
Effective months to maturity	251	100	219	103	259	98
Mortgage interest rate	4.07	1.01	3.93	0.99	4.10	1.02
Marginal tax rate (federal + state)	11.70	14.17	11.74	14.10	11.69	14.18
Marginal tax rate (mortgage interest)	11.21	12.49	10.79	12.42	11.30	12.51
Loan-month obs.	17,166,690		3,227,051		13,939,639	
Panel B: All loans (Jan 2022 - Nov 2024)						
	All		Curtailers		Non-Curtailers	
	Mean	SD	Mean	SD	Mean	SD
Origination balance	254,261	141,381	248,120	136,828	255,907	142,532
Current balance	220,539	142,181	205,486	137,893	224,572	143,039
Term (in months)	314	77	301	83	318	75
Months to legal maturity	266	91	251	94	271	90
Effective months to maturity	257	96	228	101	265	94
Mortgage interest rate	3.62	1.08	3.49	0.99	3.66	1.10
Marginal tax rate (federal + state)	11.99	13.65	12.01	13.62	11.98	13.65
Marginal tax rate (mortgage interest)	8.62	12.19	8.39	12.06	8.68	12.23
Loan-month obs.	4,621,147		976,476		3,644,671	
Panel C: Negative rate spread loans (Jan 2022 - Nov 2024)						
	All		Curtailers		Non-Curtailers	
	Mean	SD	Mean	SD	Mean	SD
Origination balance	278,811	145,307	268,048	139,589	281,903	146,762
Current balance	241,020	145,572	222,232	139,814	246,416	146,740
Term (in months)	301	83	287	88	305	82
Months to legal maturity	258	94	242	96	263	92
Effective months to maturity	251	98	224	101	258	95
Mortgage interest rate	3.09	0.60	3.04	0.59	3.10	0.60
Marginal tax rate (federal + state)	7.91	12.85	8.11	12.84	7.85	12.85
Marginal tax rate (mortgage interest)	11.89	13.27	11.33	13.09	12.06	13.31
Loan-month obs.	2,391,713		533,679		1,858,034	

Note: Authors' calculations based Fannie Mae Single-Family Loan Performance, 2010-2019 Surveys of Consumer Finances, and NBER Taxsim. Each cell is a loan-month observation. Marginal tax rate with respect to taxpayer earnings are set to 0 for households with an available retirement account.

Table 1.2: NPV losses under different assumptions

	Share with spread<0	Total NPV (in millions \$)	Average NPV (in \$)
Benchmark results	.52	-4,632.35	-86.80
No available retirement account	.39	-2,217.65	-53.60
No tax adjustments	.59	-6,263.63	-102.25
MBS	.69	-8,411.53	-119.55

Note: Column (1) reports the share of loans with a negative interest rate spread for a given set of assumptions about taxes and opportunity costs. Column (2) reports the aggregate NPV losses in millions of dollars for these loans. Column (3) reports the average NPV loss per payment for these loans, conditional on paying in excess of the monthly minimum. The first row is the benchmark results. In the second row, we assume households do not have a tax-deferred retirement account available. Row three shows results when we do no tax adjustments. Fourth row uses post-tax mortgage interest rates but mortgage-backed securities as the discount rate. Estimates from Jan 2022 - Nov 2024. Authors' calculations based on the Fannie Mae Single-Family Loan Performance dataset, FRED, 2010-2019 Surveys of Consumer Finances and NBER Taxsim. See Section 1.4.2 for more details.

Table 1.3: Prepayment heterogeneity

	$\mathbb{P}(\text{Spread} < 0)$	$\mathbb{P}(\text{Curtail} \text{Spread} < 0)$	Notional value	NPV
All	0.52	0.22	\$555.98	-\$86.80
Subgroup				
High Credit Score	0.56	0.24	\$617.70	-\$95.59
Low Credit Score	0.47	0.20	\$464.53	-\$73.78
First-time Home-buyer	0.42	0.19	\$581.59	-\$103.26
Not First-time Home-buyer	0.54	0.23	\$552.36	-\$84.48
Urban	0.52	0.22	\$565.54	-\$88.77
Rural	0.47	0.23	\$457.12	-\$66.43
Originated 2020-2021	0.70	0.22	\$519.26	-\$92.57

Note: Column (1) reports the likelihood of having a negative interest rate spread. Column (2) reports the probability of prepaying conditional on having a negative interest rate spread. Column (3) reports the average notional value of prepayment for negative rate spread loans. Column (4) represents the NPV of these payments. Each row reports estimates for a specific subgroup. "High" and "low" credit score loans are those with above- and below- median credit scores (773), respectively. Urban is defined by residing in an MSA. Estimates from Jan 2022 - Nov 2024. Authors' calculations based on the Fannie Mae Single-Family Loan Performance dataset, FRED, 2010-2019 Surveys of Consumer Finances and NBER Taxsim.

Table 1.4: Predicting prepayment from credit behavior

Panel A: Predicting mortgage curtailment

	(1)	(2)	(3)	(4)
Credit score			0.0477*** (0.000468)	0.114*** (0.000776)
Total mortgage balance (thousands)			0.00308*** (0.000162)	0.00406*** (0.000264)
Has a credit card			-5.187*** (0.122)	-8.665*** (0.146)
Total credit card balance (thousands)			-0.0785*** (0.00350)	-0.0971*** (0.00468)
Observations	1,158,679	1,158,679	1,158,679	1,158,679
Individual FEs	✓	✓		✓
Quarter FEs		✓		
Mean of outcome	22.15	22.15	22.15	22.15
R^2	0.255	0.376	0.0107	0.272

Panel B: Predicting auto loan curtailment

	$\mathbb{1}(\text{Auto_Curtail})$	Average auto loan term	Longest auto loan term
$\mathbb{1}(\text{Mort_Curtail})$	8.521*** (0.200)	- -	- -
<i>Mort_Curtail</i>	- -	-2.337*** (0.231)	-3.539*** (0.247)
Observations	890,917	102,449	102,449
Borrower controls	✓	✓	✓
Quarter FEs	✓		
Mort. orig. quarter FEs		✓	✓
Mean of outcome	16.97	66.06	68.89

Data and inference. Authors' calculations using the University of California Consumer Credit Panel (UC-CCP). Standard errors in parentheses. For binary outcomes scaled to 0/100, coefficients are percentage-point effects. In Panel B, Column (1) reports borrower-clustered SEs; Columns (2)-(3) report heteroskedasticity-robust SEs. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Panel A (mortgage prepayment). Linear probability models at the household-quarter level; the dependent variable equals 100 if the household prepaid its mortgage in that quarter. Columns (1)-(2) include no covariates (Column (2) adds quarter fixed effects). Columns (3)-(4) add credit covariates.

Panel B (auto loans). Linear regressions. Column (1): outcome is an indicator equal to 100 if any auto loan is prepaid (household-quarter level) with borrower controls and quarter fixed effects. Columns (2)-(3): outcomes are the household's average and longest auto-loan terms (months), estimated at the mortgage-origination (loan) level; controls are fixed at mortgage origination with origination-quarter fixed effects. Borrower controls are age, age squared, credit score, gender, and state fixed effects.

Table 1.5: Effects of interest rate spread on the probability of prepaying

	All			High Credit Score			Low Credit Score		
Spread	-0.01*** (0.00)	-0.00*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.00*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	0.00*** (0.00)	-0.01*** (0.00)
Loan FEs		✓			✓			✓	
Month FEs			✓			✓			✓
Mean of outcome	0.19	0.19	0.19	0.21	0.21	0.21	0.17	0.17	0.17

Note: Table reports estimates from the linear probability model in Equation (1.5). “High” and “low” credit score loans are those with above- and below- median credit scores (773), respectively. Interest rate spread is calculated using post-tax mortgage interest rate and post-tax maturity-matched Treasury yields. $N = 17,166,690$ in Columns (1) and (3). $N = 17,164,564$ in Column (2). $N = 8,571,442$ in Columns (4) and (6). $N = 8,570,528$ in Column (5). $N = 8,595,248$ in Columns (7) and (9). $N = 8,549,036$ in Column (8). Authors’ calculations based on the Fannie Mae Single-Family Loan Performance dataset, FRED, 2010-2019 Surveys of Consumer Finances and NBER Taxsim. Estimates are from Jan 2011 - Nov 2024. Standard errors (in parentheses) are clustered at the loan-level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.6: ARMs vs FRMs

	(1) Mortgage Rate	(2) $\mathbb{P}(\text{Curtail})$	(3) 20-100	(4) 100-500	(5) ≥ 500
$\mathbb{1}(\text{ARM}) \times \text{LIBOR}$	0.644*** (0.001)	-0.019*** (0.002)	-0.009*** (0.001)	-0.009*** (0.001)	-0.001 (0.001)
Observations	12,439,344	12,439,344	12,439,344	12,439,344	12,439,344
Loan FEs	✓	✓	✓	✓	✓
Month FEs	✓	✓	✓	✓	✓
Mean of outcome	4.26	0.18	0.07	0.08	0.04
R^2	0.97	0.52	0.45	0.46	0.38

Note: LIBOR is the twelve-month LIBOR rate. $\mathbb{1}(\text{ARM})$ is an indicator equal to 1 if the loan is an adjustable-rate mortgage and 0 if it is a fixed-rate mortgage. The dependent variable in Column 1 is the loan's interest rate. The dependent variable in Column 2 is an indicator equal to 1 if the loan prepays in that month. Columns 3-5 are indicators equal to 1 if the prepayment amount is in (20, 100), [100, 500), or ≥ 500 dollars, respectively. Standard errors (in parentheses) are clustered at the loan-level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.7: Effects of interest rate variation: Flexible rate vs “locked-in” loans

	(1) Mortgage Rate	(2) P(Curtail)	(3) 20-100	(4) 100-500	(5) ≥500
1 (Flexible rate)	-2.064*** (0.015)	0.058*** (0.004)	0.019*** (0.003)	0.031*** (0.003)	0.007*** (0.002)
LIBOR × 1 (Flexible rate)	0.759*** (0.003)	-0.031*** (0.002)	-0.015*** (0.001)	-0.012*** (0.001)	-0.003** (0.001)
Observations	3,079,361	3,079,361	3,079,361	3,079,361	3,079,361
Loan FEs	✓	✓	✓	✓	✓
Month FEs	✓	✓	✓	✓	✓
Mean of outcome	3.51	0.18	0.06	0.07	0.06
R ²	0.84	0.57	0.42	0.48	0.50

Note: LIBOR is the twelve-month LIBOR rate. 1 (Flexible rate) is an indicator equal to 1 if the loan is not in its introductory rate period. The dependent variable in Column 1 is the loan’s interest rate. The dependent variable in Column 2 is an indicator equal to 1 if the loan prepays in that month. Columns 3-5 are indicators equal to 1 if the prepayment amount is in (20, 100), [100, 500), or ≥ 500 dollars, respectively. Standard errors (in parentheses) are clustered at the loan level.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Chapter 2

How do Economic Transfers affect Health? Evidence from the Alaska Permanent Fund Dividend

Note: The views expressed here are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

2.1 Introduction

How sensitive are health behaviors to large, repeat, unconditional cash transfers? While, there is evidence of a persistently, positive relationship between income and health (Silver and Zhang, 2026), examinations of the impact of cash transfers on health also document heightened risky behavior around the time of cash transfers, such as Supplemental Security Income/SSI (Dobkin and Puller (2007), Finkelstein and Notowidigdo (2025)) or one-time, economic stimulus payments (Gross and Tobacman (2014)). However, among the latter studies, many focus on individuals who are either disabled or lower-income and participate in

means-tested programs, making it difficult to generalize results to a broader socioeconomic group. Thus, understanding what, if any, health behaviors are engendered by large cash transfers, particularly among a broader socioeconomic population, remains an important question. For example, do “full wallets” give rise to risky health behaviors and/or do health-improving behaviors increase because cash transfers enable individuals to receive necessary health care that they may not have otherwise received?

We study these questions by examining receipt of the Alaska Permanent Fund Dividend (**APFD**) payments. Specifically, we analyze how receipt of this repeat, large, unconditional cash transfer affects the timing and consumption of health care utilization. For context, the APFD was received by most Alaska residents (approximately 86%) over the last ten years, and the average *per-person* disbursement was \$1500 per year over this period.¹

To examine the impact on risky health behaviors, we first focus on drug and alcohol-related visits. We then examine if receipt of the APFD also results in increased utilization of more necessary health care that has flexibility in its scheduling, specifically hip and knee replacements, cataracts surgery, and tonsillectomies/adenoidectomy, with cataract surgery being the most commonly performed surgery in Ambulatory Surgery Centers (ASC) in 2019, and tonsillectomy/adenoidectomy being the most common ASC surgery for children, ages 0–17 years, in 2019 (McDermott and Liang (2023)). Specifically, we test if receipt of the APFD relaxes liquidity constraints that inhibit consumption of these types of health care.

We use event-studies to estimate changes in health care utilization in inpatient, emergency department (ED) and outpatient settings for Alaska residents between 2016 and 2023, using data from the Alaska Health Analytics and Vital Records.

As our data allow us to observe the *admission date* for health care utilization, we examine utilization around both 1) the *announcement* of the APFD payment date, which typically oc-

¹According to the 2018-2022 American Community Survey, the median household size in Alaska was approximately three people, implying the median household received, on average, approximately \$4,500 per year.

curs two to three weeks in advance of the payout and 2) the *actual* APFD payment dates.² By separately analyzing the response around these two distinct dates, we are able to test the importance of liquidity for risky behaviors (i.e., the presence of “dangerous liquidity”), in contrast to how the *expectation* of future money affects health behavior.

We find an 15.7% increase in alcohol-related emergency department visits within one week of the APFD disbursement (and 10.7% in the subsequent week), relative to the week prior to the disbursement, providing evidence that cash transfers can yield “full wallets” for those with alcohol use disorders. Thus, while the Permanent Income Hypothesis suggests that individuals should not respond to transitory, anticipated income shocks, we do not find this to be the case here. However, we do not find any changes in admission rates for amphetamine-, cocaine-, or opioid-related emergency department visits. Additionally, we do not find any effects on admissions around the APFD announcement date. Additionally, we do not find that receipt of the APFD affects utilization of cataract surgeries, hip/knee replacements, nor tonsillectomies/adenoidectomies – procedures we consider necessary but may be elective with respect to the timing of the procedure. As such, we fail to find evidence of liquidity constraints in healthcare utilization for this set of medical procedures.

As alcohol ED admissions are the largest category of admissions we examine – in any given week, the weekly alcohol ED volume is 6 to 7 times that of amphetamine and opioid admissions, and more than 40 times large as the weekly volume for cocaine admission – we examine the characteristics of those admitted to the ED for alcohol around the APFD disbursement. We fail to find significant heterogeneous treatment effects; there are no compositional changes in admission rates across gender, age, economic, or geographic subpopulations. Additionally, among those with an alcohol-ED admission the week of the APFD disbursement, we find that approximately 50 percent of individuals admitted had an ED admission for alcohol within the prior 90 days. This suggests that for a non-negligible share of the population admitted for

²We examine the likely date of direct deposit for the APFD.

alcohol, the APFD may affect the intensive margin of alcohol admissions, by contributing to the *number* of alcohol admissions that occur, but it may not affect the extensive margin — that is whether an alcohol ED admission event *ever* occurs.

To understand the importance of liquidity for the increase in alcohol-ED admissions, versus the effect of social gathering that might occur as a result of receipt of a large cash payment, we also examine alcohol-ED admissions around select holidays, specifically July 4, December 25, and December 31. This is because holidays may represent times of heightened social activity, resulting in an increased likelihood in the occurrence of risky behaviors (Blumenberg, Adam (2024)). However, holidays do not necessarily coincide with large increases in cash on-hand, as the APFD disbursement does. We do not find increased alcohol-ED admissions around these holidays, implying that liquidity appears to be a relevant factor for the observed increase in alcohol-ED admissions.

This paper contributes to three literatures. Firstly, we contribute to the literature studying the impacts of cash transfers on health by examining the impacts of a large, repeated and *near universal* cash transfer on health care utilization. Thus, our findings may be more likely to generalize to a much, broader population than those typically studied. In contrast, studies of cash transfers have typically focused on individuals participating in means-tested programs, such as SSI or Supplemental Nutrition Assistance Program (SNAP) benefits, children (Akee et al. (2013), Duncan et al. (2025)), or focus on one-time cash transfers. For example, our closest antecedent Dobkin and Puller (2007) show that drug-related hospitalizations increase sharply around the first day of each month when many welfare programs disburse funds to recipients. However, these increases are largely driven by SSI recipients. Finkelstein and Notowidigdo (2025) similarly show that drug- and alcohol-related emergency department visits spike following SSI receipt but find no such effect for in-kind transfers, namely SNAP benefits. Gross and Tobacman (2014) study how receipt of the 2008 Economic Stimulus payment affects health care utilization; such payments increase the probability of any emergency department

visit in the short- to medium-term by approximately one percent, or an additional 50,000 ED admissions in California. Like our findings, they show that this increase is driven by alcohol-related visits (and drug-related admissions). However, the stimulus payments examined are both smaller in size compared to the APFD, by a factor of two to three, and non-recurring; the APFD, in contrast, is much larger and dispersed annually. As such, we can learn how larger, recurring payments affect health behaviors. This is particularly important if individuals' responses to cash payments are non-linear in the size of the cash transfer (Andreolli and Surico (2026)). Lastly, Evans and Moore (2011) study a broad array of income shocks, including transfers, paychecks, and annual dividends, and their effects on all-cause mortality. They document an approximately nine percent increase in mortality one week after the APFD disbursement, and find that eight percent of Alaska deaths, during their sample period, are attributable to substance abuse. We expand on this by examining a broader set of health outcomes beyond mortality, and highlight the health risk of alcohol overconsumption around the time of APFD receipt.

Secondly, we contribute to literature studying the overall economic effects of the APFD. For example, a number of studies, including Watson et al. (2020) and Jones and Marinescu (2022), focus on the impacts of the APFD on crime and labor supply, respectively. Watson et al. (2020) finds an increase in police incidents related to substance-abuse after APFD disbursement, while Jones and Marinescu (2022) find no effects of the APFD on overall employment, but find increases in part-time employment. Additionally, several studies, such as Hsieh (2003) and Kueng (2018), focus on the impact of the APFD on consumer spending, and estimate marginal propensities to consume (MPC) out of the APFD. Kueng (2018) finds increased consumption of non-durables within one quarter of APFD disbursement, but no anticipation effects occurring before the disbursement.³ Related to health care, Guettabi and Witman (2023) show that receipt of the dividend has no effect on prescription drug consumption for privately insured Alaskans. Thus, we contribute to this growing literature by focusing on the effect of this cash

³A review of the overall effects of the dividend can be found Gibson et al. (2020).

transfer on risky health behaviors and by also examining if receipt of the APFD relaxes liquidity constraints that may inhibit healthcare consumption. As such, this is the first known study, to the authors, to examine the impact of the APFD on risky health behaviors and the utilization of health care, specifically health care received in an inpatient or outpatient setting.

Lastly, we contribute to growing discussions on the effects of Basic Income (BI) programs. BI programs have increased in prevalence in the US, with more than 160 pilots having taken place in more recent years.(Stanford Basic Income Lab (2026)). However, extrapolations from the findings of these programs may be limited, in part, because of smaller enrollments (often fewer than 1000 participants), income-eligibility requirements for participants, and also because these programs often operate for a few years or fewer. Thus, we contribute by examining a large-scale, long-running program that affects most residents of Alaska, and discuss the impacts of this large, near-universal, cash transfer on liquidity-sensitive health outcomes. By doing so, we show that certain risky behaviors are sensitive to the timing of cash transfers, namely the overconsumption of alcohol.

The rest of our paper will proceed as follows: Section 2 discusses the Alaska Permanent Fund background, Section 3 presents our data while Section 4 discusses the methodology. We present our results in Section 5, discussion in Section 6, conclude in Section 7.

2.2 Alaska Permanent Fund Background

The Alaska Permanent Fund Corporation is a state-owned enterprise that manages a number of funds including its namesake, the Alaska Permanent Fund (APF). The APF generates revenues primarily through oil and mining, and total fund value as of September 2025 exceeds \$84 billion dollars (Alaska Permanent Fund Corporation, 2025). Voters established the APF via an amendment to the Alaska State Constitution in 1976, and in 1982, the state began paying residents the Alaska Permanent Fund Dividend (APFD) annually. The dividend is an annual,

lump sum payment given to all eligible residents that apply and the (nominal) average amount between 2016 and 2023 was approximately \$1,500 (State of Alaska Department of Revenue, n.d.).

Eligible residents must apply in the first quarter of the year to receive the dividend. Eligibility requires individuals to meet a number of conditions. Firstly, one must be a resident of Alaska for the *entire* preceding calendar year. For example, to be eligible for the 2025 dividend, one must have been a legal resident of Alaska from *at least* January 1, 2024 - December 31, 2024. Of note, children born on or before December 31 are eligible for the dividend in the following calendar year. Secondly, one must both intend to remain an Alaska resident indefinitely and not claim residency in any other jurisdiction in the previous calendar year. Those who were sentenced due to or were incarcerated for a felony in the previous calendar year are ineligible.⁴ Lastly, those absent from the state for more than 180 days must be on an “allowable absence,” and residents must physically reside in Alaska for at “at least 72 consecutive hours at sometime during” either of the two preceding calendar years (State of Alaska, Department of Revenue, 2024). While we cannot observe individuals’ eligibility status in our data, the dividend take-up rate is high; for example, we estimate that nearly 85% of the entire population were paid a dividend in 2023.⁵

The dividend amount for each year is calculated using the number of eligible applicants in that year and the total amount of the Permanent Fund that is available for distribution, as determined by the AK state legislature.⁶ The AK state legislature takes multiple factors into account when determining the amount available, including the financial performance of the total Fund itself. Available funds are also reduced by prior-year dividend obligations, operation expenses, and other state agency program appropriations. For example, the Alaskan legislature

⁴Those incarcerated in for a misdemeanor in the prior calendar year may also be ineligible if they have been previously convicted of a felony or multiple misdemeanors.

⁵The remaining 15% includes both ineligible persons and non-applicants.

⁶Prior to 2018, the dividend calculation was made using a statutory formula. Dividend amounts from 2018 to 2024 have not used this formula.

has used a portion of the Permanent Fund earnings over the past ten years to cover the cost of government services such as the Alaska Department of Health and Social Services and the Department of Corrections.

Table 2.1 provides information on the APFD disbursements for our sample. Each year, the dividend amount and payout date is announced anywhere from 2 to 4 weeks prior to the payout date. Payout dates typically occur in early October or late September save 2020. The overwhelming majority of applicants receive the dividend via direct deposit; this allows us to precisely estimate the effects of the dividend timing on health behaviors since nearly all recipients receive their disbursement on the same day. Lastly, the average dividend amount is about \$1,500 but can vary anywhere from less than \$1,000 to nearly \$3,300 per person per year.

2.3 Data

We use admission-level data provided by the Alaska Health Analytics and Vital Records.⁷ Our sample includes the universe of inpatient, emergency department (ED), and other outpatient admissions to Alaska health care facilities occurring between 2016-2023. The data include medical information such as primary and secondary diagnoses and procedures; demographic information such as patient age, gender, race and ethnicity, and zip code of residence; and insurance information such as primary insurer and hospital charges. Lastly, we observe a time-invariant, person-specific identifier that allows us to identify if a person has had multiple visits to a health care facility in our data.

We categorize drug and alcohol admissions based on ICD-10 diagnoses codes and identify cataract, hip and knee, and tonsillectomy/adenoidectomy surgical procedures using or Current Procedural Terminology (CPT) codes. For drug and alcohol-related admissions, we follow

⁷Data is provided by the Alaska Department of Health, Division of Public Health, Health Analytics and Vital Records Section.

Dobkin and Puller (2007), and classify an admission type as being related to drug or alcohol if either the principal diagnosis code or any of the 29 secondary diagnoses codes are recorded as the relevant ICD-10 code. We use the following ICD-10 codes to categorize admissions belonging to the admission category alcohol (F10 and T510X), amphetamines (F15, T43.62, and T43.65), cocaine (F14 and T405X), and opioids (F11, T40.0X, T40.1X, T40.2X, T40.3X, T40.41, T40.42 and T40.49).

Dates of the APFD direct deposit and check mailings are collected from the *Permanent Fund Dividend Division Annual Report*, which is provided by the State of Alaska Department of Revenue and is published each fiscal year. We determine the initial announcement date of the APFD, which provides the public with the [future] date of APFD payout and the amount to be paid out, by determining either 1) the date listed on The State of Alaska Department of Revenue, Permanent Fund Dividend Announcement, which states, *For Immediate Release* or 2) the date listed on an online Alaska news articles stating that the dividend amount (or payout date) was announced *today* (e.g., “Today, the governor announced how much Alaska citizens will be getting” or “Today, the Department of Revenue announced the amount of the Alaska Permanent Fund Dividend”).⁸ Lastly, we define the date of the APFD payout, our benchmark for the disbursement date, as the date direct deposits are scheduled to be paid out. This is because, during our sample period, 88 percent of recipients, on average, receive their disbursement through this method. The remainder receive their disbursement via check, which is typically mailed out within two weeks of the direct deposit date.

2.3.1 Summary Statistics

We present summary statistics of patient data in our sample in Table 2.2 and benchmark them to both the Alaska and US population as a whole using American Community Survey (ACS) 5-year estimates from 2018-2022, the approximate midpoint of our sample period. Panel A shows

⁸These two dates are confirmed to coincide.

the demographic breakdown. Our sample of patients is slightly older than both the Alaska and US average. Women constitute a disproportionate share of our sample (55-60%), especially relative to the Alaska population, which has more males (52.6%). The American Indian/Alaska Native (AI/AN) population is overrepresented in the hospitalization data; while just under 14% of the Alaska population identifies as AI/AN, about one third of all hospitalizations are from the AI/AN population. Rural patients also constitute a disproportionate share of admissions; about half of our sample is rural patients, whereas only 15.2% of the Alaska population lives in a rural area. For additional context, we include median family size and median household income for Alaskans and all Americans; both populations have a median family size of 3 persons, and Alaskans have higher household incomes (\$100,309 vs \$88,612).

In Panel B, we calculate the primary payer share of admissions – government, commercial, self-pay, and other, where these categories are mutually exclusive. In Panel C, we calculate the share of people covered by government and commercial insurance and those without any coverage. These categories are not mutually exclusive as some people are dually-covered through marriage, Medicare/Medicaid etc. Public insurers are the primary payer for between half and over two thirds of admissions. However, only about 35% of Alaskans and Americans more broadly are covered by a public insurer. That public insurance constitutes such a large proportion of the primary payer in our data is due in part to the fact that those aged 65 and older (i.e., those who are likely covered by Medicare) are overrepresented in the data; when we remove admissions from patients aged 65 and over, the share of inpatient/ED/outpatient admissions whose primary payer is a government insurer falls to 60%, 59% and 42%, respectively. Additionally, that AI/AN admissions are overrepresented in our sample relative to the overall population explains this discrepancy; nearly half of all AI/AN children and nearly one quarter of all AI/AN adults are enrolled in Medicaid (Raimondi, 2025).

2.4 Methodology

To investigate the effects of the dividend on health behaviors, we estimate the following event study:

$$y_t = \sum_{k=-30; k \neq -1}^{30} \beta_k \mathbf{1}(t - T_0 = k) + \varepsilon_t \quad (2.1)$$

where y_t is admission counts per 10,000 people in day t and $\mathbf{1}(t - T_0 = k)$ is an indicator equal to 1 if day t is k days before/after the direct deposit date, T_0 . We include only days that fall within the ± 30 day window around the direct deposit date each year. Each β_k represents the average daily change in admission rates relative to the day before dividend disbursement.⁹

We rely on exogeneity in the annual timing of the APFD disbursement to identify the effect of dividend receipt on health behaviors (Guettabi and Witman (2023)). That is, we exploit the exogenous variation in the exact day the dividend is released each year, where the dividend disbursement varies in calendar date as well as day of the week. The spirit of our strategy is to make comparisons between the days just after the APFD disbursement is made, in contrast to the day just prior to the dividend disbursement, under the assumption that health behaviors in the days just before the disbursement are a good counterfactual for health behaviors just after the disbursement. This presumes that within a local window of the disbursement, liquidity-sensitive health behaviors tend to trend smoothly in the absence of the cash disbursement. While this assumption cannot be tested directly, we provide evidence in support of our assumption by 1) examining and confirming smooth trends in utilization for health events that are not likely to be liquidity-sensitive or events whose timing is difficult to manipulate, and thus unlikely to be affected by the APFD (see discussion in Section 2.5.3) and 2) showing that in the two years the APFD was paid out in a non-typical disbursement (i.e., July 2020 and September 2022 versus the typical October disbursement), health responses, if any, occur around the respective payout dates in these atypical payout years. This allows us to rule out

⁹We cluster standard errors at the calendar week level to account for serial correlation in the dependent variable.

an October-specific effect as the explanation for changes in health behaviors.

2.5 Results

2.5.1 Drug and Alcohol Admissions

Descriptive figures for drug and alcohol utilization in the aggregate are shown in Figure 2.2, where we present weekly totals for alcohol-, amphetamine-, cocaine-, and opioid-related visits.¹⁰ Panel (a) shows ED visits and panel (b) shows inpatient visits. There is a pronounced spike in alcohol-related ED admissions in the first two weeks after the dividend is paid out; total admissions increase from 2,345 admissions in the week prior to 2,842 – an approximate 21% increase. However, we find no visual evidence of changes in ED admissions for other drugs and no evidence of changes in inpatient admissions for *any* drug categories.

We next test for dangerous liquidity by estimating Equation (3.2) for admissions related to alcohol, amphetamine, cocaine, and opiates. We exclusively focus on emergency department (ED) encounters because ED admissions more appropriately isolate acute, high-risk behaviors that would not be observed in scheduled or elective care settings. Results are shown in Figure 2.3. Panels (a), (b), (c), and (d) show the effect of the APFD on admission rates per 10,000 people for alcohol-, amphetamine-, cocaine-, and opioid-related admissions, respectively. We find an immediate and pronounced short-term increase in alcohol-related ED visits. The average increase in daily admissions across the first seven days ($\frac{1}{7} \sum_{k=0}^6 \widehat{\beta}_k$) is approximately 0.09 per 10,000 people; the average alcohol ED admission rate for $t_{-1} = 0.6$, implying a 15.7% ($p = .001$) relative average increase in the week of the dividend and a 10.7% ($p = .018$) increase in the subsequent week. The effects return to statistically indistinguishable from zero after just two weeks. However, we find no statistically significant change in admission rates for amphetamines, cocaine, or opioids; the average change in admission rates in the first week

¹⁰Note that in graphs where data points are based on fewer than 20 events, the rates plotted may be statistically unreliable.

for each admission type are 0.005, -0.008, and 0.004 admissions per 10,000, respectively.

The previous results tested *if* excessive drug and alcohol consumption are sensitive to large, lump-sum payment. We next examine *how* sensitive these behaviors are to variation in the size of the APFD payment. Column 6 of Table 2.1 shows that there is substantial variation in the dollar amount of the dividend in our sample, ranging from less than \$1,000 (2020) to over \$3,000 (2022). To test for sensitivity along the intensive margin, we estimate:

$$y_t = \sum_{k=-4, k \neq -1}^{k=4} \beta_k Div_{kt} + \varepsilon_t \quad (2.2)$$

where Div_{kt} is an indicator equal to 1 if the day t is in the k th week before/after the disbursement of the dividend.¹¹ Week $k = 0$ begins on the direct deposit date and includes the next six days. We estimate this equation for each year 2016-2023 and plot each estimated coefficient of β_0 – the average weekly change in daily alcohol admissions for the week of the dividend – against the respective dividend amount. Results are shown in Figure 2.4. While most of the yearly coefficients are statistically distinguishable from zero, we do not see a clear monotonic relationship between the dividend amount and alcohol admissions.

2.5.2 Heterogeneous effects on alcohol-related admissions

Given the sharp increase in alcohol- related admissions around the APFD and its large admission volume, we explore if there are heterogeneous effects of the dividend on alcohol consumption across different subpopulations. This is motivated by a body of literature documenting how risky alcohol consumption varies across different socioeconomic factors (Chartier and Caetano (2010) and Patrick et al. (2012)), gender (Yakovlev (2018) and Nolen-Hoeksema (2004)), and across age (Bohm et al., 2021), among other dimensions. We estimate Equation (3.2) using the share of alcohol admissions that come from the specific subgroup of interest as the dependent

¹¹We use this approach rather than estimating Equation (3.2) for each year because the latter approach would force each coefficient to be estimated off precisely one data point.

variable. That is,

$$y_t = \frac{\text{\#admits from subpopulation}_t}{\text{\#total admissions}_t}. \quad (2.3)$$

This approach allows – holding overall alcohol admissions fixed – to assess whether the *composition* of alcohol-related ED admissions has shifted across subgroups. The estimated coefficients are interpreted as the average daily change in the share of total alcohol ED admissions coming from a particular subpopulation. We also graph daily admission levels by subgroup in Appendix Figures B.1 - B.4.

Gender Historically, men have engaged in binge drinking at much higher rates than women. However, in recent years, the gap has narrowed, driven in large part by the binge drinking rate of women in their 30s and 40s doubling between 2006-2018 (McKetta and Keyes, 2020). Given these two facts, it is *ex ante* unclear whether spikes in alcohol admissions should be differentially driven by men or women. Results are shown in Figure 2.6, where we use the male share of alcohol ED admissions as the dependent variable. While overall, males constitute the majority of these types of admissions in our sample (approximately 55% vs. 45%), we do not find a compositional shift in admissions around the time of APFD disbursement; all coefficients are statistically indistinguishable from zero. Alternatively put, men are not *differentially* affected by the dividend with respect to alcohol consumption relative to women.

Age Excessive alcohol consumption is also more highly concentrated among young people, in particular those aged 18-34 (Keyes and Miech, 2013). To test for heterogeneous effects across age, we divide our sample into six different age categories, allowing for variation in alcohol consumption likelihood across the age groups – those younger than 18, 18-34, 35-44, 45-54, 65-74, and those aged 75 years and older. Results are shown in Figure 2.7. We do not detect compositional changes in age for alcohol related ED admissions. Despite prior literature suggesting that binge and heavy episodic drinking is more popular among younger people, we find that APFD-induced alcohol-related admissions are more or less proportional

across age groups. Further, our results are robustly null to alternative age groupings (e.g., dis-aggregating the 18-34 bins into 18-24 and 25-34).

Geography The literature provides mixed evidence on the geographic determination in excessive alcohol consumption risk; conclusions are sensitive to definitions of urbanicity and measures of alcohol consumption. However, the evidence leans more toward peoples from rural communities being at higher risk of alcohol-related harm than those from urban communities (Friesen et al., 2022). We test if those from urban or rural communities are differentially affected by the APFD in our setting. We adopt the approach of Evans and Moore (2011) and define an admission to be “urban” if the patient zip code is from a city with more than 10,000 residents (Anchorage, Fairbanks, and Juneau). Results are shown in Figure 2.8, where the outcome is the share of alcohol admissions from urban zip codes. We cannot conclude that the geographic composition of alcohol-related admissions is affected by the APFD disbursement.

Income We next examine heterogeneity by income. Firstly, economic theory predicts that responses to income shocks vary depending on baseline resources; households with lower incomes face tighter budget constraints and may therefore exhibit larger consumption responses to income shocks. Secondly, the literature finds evidence of heterogeneous alcohol consumption and harm by socioeconomic status; for example, Collins (2016) finds that individuals from higher socioeconomic status consume “similar or greater amounts of alcohol” compared to those with lower socioeconomic status but the latter group faces higher rates of alcohol-related harm. Using data from the American Community Survey, we categorize admissions into quintiles of income based on patient zip code. That is, we take each patient’s zip code’s median household income and map it to the statewide distribution of household incomes. We then assign each admission to an income quintile, where the first quintile corresponds to zip codes in the bottom 20% of Alaska’s income distribution, and the top quintile corresponds to those in the top 20%. Results are shown in Figure 2.9. We find a precise null effect. While the majority (between 80-85%) of alcohol ED admissions come from lower-income zip codes, we

find virtually no effect of the APFD disbursement on the composition of such hospitalizations.

Insurance status Lastly, we examine heterogeneity by patient insurer – more specifically, the primary payer. Insurance coverage shapes both the cost of medical care and the underlying health and socioeconomic composition of the patient population. As such, one may expect differential health utilization and outcomes across different insurer types; for example, privately insured patients facing higher out-of-pocket (OOP) costs may be more sensitive to the receipt of a large cash transfer than a patient with a low OOP max. To test for heterogeneous responses, we categorize primary payer into four groups: commercial, government, self-pay, and other. Results are shown in Figure 2.10. We fail to find evidence of a differential effect of the APFD on alcohol ED admissions by insurer type.

2.5.3 Placebo admissions and robustness checks

As our setting does not provide a clear untreated comparison group, since the dividend is nearly universally received, we rely on placebo outcomes that should not plausibly be affected by the dividend to assess the validity of our research design. We focus on appendicitis, asthma, and birth admissions where the first two events are often acute with limited notice as to their occurrence and where there should be little incentive to time births to the APFD disbursement since children are eligible for the current years' fund payment so long as they are born by December 31 in the preceding calendar year. As such, these placebo outcomes act as a diagnostic test; if the increase in alcohol admissions we see is a spurious result due to factors other than the dividend, we should expect to see similar changes in the placebo outcomes.

We estimate Equation (3.2) where the dependent variable is admissions per 10,000 people for appendicitis, asthma, or birth admissions. Appendicitis and asthma admissions come from inpatient, ED, and outpatient datasets, and birth admissions come solely from the inpatient dataset.¹² Results are shown in Figure 2.5 Panel (a) shows the estimated coefficients and

¹²Nearly every birth admission in our sample is in an inpatient setting.

confidence intervals for appendicitis, and panels (b) and (c) show the results for asthma and births, respectively. We do not detect statistically significant changes in admissions across all three placebos around the timing of the APFD disbursement. Further, the estimated coefficients are small in magnitude and are near zero. The absence of effects is consistent with our interpretation that the observed increases in alcohol admissions are due to receipt of the APFD rather than a spurious correlation.

We further evaluate the credibility of our design by ruling out calendar-specific shocks or seasonal factors. Because the dividend is usually paid out in the first week of October each year during our sample period, we estimate Equation (2.2) using only observations from 2020 and 2022, when the dividend was paid out in July and September, respectively.¹³ Results are shown in Figure 2.11. We find a respective 24.5%, 26.8% and 17.6% increase in alcohol admissions in the week of the dividend and subsequent two weeks. These estimated coefficients are qualitatively similar to our benchmark estimates, suggesting the latter are not confounded by October-specific factors.

2.5.4 Liquidity Constrained Health Behavior

We next test for liquidity constraints in health care consumption. If individuals are liquidity constrained, short-term increases in cash could affect both the timing and likelihood of seeking health care. Conversely, if individuals face no such liquidity constraints (e.g., by borrowing against future income), then temporary income shocks should have little to no effect on utilization. We focus on more necessary forms of health care with flexibility in scheduling; specifically, our outcomes of interest are hip and knee replacements and cataract surgeries, the latter being the most commonly performed surgeries in outpatient settings (McDermott and Liang, 2023). Results are shown in Figure 2.12. Panel (a) shows results for cataract surgery and panel (b) shows results for hip or knee replacement surgery. The event study for cataract surgery yields

¹³We again use this methodology to ensure sufficient sample size.

noisy, statistically insignificant effects. This is in part due to the fact that of a limited sample size – $N = 488$ days within ± 30 day window across 8 years. The results in panel (b) show a tight null effect of the APFD on hip/knee replacement surgeries.

While hip/knee replacement and cataract surgeries are popular procedures, they are primarily concentrated among those aged 65 and older where the overwhelming majority of those aged 65 and older are insured via Medicare, and thus may face relatively lower cost-sharing for health care. Given this, we expand our examination to tonsillectomy and adenoidectomy procedures to 1) focus on a younger demographic and 2) isolate individuals for whom a medical procedure may be more likely to bind their budget constraint, given that individuals receiving this health care are unlikely to have more generous Medicare insurance. Further, these are the most common major ambulatory surgeries performed in hospital-owned facilities in 2019 for those aged under 18 (McDermott and Liang, 2023). Results are shown in panel (c) of Figure 2.12. Again, we find no statistically significant changes in tonsillectomy/adenoidectomy procedures around the timing of the dividend. In sum, we fail to find evidence in support of liquidity constraints in health care utilization for these procedures.

2.5.5 Responses around the Announcement Date

We next investigate the sensitivity of the examined health behaviors to the *announcement* of the APFD. In doing so, we examine the importance of liquidity versus the expectation of future money on health behaviors. To provide context to our examination, the predictions of the Permanent Income Hypothesis (Friedman, 1957) imply that individuals should not respond to transitory, anticipated income shocks, and should only react to *unanticipated* income shocks, where consumption increases at the announcement of such income shocks (assuming that individuals are not liquidity constrained).

To test if the future expectation of money engenders a similarly sized response to the actual receipt of cash, we examine whether there are changes in health care utilization around the

announcement date of the dividend payout (i.e., the earliest date announcing the [future] payout date of the dividend and the amount that will be paid). We reestimate Equation (?? but create relative time dummies around the *announcement date* rather than the direct deposit date.

Results for alcohol and drug-related outcomes are shown in Figure 2.13. We do not detect statistically significant changes in admissions for amphetamines, cocaine, or opioids. We do estimate significant increase sin alcohol admissions three to four weeks after the announcement; however, this is likely due to the fact that the direct deposit date is typically three weeks after the announcement. Therefore, it appears that alcohol consumption is sensitive to the actual receipt of cash, rather than the expectation of future cash.

2.6 Discussion

To give a broader context to our findings, we consider why we do not find effects of the APFD outside of alcohol admissions, while also considering the economic significance of our alcohol estimate.

It may be the case that healthcare spending after receipt of the Permanent Fund dividend is not concentrated in major procedures, such as knee replacements or cataract surgeries, but rather in more routine care. Studies by the JPMorgan Chase Institute (JPMCI) provide strong suggestive evidence that this may be the case. Using linked bank account and credit card data for 1.2 million JPMorgan Chase Bank customers from 2014 to 2016, Farrell et al. (2018) show that around the time of tax refund repayments, consumers increase out-of-pocket (OOP) healthcare spending by 60 percent within the first week of receiving their tax refund (relative to a typical week prior to receiving the tax refund). 32 percent of all OOP payments made from the tax refund went to dentists' offices while 23 percent went to doctors' offices. Results were also greater for individuals in the lowest quintile of checking balances: for these individuals,

the increase in OOP healthcare spending was 220 percent greater than in a typical week prior to the refund. For high checking account balance individuals (in the top quintile of balances), the post-refund increase was only 11 percent. These results suggests that liquidity constraints do play a major role in the timing and likelihood of receiving care. When combined with our results, it is also likely the case that liquidity constraints also matter for the *type* of care received when receiving a cash transfer.

In addition to increases in OOP healthcare spending, Farrell et al. (2018) also show that individuals receiving tax refunds also spent a large percentage of their increased liquidity to pay outstanding medical bills. 37 percent of the additional OOP spending triggered by the tax refund goes to remote bill pay (i.e., payments not made in-person), with 31 percent of that 37 percent going to hospitals. This suggests that the Permanent Fund dividend may not increase the likelihood of new hospital procedures around the timing of dividend receipt, but may instead lead to bill payments for prior hospitalizations or other prior healthcare procedures.

As discussed, our estimates imply that receipt of the APFD increases alcohol ED admissions nearly 16% in the week of the dividend, followed by nearly 11% the following week. These relative magnitudes are comparable to Finkelstein and Notowidigdo (2025) and Dobkin and Puller (2007). The former estimate a 30% increase in drug/alcohol admissions among SSI recipients following the receipt of the transfer, and the latter estimate a 23% increase among SSI recipients. While these relative effects are large, they are benchmarked to underlying hospitalization rates that are low; for example, the baseline average daily alcohol ED admission rate per 10,000 people in Alaska in the week preceding the dividend is .57 admissions per day per 10k. For Finkelstein and Notowidigdo (2025), the baseline drug/alcohol admission rate is 2.36 per day per 10k; this is substantially higher than the rate in our setting, but our denominator includes the entire population of Alaska, whereas their denominator only includes South Carolina SSI recipients. Because the underlying hospitalization rates are low, this implies that the APFD begets only modest increases in individual risk. However, our results are additionally

consistent with Finkelstein and Notowidigdo (2025) in that our results are largely driven by a smaller but substantially higher-risk subpopulation. For those with an alcohol-ED admission within one week of the APFD, around 50% of these individuals had an ED admission for alcohol within the prior 90 days. Likewise, they estimate that 60% of the increases in drug/alcohol ED admissions due to the receipt of SSI is driven by those with prior behavioral health issues.

2.7 Conclusion

We examine the effects of receipt of the Alaska Permanent Fund Dividend (APFD) on health behaviors and health care utilization. We examine if the timing of certain risky health behaviors, namely drug and alcohol consumption that results in a hospital visit, is sensitive to the timing of the APFD disbursement. We also examine if the funds' disbursement relaxes liquidity constraints for health care consumption, such that there is an increase in the receipt of health care that has flexibility in its timing, where we focus on hip and knee surgeries, cataract surgery, and tonsillectomy/adenoidectomies.

We find evidence that receipt of the APFD results in "dangerous liquidity" for certain risky behaviors, specifically for alcohol consumption, with an approximately 15.7 percent increase in the number of alcohol-related emergency department admissions within one week of APFD disbursement. Additionally, we do not find evidence of changes in the consumption patterns of the above mentioned health procedures. Lastly, we find that the timing and volume of alcohol emergency department admissions is only sensitive to the timing of APFD payout, and not to the announcement of the future payout date nor future payout amount. This suggests that liquidity, rather than the future expectation of money, is an important contributor for certain risky behaviors.

Figures and Tables

Figure 2.1: Total weekly hospitalizations relative to the dividend






Figure 2.2: Total weekly drug-related hospitalizations relative to the dividend




(a) Emergency Department




(b) Inpatient


Figure 2.3: Effects of APFD on drug-related emergency department visits




(a) Alcohol



(b) Amphetamines



(c) Cocaine



(d) Opioids

Figure 2.4: Effects of APFD admissions by dividend amount







Figure 2.5: Effects of APFD on placebo admissions



(a) Appendicitis



(b) Asthma



(c) Births

Figure 2.6: Heterogeneous effects of APFD on alcohol-related admissions: Male share of admissions







Figure 2.7: Heterogeneous effects of APFD on alcohol-related admissions: Age




(a) < 18




(b) 18-34




(c) 35-44



(d) 45-54



(e) 65-74



(f) 75+

Figure 2.8: Heterogeneous effects of APFD on alcohol-related admissions: Urban share of admissions







Figure 2.9: Heterogeneous effects of APFD on alcohol-related admissions: Income




(a) Bottom quintile




(b) Second quintile



(c) Third quintile




(d) Fourth quintile




(e) Top quintile


Figure 2.10: Heterogeneous effects of APFD on alcohol-related admissions: Insurance type




(a) Commercial



(b) Government



(c) Self-pay



(d) Other

Figure 2.11: Effects of the APFD on Alcohol ED admissions: July 2020 and September 2022 only







Figure 2.12: Effects of APFD on elective surgery procedures



(a) Cataract Surgery




(b) Hip/Knee Replacement




(c) Tonsillectomy/adenoidectomy


Figure 2.13: Effects of APFD announcement on drug-related emergency department visits




(a) Alcohol



(b) Amphetamines



(c) Cocaine



(d) Opioids

Table 2.1: APFD Summary, 2016-2023

Year	Announcement Date	Direct Deposit Date	Check Date	% Direct Deposit	Dividend \$
2016	Sep 14	Oct 06	Oct 06	84.17%	\$1,022
2017	Sep 15	Oct 05	Oct 05	84.17%	\$1,100
2018	Sep 04	Oct 04	Oct 04	85.22%	\$1,600
2019	Sep 27	Oct 03	Oct 24	87.13%	\$1,606
2020	Jun 12	Jul 01	Jul 01	88.58%	\$992
2021	Sep 30	Oct 11	Oct 25	89.93%	\$1,114
2022	Sep 08	Sep 20	Oct 06	90.94%	\$3,284
2023	Sep 21	Oct 05	Oct 26	90.66%	\$1,702

This table reports the announcement date, the direct deposit date, the mailed check date, the share of recipients who received direct deposit dividends, and the dividend amount each year. Dollar values are nominal. Information provided by State of Alaska Department of Revenue PFD Annual Reports (<https://pfd.alaska.gov/Division-Info/annual-reports>)

Table 2.2: Summary Statistics

Variable	Sample			ACS	
	Inpatient	ED	Outpatient	Alaska	U.S.
Panel A: Demographics					
Age (years)	43.0	38.2	46.2	36.5	39.1
Male (%)	44.5	46.2	40.0	52.6	49.6
White (%)	52.3	41.7	43.5	58.3	58.9
AIAN (%)	27.6	36.9	40.5	13.8	0.6
Black (%)	3.2	4.3	1.7	3.2	12.1
Asian (%)	3.7	2.6	2.4	7.6	5.8
Other (%)	7.3	6.9	8.6	9.6	4.0
Hispanic (%)	3.3	4.2	2.6	7.5	18.6
Urban (%)	52.0	54.4	49.8	84.8	92.1
Family size (median)	-	-	-	3	3
HH income (median)	-	-	-	100,309	88,612
Panel B: Primary Payer					
Government (%)	70.0	62.9	53.5	-	-
Commercial (%)	24.7	23.6	33.9	-	-
Self-pay (%)	5.4	9.9	10.7	-	-
Other (%)	0.9	3.5	1.8	-	-
Panel C: Health Insurance Coverage					
Government (%)	-	-	-	35.3	36.2
Commercial (%)	-	-	-	66.3	66.9
No insurance (%)	-	-	-	11.3	9.0
Observations	477,123	2,228,227	8,480,077	33,065	15,721,123

Note: Each cell reports variable's mean value unless otherwise specified.. The first three columns are from the 2016-2023 Alaska Health Analytics and Vital Records hospitalization data. The fourth and fifth columns are from the 2018-2022 5-year American Community Survey. Race categories exclude Hispanics: respondents of Hispanic ethnicity are coded as Hispanic irrespective of race, while non-Hispanics are classified by the reported race. Urban is defined as being from a city of 10,000 or more in the hospitalization data and as being from a metro area using ACS data.

Chapter 3

Dope and Disorder: The effects of the reformulation of OxyContin on crime and homelessness

3.1 Introduction

Drug use is frequently believed to impose large social costs. Two of the most salient of these costs are crime and homelessness, and policymakers will often justify expensive interventions on the basis that such drug use destabilizes communities. For example, in 2024 alone, the United States federal government spent nearly 44 billion dollars on drug control measures, half of which was spent on prevention, prohibition, and domestic law enforcement (Statista, 2024). Additionally, in 2022, the Drug Enforcement Agency (DEA) announced its Operation Overdrive initiative, which was aimed at curbing increasing rates of drug-related violent crime. The stated purpose was “to make our communities safer and healthier, and to reverse the devastating trends of drug-related violence” (Drug Enforcement Administration , DEA). President Trump’s 2025 Executive Order *Ending Crime and Disorder on America’s Streets* redirected federal grant

money toward civil commitment of homeless persons, “the overwhelming majority of [whom] are addicted to drugs” (The White House, 2025). Given such substantial resources spent on prevention, it is important to understand if and how drug use can generate these externalities.

This paper studies if and how drugs – specifically illegal street opioids – contribute to both criminal activity and the incidence of homelessness. To do so, I use a supply shock to the stock of prescription opioids, which induced abusers to substitute heroin and later fentanyl. In 2010, Purdue Pharma reformulated its name-brand oxycodone product, OxyContin, and introduced a new version that was designed to be resistant to abuse or misuse. This reformulation made it difficult for individuals to abuse the pill to snort or inject it, for example. Prior literature (Alpert et al. (2018); Evans et al. (2019); Powell and Pacula (2021)) has documented that this reformation not only exacerbated the opioid crisis, but begot its “heroin-wave,” leading to massive spikes in use and heroin-related mortality. I use a difference-in-differences design, making comparisons of areas that were “highly exposed” to OxyContin to areas that were less exposed to OxyContin to estimate the causal effect of illicit opioid use on crime/homelessness rates.

This massive substitution toward more potent, dangerous, and harmful alternatives (Collier, 2013) could increase crime via substance-financing activity (i.e., users commit income generating crime to sustain their habit); it could also increase crime via pharmacological effects, whereby users commit crimes while intoxicated due to an altered state of mind. However, I find that illicit opioid use has essentially no effect on any of the most serious offenses. This result is both surprising as it goes against conventional wisdom and important in that it suggests crime-control rationales for some anti-opioid interventions may be overstated.

On the other hand, opioid use may cause users to experience homelessness. The correlation between drug use and homelessness has been well documented in the medical, public health, and sociological literature (Doran et al. (2018); Johnson and Chamberlain (2008)). Among adults seeking specialty treatment for opioid use disorder, about one in eight reported experi-

encing homelessness (Han et al., 2022). Motivated by this relationship, I study the effects of the opioid epidemic on homelessness. I find that areas with the largest pre-reformulation OxyContin prescription rates saw the largest post-reformulation increases in homelessness rates. A one standard deviation increase in exposure to OxyContin is associated with an 8% increase in the overall homeless population.

This paper contributes to the emerging literature studying the negative externalities of the opioid epidemic beyond health and mortality. Some authors identify the effects of opioid use – via variation in state prescription limiting laws – on outcomes such as earnings and marriage prospects (Kaestner and Ziedan, 2023), job loss and labor force exit (Harris et al., 2020), and home prices (Custodio et al., 2024). (Kaestner and Ziedan, 2023) document that adoption of modern Prescription Drug Monitoring Programs (PDMPs) reduced employment and earnings by about 2%. Likewise, Harris et al. (2020) exploit state variation in PDMP adoption and document that a 10% increase in per capita opioid prescription rates are associated with a 0.53 point reduction in labor force participation. Custodio et al. (2024) similarly exploit state variation in opioid-limiting laws and show that *decreases* in opioid abuse from such laws are correlated with increases in county-level home prices. Another set of papers have used the reformulation of OxyContin as a setting to study the epidemic’s downstream effects, but use state-level variation in reported OxyContin misuse from the National Survey on Drug Use and Health (NSDUH) to study outcomes such as disability receipt (Park and Powell, 2021). Lastly, (Karimli, 2022) uses the same data, setting, and methodology as this paper to study how the opioid epidemic increased mortgage default; I improve on this measure by using a more precise definition of exposure (see 3.3.1 for more details).

It also contributes to the vast literature studying the relationship between drug use and criminal activity. In the context of opioids, Moore and Schnepel (2024) document that a massive reduction in the heroin supply in Australia caused a short-term 41% increase in violent crimes, largely driven by homicides and assaults. Further, they find income-generating crimes such

as robbery and burglary increase immediately following the supply shock. This finding supports the substance-financing hypothesis; the reduction in supply caused a 400% increase in the purity-adjusted price of heroin, thereby incentivizing users to engage in crime to support their use. Tan (2024) also studies the effects of the reformulation on homicides and finds that states with higher pre-reformulation OxyContin misuse rates saw a modest increase in homicides post-reformulation, but this effect quickly dissipates. I also study homicide rates but use policing data to measure homicide rather than mortality data as Tan (2024) does.

Other papers study how state-laws targeting opioid use, both supply and demand laws, affect crime. Deiana and Giua (2021) show that supply-side policies targeted at curbing prescription misuse such as pain management clinic laws (PMCL) increase illicit opioid sale related crimes. Mallatt (2022) also shows that supply-side policies such as and prescription drug monitoring programs (PDMPs), coupled with the reformulation of OxyContin, induced substitution toward heroin as evidenced by increases in heroin possession and sale offenses. Sim (2023) finds that states with tighter prescription laws – namely, triplicate prescription laws – in place at the introduction of OxyContin have 25% higher crime rates on average than those without such laws. Being able to better regulate the dispensing of OxyContin, states with triplicate prescription laws saw less consumption of OxyContin, and therefore less crime. Where my paper differs is I focus explicitly on the second and third waves of the epidemic (heroin- and synthetics-waves). Given the illicit nature of heroin and synthetic opioids such as fentanyl, the policy implications with respect to the opioid use-crime relationship are different. For example, policy aimed at curbing illicit opioid use cannot operate through standard regulatory channels such as prescribing rules or physician monitoring programs.

In terms of the *demand* for opioids, Doleac and Mukherjee (2022) find that increased naloxone access is associated with increases in opioid-related crimes and, albeit imprecisely estimated, overall increases in theft. With the adoption of naloxone access laws, the risk of fatal overdose is smaller, which could potentially encourage use and ultimately increase use-sustaining crime.

Outside the context of opioids, Fryer et al. (2013) show that the crack epidemic caused 5% increases in violent and property crimes in the largest US cities in the 1980s. Again, not all drugs are the same, so it is important to characterize the relationship between drug use and crime for different classes of drugs.

Lastly, despite the aforementioned well-documented relationship between drug use and homelessness, the literature largely lacks a causal link. To my knowledge, there is only one paper that uses rigorous causal methods studying the effects of the opioid epidemic on the incidence of homelessness. Like the setting of this paper, Olvera et al. (2024) also study how policies designed to curb prescription opioid abuse may have real economic effects downstream. Specifically, they exploit the up-scheduling of hydrocodone in 2014, which made prescribing and dispensing the drug more stringent. Using a similar design as this paper, they find that areas that were mostly highly exposed to hydrocodone prior to the up-scheduling saw the largest *reductions* in homelessness rates. In other words, by making the potentially addictive drug harder to obtain, they argue this decreased abuse and ultimately decreased the likelihood a hydrocodone user would eventually become homeless. In sum, given the positive relationship between opioid use and homelessness, the up-scheduling of hydrocodone had positive externalities. I substantiate their results using a different setting and supply shock, also finding a positive causal association between opioid use and homelessness.

This paper's main contributions are twofold. First, this paper studies *illicit* opioid abuse, rather than just overall drug abuse. The Centers for Disease Control and Prevention break the opioid epidemic into three distinct waves. The first wave, from the 1990s to 2010, were a result of prescription opioid use. The second ("heroin wave") and third ("synthetic opioid wave") waves began in 2010 and 2013, respectively. The latter two are the waves of interest. Very little of the previous literature has this narrowed focus. Focusing on illicit use is important because the relationship between crime and drug use may differ along different classes of drugs.¹ Illicit

¹Consider, for example, the propensity to commit serious crimes for a recreational marijuana user vs. a crystal methamphetamine user.

opioids are sourced through black markets that may generate other criminal activity; as a result, our understanding of the drug-crime link may not generalize to illicit drug use. The levers policymakers have to curb prescription opioid abuse and illicit drug use are entirely different. Secondly, much of the extant literature makes cross-state comparisons; this paper improves by using a finer geography (county). Given the fact that there is significant within-state variation in heroin/fentanyl use and criminal activity, state-level aggregates can obscure meaningful spatial heterogeneity. A county-level analysis allows us to study how opioid use and its externalities co-vary across counties within state rather than averaging over often disparate local areas.

3.2 Background

The opioid epidemic refers to the increasing rate of use, abuse, and fatal overdoses of prescription or illicit opioids since the late 1990s to the present day. Between 1999 and 2023, nearly 1 million Americans have died of a fatal opioid overdose (Centers for Disease Control and Prevention, 2025). Opioids are analgesics – i.e., pain killers – whose pain-relieving properties are derived from or mimic the opioid poppy plant. This class of drugs became particularly popular in the late 1990s as doctors began commonly prescribing opioids to manage chronic pain (Congressional Research Service, 2022). Prior to this, opioids were mostly used for acute pain from cancer or terminal illnesses. Around the same time, pharmaceutical companies heavily promoted these drugs; these promotions often asserted that they were non-addictive or that the risk of addiction was small (Van Zee, 2009).

The consequences of this rapid spike in prescriptions are still felt today. Almost 650,000 Americans died from an opioid-related overdose between 1999-2021. In 2021 alone, three out of every four drug overdose deaths included an opioid (Centers for Disease Control and Prevention, 2025). As of 2021, almost 3 million Americans struggle with opioid use disorder, with nearly 80% of those going untreated (Jones et al., 2023). The epidemic is also one of the pre-

eminent causes of the rising mortality rates for low-education non-Hispanic whites in the US (Case and Deaton, 2015).

3.2.1 OxyContin and its reformulation

Purdue Pharma first developed OxyContin in 1996. It is Purdue's trade name product for oxycodone hydrochloride, one of many common generic class of opioids prescribed for pain-relief. It quickly became not only one of the most popular opioids, but the fifteenth most popular prescribed drug overall in the US (Bartholow, 2011). OxyContin's relative potency and duration is one of the reasons for its popularity among users and abusers² alike: its original formulation was available in tablets of dosage strength ranging from 10 milligrams to 80 milligrams³. Comparatively, Percocet typically contains only 5 milligrams of oxycodone. In addition to its relative dosage strength, OxyContin also had the longest duration (12 hours) of all oxycodone products at the time of introduction (National Drug Intelligence Center, 2001). Another – and a perhaps more important – reason that OxyContin was popular among abusers was that if the tablet were dissolved or crushed, the user can experience an immediate and euphoric high. In essence, the user can experience the 12-hour slow-release effect all at once (Alpert et al., 2018).

However, Purdue wanted to strip its association with “this abuse profile” and produce an alternative, abuse-deterrent version of the drug (Cicero and Ellis, 2015). Fourteen years after its introduction, the company announced it reformulated OxyContin and by August 2010, it ceased all shipments of the drug. Purdue executives and the Food and Drug Administration celebrated the formulation as a step in the right direction in curbing the opioid epidemic.⁴ The stock of one the most popular drugs to abuse disappeared practically overnight. This meant that people who had already developed an addiction to the drug had two options to get high:

²Abuse of a prescription opioid should be thought of as any use of the drug beyond the prescribed instructions.

³Just four years later, Purdue released a 160 mg version.

⁴Another key fact about the reformulation is that there was no notification made to the public nor to prescribers, and prices of OxyContin did not change (Beachler et al., 2022).

increase consumption but sacrifice the immediacy they received via tampering with the tablet pre-reformulation, or substitute their use to other opioids. Self-reported OxyContin misuse plummeted post-reformulation; a third of OxyContin users replaced it with other drugs, and almost three quarters of these switched to heroin (Cicero and Ellis, 2015). All else equal, there was an immediate increase in use of street opioids, and given their addictive properties, a consequent increase in the demand for *more* street opioids.

3.2.2 Conceptual framework

Opioid use and crime Since Becker (1968), economists have modeled criminal activity as a decision an agent makes by weighing the potential rewards and punishments he receives for committing a crime. Other researchers (e.g., Maclean et al. (2022)) have extended this model in the context of opioid use. I adopt this framework and posit two channels that I test. Firstly, an agent that uses illegal opioids may commit a crime to fund his or her addiction. Users quickly develop a tolerance to the drug, needing higher doses at higher frequencies to maintain a high. With an increase in the user's quantity demand for the drug with a fixed budget constraint, the user might commit crimes for external financing. On the other hand, opioid use – and drug use more generally- alters users' discretion. While intoxicated, the expected costs and benefits of criminal activity can change; opioids have been shown to alter users' intertemporal preferences, discounting future rewards (or punishments) far higher than non-users (Kirby et al. (1999); Madden et al. (1997)).

If heroin causes users to commit crimes that fund their addictions, we should expect to see increases in *acquisitive* or income-generating crimes. Borrowing definitions from Freedman and Owens (2016), acquisitive crimes include robbery, burglary, larceny, and motor-vehicle theft.⁵ If, on the other hand, opioid use causes users to commit crimes simply due to its pharmacological effects, we should expect to see increases in *non-acquisitive* crimes – murder, rape, and

⁵Freedman and Owens (2016) also include intrafamily assault in this definition, but I do not observe such crimes in the data.

assault; or, at the very least, we should expect there to be no differential changes in crime rates between acquisitive and non-acquisitive crime (Goldstein, 1985).

Of course, these drug-crime channels are not necessarily unique to street opioids nor prescription opioids more broadly. But what makes the reformulation of OxyContin an attractive setting is that heroin is one of the most addictive substances on the planet. An estimated 23-38% of first-time heroin users develop a physical dependence to the drug, whereas only 3-19% of people who take prescription painkillers are estimated to develop an addiction (AMA Alliance (2019); Santiago Rivera et al. (2018)). So if the number of users of heroin increases, we should also see the number of *chronic* users of heroin increase, which could increase the number of people at risk to engage in criminal behavior.

Opioid use and homelessness The relationship between opioid use and homelessness can propagate through a number of channels. Firstly, opioid use can reduce one's labor supply due to job loss (Harris et al., 2020) or health deterioration, for example. Lack of stable work can cause a user to miss rent payments, which may result in eviction. Users may also be locked out of housing altogether; a prior drug arrest or conviction may cause landlords to reject users' rental applications (Evans and Porter, 2015). This sort of market friction may cause users on the margin to lose housing access altogether. Opioid use may also cause interpersonal relationship issues among the user's family and friends, which can reduce her informal safety net; the loss of such a safety net may cause a user to transition into homelessness (Jurewicz et al., 2022).

3.3 Data and methodology

Opioid data Administrative data on opioid prescriptions come from the Drug Enforcement Agency's (DEA) Automated Reports and Consolidated Ordering System (ARCOS). This dataset contains the universe of prescription opioid transactions between manufacturers/distributors

to buyers in the US from 2006-2019. My analysis sample only includes transactions where the buyer is a pharmacy (retail, chain, etc.) or a practitioner/nurse practitioner/physician's assistant office. I follow the literature (e.g., Cornaggia et al. (2022)) and subset on these types of buyers to isolate *consumer* use. Opioids shipped to research universities or nursing homes are unlikely end up in the hands of prescription drug abusers, so such buyers are not included in my sample. Theft from these institutions into the hands of drug abusers is possible, but I assume that these sorts of leakages are negligible. I aggregate total OxyContin volume to the county-year level in morphine milligram equivalents (MME).

Crime and arrest data Crime data come the FBI's Uniform Crime Reporting Database (UCR). I use two files: (1) Offense Known and Clearances by Arrest ("Offenses Known") and (2) the Persons Arrested datasets. The former is an agency-month panel that contains counts of crimes either reported to or known to the police; i.e., the data include both events where there was an arrestee and those that were not cleared. However, the data observed in the file are only the most serious crimes (i.e., "index" or Part I crimes). This list includes homicide, rape, robbery, aggravated assault, burglary, auto theft, and arson. This means that drug possession or sale offenses, for example, are not reported in the dataset.

The Persons Arrested dataset is similar, in that it is an agency-year panel that contains the number of *arrests* by crime type. One benefit of this dataset is that it includes arrest counts for less serious crimes (Part II crimes) including the possession and sale of certain drugs. That I only observe drug possession and sale *arrests*, a subset of criminal activity that is actually detected by law enforcement, has implications for the interpretation of the results. Estimates using drug arrests should be interpreted as capturing changes in opioid-related criminal justice contact, rather than the change in total incidence of drug crime.

There are a few issues with UCR. Firstly, not all police agencies report to the FBI. There is no statute mandating agencies report – it is only on a voluntary basis. These leads to a potential

undercounting of crimes and a possible sample selection issue. Fortunately, most agencies *do* report (Federal Bureau of Investigation, 2010). Another potential problematic issue with the datasets is that crime reports are subject to the Hierarchy Rule. That is, in a multi-crime event, only the most serious crime is reported. For example, if a criminal burglarizes a home and murders the inhabitants, this crime is coded *only* as a homicide. Specifically, the general hierarchy is: homicide > rape > robbery > aggravated assault > burglary > theft > auto theft > arson > simple assault. Because of the crime type hierarchy, the data will mechanically undercount less serious offenses. While criminologists often point to the hierarchy rule as an inherent flaw in the Offenses Known data, Kaplan (2021) points out the degree of undercounting is almost negligible.

3.3.1 Methodology

To capture the measure of “exposure” to OxyContin, for each CoC, I construct the following variable:

$$OxyExp_c = \frac{\sum_{t=2006}^{2009} OxyContin_{c,t}}{\sum_{t=2006}^{2009} Population_{c,t}} \quad (3.1)$$

where the numerator, $OxyContin_{c,t}$ is the total sum of all OxyContin prescriptions in morphine milligram equivalents. For easier interpretation of coefficients, I normalize $OxyExp$ such that $E[OxyExp] = 0$ and $V[OxyExp] = 1$; i.e., I subtract the sample mean and divide this difference by the sample standard deviation. This variable is similar to the measure in Karimli (2022) – who studies the effects of the reformulation on mortgage default – but I specifically use OxyContin rather than just any oxycodone. Also, I use morphine milligram equivalents rather than dosage units. This paper’s measure of exposure is an improvement in two ways. Firstly, by only using OxyContin prescriptions rather than any oxycodone, this measure better targets areas that were more exposed to OxyContin. Consider the extreme example where a given area consumes more oxycodone per capita than any other area but *none* is OxyContin. This area would *not* be exposed to OxyContin, but the pre-existing measure of exposure would

be high for this area. Secondly, using morphine milligram equivalent rather than dosage units is an improvement in terms of precision. For example, 10mg of one opioid can be less potent than 20mg of another opioid if the mme of the former is greater than the latter.

An obvious concern is that OxyContin prescription rates are correlated with local economic conditions, and therefore correlated with criminal activity. However, the spatial variation in prescription rates is largely an artifact of Purdue’s marketing campaigns of OxyContin in the 1990s. Further, there is substantial spatial variation in physician prescribing (Cutler et al. (2019); Beheshti and Neller (2023)). Importantly, this marketing and physician discretion – and thus an area’s exposure to OxyContin – has been shown to be uncorrelated with economic conditions (Currie and Schwandt, 2021).

3.3.2 Estimating equation

To identify the causal effects of illicit opioid use on crime, I estimate the following two-way fixed effects model:

$$y_{ct} = \beta OxyExp_c \times \mathbf{1}(Year > 2010)_t + X'_{ct}\Omega + \alpha_c + \delta_t + \nu_{ct} \quad (3.2)$$

where y_{ct} is either crime rate (crime type per 100,000 people) or homelessness rate in county c in year t . $\mathbf{1}(Year > 2010)$ is an indicator equal to one when $t > 2010$, the year of the reformulation. Thus, β identifies the effect of a one standard deviation increase in pre-reformulation OxyContin’s intensity on post-reformulation outcomes. It is interpreted causally under the assumption that absent for reformulation, outcomes in counties with ex-ante OxyContin exposure would have followed parallel trends conditional on observables. X is a vector of demographic, economic, and policing characteristics⁶. To account for cross-sectional differences across coun-

⁶X includes unemployment rate; labor force participation; average home price; shares of population that is disabled, black, married, and with a HS degree or less; shares of population with high school degree, some college, and bachelors degree or more. It also includes policing characteristics such as police officers per capita and the female share of the police force.

ties and common shocks across time, I include county and year fixed effects – α_c and δ_t , respectively. To account for serial correlation within county, standard errors are clustered at the county level.

To understand the potential dynamic treatment effects, I also estimate event studies using a modified specification:

$$y_{ct} = \sum_{t=2007; t \neq 2009}^{t=2019} \beta_t OxyExp_c \times \mathbf{1}(Year = t) + X'_{ct} \Omega + \alpha_c + \delta_t + \nu_{ct} \quad (3.3)$$

In each model, I also include time-varying indicators for the adoption of state laws targeting prescription opioid abuse.⁷ These include both supply-side laws (Prescription Drug Monitoring Programs, Pain Management Clinic Laws, Doctor Shopping Laws) and demand-side laws (Naloxone Access and Good-Samaritan Laws). I interact these with OxyContin intensity to account for differential treatment effects across different policy interventions.

3.4 Results

I first substantiate the literature and verify that exposure to OxyContin did indeed induce users to substitute to heroin and fentanyl. Figure 3.1 shows state average age adjusted death rates by quartiles of OxyContin exposure. These are calculating using deaths where the underlying cause of death is some form of drug use, and the multiple cause of death code includes heroin, synthetic opioids, or methadone.⁸ Prior to the reformulation, there is virtually no difference in death rates regardless of OxyContin's penetration. Once the abuse-deterrent formulation is introduced, there are clear divergent trends; states with more (less) OxyContin prescriptions see larger increases in illicit opioid ADRs. The highest quartile of exposed states have an

⁷I use Deiana and Giua (2021) to code the timing of adoption for each law for each state.

⁸Specifically, I query deaths where the IDC-10 underlying cause of death codes are X40–X44, X60–X64, X85, or Y10–Y14 (accidental poisonings by or exposure to, intentional self-harm from, or assaults by non-alcoholic drugs), and multiple cause of death codes T40.1 (heroin) and/or T40.4 (synthetic opioids excluding methadone).

average ADR that is almost 5x greater than that of the lowest quartile by 2019.

As mentioned, the main independent variable of interest is *OxyExp*, which has been normalized to be mean 0 with standard deviation 1. Because of this normalization, the coefficients in all the following tables and figures should be interpreted as a one standard deviation increase in exposure to OxyContin in 2007-2009. To put these coefficients in context, the standard deviation of *OxyExp* is about 50 mme/person in my sample. Oxycodone's mme conversion factor is 1.5, meaning 1 milligram of oxycodone is equivalent to 1.5 milligrams of morphine. A typical OxyContin tablet has a dosage strength starting at 10mg. Thus, it may be apt to think of a one standard deviation increase in exposure as giving 3.3 10mg OxyContin tablets to every single person in each county in the pre-reformulation period.

How does pre-reformulation exposure translate into actual illicit opioid use? I attempt to do so by estimating Equation (3.2) at the state level, where the dependent variable is deaths via heroin/synthetics/methadone per 100k individuals. The coefficient on the $OxyExp \times 1(Year > 2010)$ interaction term is 2.53, meaning a one standard deviation exposure increase in OxyContin prior to reformulation caused an average of 2.53 illicit opioid deaths per 100k. So if we assume the extreme and morbid possibility a one-to-one mapping of use to death, then this is the absolute lower bound on increases in heroin use. In sum, a one standard deviation in exposure to OxyContin could be thought of as increasing the rate of heroin users in a county by *at least* 2.53 per 100k.

3.4.1 Crime Results

Drug possession and sale Figure 3.2 shows results from (3.3) using cocaine/opium related crimes per 100,000 individuals. UCR data lumps together cocaine and opium-related drugs (e.g., heroin, morphine and codeine). While an imperfect measure of illicit opioids, it is the best available in the data. Panel (a) shows the results for possession arrests, while Panel (b) shows the results for sale arrests. I find large increases in the possession arrests and precisely no

effect on the sale of these drugs. Surprisingly, there is a lagged effect on per capita possession arrest; there is precisely 0 impact immediately following the reformulation, but by 2017, I estimate an increase of about 11 possession arrests per capita.

Acquisitive crimes Table 3.1 reports estimates from (3.2) using acquisitive (or income-generating) crime rates as the dependent variable. Panel A reports results using offenses known (i.e., any reported crime known to police), whereas Panel B reports results using arrests. Column (1) in both panels aggregates all crime types together, and Columns (2)-(5) reports each component crime. Surprisingly, every single point estimate is *negative*. However, with the exception of motor vehicle theft, all results are indistinguishable from zero. I include the sample mean of each outcome from 2006, the first year in my dataset. Relative to the respective means, the point estimates are extremely small. For example, I estimate that a one standard deviation increase in exposure led to a .03% reduction in all acquisitive crimes per 100k reported relative to the benchmark average. With 95% confidence, I can rule out an effect size above (below) a .02% increase (.08% decrease). Compare these results to Fryer et al. (2013), who estimate that increases in crack cocaine use in the 1980s are associated with about a 5% increase in violent and property crimes; these effect sizes are an order of magnitude larger than those that I find. Overall, these results do not support the “substance-financing” hypothesis. In other words, I find no evidence to suggest that increases in illicit opioid use induce users to commit crimes to fund their addictions.

Non-acquisitive crimes Table 3.2 reports estimates from Equation (3.2), but the dependent variable are non-acquisitive crimes rates. Again, Panel A reports results using known offenses, whereas Panel B reports results using arrests. Column (1) uses an aggregate crime rate, and Columns (2)-(4) report results for each component crime. By and large, I again find virtually no effect of illicit opioid use on any type of crime rate; each point estimate is small, and all are indistinguishable from zero. Again, even taking the point estimates seriously, a one standard deviation increase in OxyContin exposure *at most* increases non-acquisitive crime rates

by .07%. An interesting pattern is that all point estimates for known offenses are positive, whereas all point estimates for arrests are negative; however, I am unable to determine exactly why such a pattern arises. These results also do not support the “pharmacological effects” hypothesis; increases in heroin use do not cause increases in serious but otherwise non-income generating crimes.

Dynamic effects To test for the plausibility of parallel trends and to understand the dynamic effects of the reformulation over time, I estimate Equation (3.3) for aggregated acquisitive and non-acquisitive offense rates. Results are shown in Figure 3.3. The blue dots represent estimates of each β_t using acquisitive offenses known per 100k as the dependent variable, and the red dots represent non-acquisitive offenses known per 100k. All pre-reformulation point estimates of β_t are statistically null, supporting the parallel trends assumption. Further, for acquisitive offenses, I am unable to detect any statistically significant effect; the point estimates bounce around ± 0 and are noisily estimated. For non-acquisitive offenses, the results show mostly no effect until about 5 years after the reformulation; there is a modest up-tick in crime but this quickly returns to zero.

Alternative measure of exposure As a robustness check, I calculate an alternative measure of OxyContin exposure. Specifically, I use the OxyContin *share* of all oxycodone; i.e., I calculate:

$$OxyShare_c = \frac{\sum_{i=2006}^{2009} OxyContin_{ct}}{\sum_{i=2006}^{2009} Oxycodone_{ct}} \quad (3.4)$$

Here, exposure is measured by the proportion of all oxycodone purchases that is specifically OxyContin. The logic is similar to that of *OxyExp*; areas where OxyContin is more popular relative to other forms of oxycodone should also see higher increases in heroin/fentanyl use. I normalize this to be mean 0 and variance 1 to keep the interpretation of coefficients consistent.

Event study results for cocaine/opium possession and sales arrest rates are shown in Figure 3.5.

Panel (a) shows possession arrests per 100k, while Panel (b) shows sales arrests per 100k. The results largely substantiate those found using the default measure of exposure. What differs is I detect a slight dip in possession arrests post 2010, but there is an increase in later years. All coefficients are insignificant, however. For sale arrest rates, all point estimates are insignificant, but the pattern shows a mild and steady increase in the number of arrests per capita.

Static difference-in-differences results for known offense rates are shown in Table 3.3. The results using this alternative measure of exposure exhibit more or less the same patterns as those using the default measure; however the coefficients are larger in magnitude but the results are noisier. I am unable to statistically detect any sort of effect of an increase in OxyContin exposure on crime rates.

Placebo test A natural concern about *OxyExp* is that it does not necessarily capture the intensity of OxyContin, but rather the demand for opioids more generally, which could potentially bias the results. This placebo measure should not predict heroin/fentanyl use after the reformulation in the way that OxyContin exposure does. I provide suggestive evidence that this is not the case. I calculate (3.1) but instead using all non-oxycodone opioid prescriptions mme units in the numerator. I then plot state average ADR by quartile of the placebo. Results are shown in Figure 3.4. Compared to Figure 3.1, the degree of divergence in trends is particularly muted. Further, the relationship between placebo exposure and mortality is not monotonic as it is in the case with OxyContin exposure. For example, by 2019, the second quartile has the highest ADR rather than the fourth quartile. That the placebo measure does not strongly predict heroin and fentanyl use after the reformulation, but the OxyContin measure does, suggests that *OxyExp* truly does capture the intensity of OxyContin usage and not the demand for opioids more generally.

I estimate Equation (3.3) with this placebo measure of intensity using known offense rates as the dependent variable. Results are shown in Figure 3.6. Again, the blue dots represent

estimates of each β_t using acquisitive offenses known per 100k as the dependent variable, and the red dots represent non-acquisitive offenses known per 100k. The results are noisily estimated but there is no clear pattern in outcomes. As expected, I find no effect on either type of offense rate.

Discussion Overall, I conclude that increases in illegal opioid use have at most no effect on income-generating crime and at least a modest *negative* effect – primarily through reductions in motor vehicle theft. I argue that illegal opioid use also has no to a small positive effect on non-acquisitive crimes. That there is also no effect on arrests for such crimes rules out the possibility that heroin use could cause crime but law enforcement have failed at properly investigating.

It could be possible that many heroin users obtain financing through purely licit means. From 2008-2013, an estimated 72% of Americans with substance use disorder were employed either full or part-time (Substance Abuse and Mental Health Services Administration, SAMHSA). Perhaps income earned through legal employment is enough to sustain a user's addiction. However, this is at odds with previous papers that show *prescription* opioid use leads to depressed earnings, job loss, and labor force exit (Harris et al. (2020); Kaestner and Ziedan (2023)). For non-acquisitive crimes, it's possible that the intoxication from heroin or fentanyl is simply not conducive to committing serious offenses. The high from opioids, including heroin and fentanyl, includes severe drowsiness and cognitive clouding. Heart function and breathing become depressed, and users often experience nausea (National Institute on Drug Abuse, NIDA). It is harder to imagine a person under the influence of heroin committing a rape, for example, than someone under the influence of crack cocaine, which increases alertness and paranoia, and is often associated with aggression and violence (Miller et al., 1991).

One possible explanation is that the mortality from increased heroin use has offset the possible crime effects. The share of overdose deaths attributed to either heroin or synthetics constitutes

the majority of opioid-related deaths post-reformulation (Centers for Disease Control and Prevention, 2025). Unfortunately, I am not able to test this hypothesis with the data available.

3.4.2 Homelessness Results

As mentioned, Olvera et al. (2024) study how the up-scheduling of hydrocodone, in the absence of dangerous substitutes, caused areas most-highly exposed to hydrocodone see the largest reductions in hydrocodone use and ultimately the largest reductions in homelessness rates. I hypothesize that the areas mostly highly exposed to OxyContin prior to its reformulation will see the largest *increases* in homelessness rates. Note that my hypothesis and the results from Olvera et al. (2024) both support a positive connection between opioid use and homelessness. I argue that OxyContin, given its highly addictive qualities and popularity among abusers, induced users to substitute toward illicit drugs whereas no such substitution (or at least to a lesser degree) occurred for hydrocodone users. In simple terms, OxyContin users substituted to heroin – hydrocodone users did not.

To test this hypothesis, I construct a new panel of homelessness rates at Continuum of Care - year level. A Continuum of Care (CoC) is a geographic unit designated by the Department of Housing and Urban Development (HUD). It is often a collection of counties but sometimes is coterminous with a single county. There are over 400 CoCs in the US. Homelessness data come from HUD's Continuum of Care Homeless Assistance Programs Homeless Populations and Subpopulations Reports. The data provide aggregate counts of homeless persons for each CoC and year. The reports come from Point-in-Time (PIT) estimates; specifically, the PIT estimates are "a count of sheltered and unsheltered homeless persons on a single night during the last ten days in January" (U.S. Department of Housing and Urban Development, 2024). CoCs also partition counts into "sheltered" and "unsheltered" homeless individuals. A sheltered homeless individual defined as one living in a supervised shelter run by governmental or charitable entities intended on providing temporary housing; an unsheltered homeless individual is one "with

a primary nighttime residence that is a public or private place not designed for or ordinarily used as a regular sleeping accommodation for human beings, including a car, park, abandoned building, bus or train station, airport, or camping ground” (U.S. Department of Housing and Urban Development, 2016).

One issue with the PIT counts is that while CoCs are mandated to report sheltered counts every year, they are only required to report unsheltered counts every odd year. However, it is still common for CoCs to count the unsheltered homeless population each year. To account for this, I follow Olvera et al. (2024) and include an indicator in models where the dependent variable is unsheltered homeless counts. Additionally, data collection began in 2007 so I lose one pre-reformulation study period.

I estimate (3.3) using the number of homeless individuals per 100k in CoC c in year t as the dependent variable. Results are shown in Figure 3.7. Panel (a) shows results for all homeless individuals, whereas (b) and (c) show results for unsheltered and sheltered, respectively. I find that areas highly exposed to OxyContin also see the largest increases in their homeless populations. The treatment effect estimated in the static difference-in-differences model is about 69; relative to the sample mean, this implies that a one standard deviation increase in OxyContin exposure is associated with an 8% increase in the homeless population. The results are almost exclusively driven by the unsheltered homeless population. For unsheltered homeless, the implied effect size is about a 15% increase. This is consistent with the literature on different homeless populations’ relationships with drug and alcohol abuse: sheltered individuals use drugs far less than unsheltered individuals (Rountree et al. (2019); Kuhn and Culhane (1998)).

To put these effect sizes into perspective, Olvera et al. (2024) find that the up-scheduling of hydrocodone decreased the rate of number of individuals experiencing homelessness by over 25%. They too find this reduction is largely driven by decreases in the unsheltered population; up-scheduling reduced the rate for this population by 32%. My results show that the effects

of opioid use on homelessness are symmetric in sign but not necessarily in magnitude, as I estimate much more “conservative” effect sizes.

In sum, the results both confirm the positive causal direction of opioid use and homelessness, but my findings highlight an important policy implication: laws aimed at curbing prescription/dispensing drugs *can* have positive downstream effects if there are no harmful substitutes.

3.5 Conclusion


This paper studies how the reformulation of OxyContin, which ignited the illicit waves of the opioid epidemic, affected crime. I find that areas most exposed to OxyContin saw the greatest heroin- and fentanyl-related mortality rates post reformulation, and also saw the largest increases in heroin possession arrests. Contrary to conventional wisdom, I find virtually no effect on heroin sales, and more importantly, no effect on serious income-generating and non-income generating crime. However, it is important to note that the data do not capture the full universe of criminal activity, but only offenses known to law enforcement or those resulting in arrest. As such, the null effects should be interpreted as pertaining to justice-contacted crime rather than total criminal behavior. These results rule out the hypotheses that opioid users may commit crimes that result in arrests to sustain their habits or because of the drugs’ intoxicating effects. One negative externality that did result from the reformulation is a major increase in the incidence of homelessness. I estimate that a one standard deviation increase in exposure to OxyContin pre-reformulation is associated with an 8% increase in the homeless population post-reformulation, largely driven by unsheltered homeless individuals.

Taken together, these findings have implications for the allocation of public resources. Policies such as Operation Overdrive and others designed to curb drug-related crime may be ineffective in addressing the broader externalities of opioid use. Instead, the results suggest that homelessness is a more pronounced externality of the opioid epidemic, and that there is potential

value in reallocating resources toward supporting the homeless population.

Figures and Tables


Figure 3.1: Age adjusted death rates by exposure to OxyContin




Note: Figure shows state-level average age-adjusted death rates by state quartile of OxyContin exposure. Exposure is calculated using Equation (3.1) but at the state-level. The vertical dashed line is August 2010, the month when all pre-reformulation OxyContin was removed from pharmacy shelves. Author's calculations based on ARCOS and CDC Multiple Cause of Death datasets.

Figure 3.2: Effects of OxyContin reformulation on drug possession/sale arrests

(a) Cocaine/Heroin Possession




(b) Cocaine/Heroin Sale




Note: Figure shows results from Equation 3.3. The dependent variable in Panel (a) is cocaine/heroin possession arrests per 100k. The dependent variable in Panel (b) is cocaine/heroin sale arrests per 100k. The vertical dashed line is August 2010, the month when all pre-reformulation OxyContin was removed from pharmacy shelves.

Figure 3.3: Dynamic effects of OxyContin reformulation on crime rates



Note: Figure shows results from Equation 3.3. The blue estimates/confidence intervals are results using all acquisitive crimes known per 100k individuals, and the red estimates/confidence intervals are results using all non-acquisitive crimes per 100k individuals. The vertical dashed line is August 2010, the month when all pre-reformulation OxyContin was removed from pharmacy shelves.


Figure 3.4: Age adjusted death rates by placebo exposure




Note: Figure shows state-level average age-adjusted death rates by quartile of placebo exposure. Exposure is calculated using Equation (3.1) using non-oxycodone mme units but at the state-level. The vertical dashed line is August 2010, the month when all pre-reformulation OxyContin was removed from pharmacy shelves. Author's calculations based on ARCOS and CDC Multiple Cause of Death datasets.

Figure 3.5: Drug Possession/Sales Arrest – Alternative Measure

(a) Cocaine/Heroin Possession




(b) Cocaine/Heroin Sale



Note: Figure shows results from Equation 3.3. The dependent variable in Panel (a) is cocaine/heroin possession arrests per 100k. The dependent variable in Panel (b) is cocaine/heroin sale arrests per 100k. The vertical dashed line is August 2010, the month when all pre-reformulation OxyContin was removed from pharmacy shelves.


Figure 3.6: Placebo effects of OxyContin reformulation on crime rates




Note: Figure shows results from Equation 3.3 using a placebo exposure measure. The blue estimates/confidence intervals are results using all acquisitive crimes known per 100k individuals, and the red estimates/confidence intervals are results using all non-acquisitive crimes per 100k individuals. The vertical dashed line is August 2010, the month when all pre-reformulation OxyContin was removed from pharmacy shelves.

Figure 3.7: Effects of illicit opioid use on homelessness rates


(a) All homeless



(b) Unsheltered Homeless



(c) Sheltered Homeless



Note: Figure shows results from Equation 3.3. The dependent variable in Panels (a), (b), and (c) are overall, unsheltered, and sheltered homeless individuals per 100k people, respectively. The vertical dashed line is August 2010, the month when all pre-reformulation OxyContin was removed from pharmacy shelves.

Table 3.1: Acquisitive crimes

Panel A: Offenses Known					
	(1)	(2)	(3)	(4)	(5)
	All acquisitive	Burglary	Larceny	Motor vehicle theft	Robbery
$OxyExp \times Post$	-6.998 (12.32)	-0.656 (3.849)	-1.453 (8.666)	-4.414** (1.353)	-0.475 (0.433)
Observations	40792	40792	40792	40792	40792
County FEs	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓
$\bar{Y}_{2006}$	2193.02	526.61	1464.717	159.54	41.92

Panel B: Arrests					
$OxyExp \times Post$	-2.037 (3.806)	-0.300 (0.933)	-0.658 (2.987)	-0.646 (0.383)	-0.433 (0.324)
Observations	40726	40726	40726	40726	40726
County FEs	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓
$\bar{Y}_{2006}$	390.73	86.98	260.88	25.63	17.24

Note: Each coefficient is the estimate of β from equation 3.2. Outcomes in Panel A are outcomes known to police, while outcomes in Panel B are actual arrests made. Each outcome is per 100k individuals. The dependent variable in Column (1) is all acquisitive crimes. The dependent variable in in Columns (2) - (5) are the component acquisitive crimes: burglary, larceny, motor vehicle theft, and robbery, respectively. $\bar{Y}_{2006}$ is the sample average of the dependent variable in 2006. Standard errors clustered at the county-level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.2: Non-acquisitive crimes

Panel A: Offenses Known				
	(1) All non-acquisitive	(2) Murder	(3) Rape	(4) Assault
$OxyExp \times Post$	1.078 (6.275)	0.0225 (0.0694)	0.204 (0.358)	0.851 (6.119)
Observations	40792	40792	40792	40792
County FEs	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓
$\bar{Y}_{2006}$	994.84	3.24	23.80	967.67

Panel B: Arrests				
$OxyExp \times Post$	-1.961 (1.114)	-0.0272 (0.110)	-0.00157 (0.122)	-1.932 (1.031)
Observations	40792	40726	40726	40726
County FEs	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓
$\bar{Y}_{2006}$	150.72	3.26	7.14	107.21

Note: Each coefficient is the estimate of β from equation 3.3. Outcomes in Panel A are outcomes known to police, while outcomes in Panel B are actual arrests made. Each outcome is per 100k individuals. The dependent variable in Column (1) is all non-acquisitive crimes. The dependent variables in Columns (2) - (4) are the component non-acquisitive crimes: murder, rape, and assault, respectively. $\bar{Y}_{2006}$ is the sample average of the dependent variable in 2006. Standard errors clustered at the county-level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.3: Effects of the reformulation on offenses known (Alternative Measure)

Panel A: Acquisitive Offenses					
	(1) All acquisitive	(2) Burglary	(3) Larceny	(4) Motor vehicle theft	(5) Robbery
$OxyShare \times Post$	-12.39 (11.92)	-3.395 (3.811)	-7.369 (8.038)	-1.743 (1.602)	0.112 (0.541)
Observations	40792	40792	40792	40792	40792
County FEs	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓
$\bar{Y}_{2006}$	2193.02	526.61	1464.71	159.54	41.92
Panel B: Non-acquisitive Offenses					
	(1) All non-acquisitive	(2) Murder	(3) Rape	(4) Assault	
$OxyShare \times Post$	10.69 (7.956)	0.0722 (0.0759)	0.619 (0.334)	9.997 (7.829)	
Observations	40792	40792	40792	40792	
County FEs	✓	✓	✓	✓	
Year FEs	✓	✓	✓	✓	
$\bar{Y}_{2006}$	994.84	3.24	23.80	967.67	

Note: Each coefficient is the estimate of β from equation 3.2. Each dependent variable is offense known per 100k individuals. Panel A is acquisitive offenses and Panel B is non-acquisitive offenses. The dependent variable in Column (1) is all crimes, and the dependent variables in the other columns are the component crimes. $\bar{Y}_{2006}$ is the sample average of the dependent variable in 2006. Standard errors clustered at the county-level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Bibliography

- Adelman, Saul, Mark Cross, and David Shrider**, “Why Do Homeowners Make Mortgage Curtailment Payments?,” *Journal of Housing Research*, 2010, 19 (2), 195–212.
- Agarwal, Sumit, John C Driscoll, and David I Laibson**, “Optimal mortgage refinancing: a closed-form solution,” *Journal of Money, Credit and Banking*, 2013, 45 (4), 591–622.
- Akee, Randall, Emilia Simeonova, William Copeland, Adrian Angold, and E Jane Costello**, “Young adult obesity and household income: effects of unconditional cash transfers,” *American Economic Journal: Applied Economics*, 2013, 5 (2), 1–28.
- Alaska Permanent Fund Corporation**, “Financial Statements: September 30, 2025,” PDF document 2025. Accessed via user-provided file.
- Alpert, Abby, David Powell, and Rosalie Liccardo Pacula**, “Supply-Side Drug Policy in the Presence of Substitutes: Evidence from the Introduction of Abuse-Deterrent Opioids,” *American Economic Journal: Economic Policy*, 2018, 10 (4), 1–35.
- AMA Alliance**, “Prescription Opioid Epidemic: Know the Facts,” 2019. Accessed: 2025-02-23.
- Amromin, Gene, Jennifer Huang, and Clemens Sialm**, “The tradeoff between mortgage prepayments and tax-deferred retirement savings,” *Journal of Public Economics*, 2007, 91 (10), 2014–2040.
- , **Neil Bhutta, and Benjamin J Keys**, “Refinancing, monetary policy, and the credit cycle,” *Annual Review of Financial Economics*, 2020, 12, 67–93.
- Andersen, Steffen, John Y Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai**, “Sources of inaction in household finance: Evidence from the danish mortgage market,” *American Economic Review*, 2020, 110 (10), 3184–3230.
- Andreolli, Michele and Paolo Surico**, “‘Less is More’: Consumer Spending and the Size of Economic Stimulus Payments,” *American Economic Journal: Macroeconomics*, 2026, 18 (1), 34–68.
- Anenberg, Elliot, Therese C Scharlemann, and Eileen Van Straelen**, “Borrowing and Spending in the Money: Debt Substitution and the Cash-out Refinance Channel of Monetary Policy,” 2023.

- Bartholow, Michael**, “Top 200 Drugs of 2010,” 2011. Accessed: 2024-06-04.
- Beachler, D.C., K. Hall, R. Garg, G. Banerjee, L. Li, L. Boulanger, H. Yuce, and A.M. Walker**, “An Evaluation of the Effect of the OxyContin Reformulation on Unintentional Fatal and Nonfatal Overdose,” *Clinical Journal of Pain*, Jun 2022, 38 (6), 396–404.
- Becker, Gary S.**, “Crime and Punishment: An Economic Approach,” *Journal of Political Economy*, March-April 1968, 76 (2), 169–217.
- Beheshti, David and Seth Neller**, “What Causes Geographic Variation in Drug Prescribing? Evidence from Physician Migration,” 2023. Working paper.
- Berger, David W, Konstantin Milbradt, Fabrice Tourre, and Joseph S Vavra**, “Optimal Mortgage Refinancing with Inattention,” Technical Report, National Bureau of Economic Research 2024.
- Berkouwer, Susanna B. and Joshua T. Dean**, “Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households,” *American Economic Review*, 2022, 112 (10), 3291–3330.
- Bhutta, Neil, Jesse Bricker, Andrew C Chang, Lisa J. Dettling, Sarena Goodman, Joanne W Hsu, Kevin B Moore, Sarah Reber, Alice Henriques Volz, and Richard Windle**, “Changes in US Family Finances from 2016 to 2019: Evidence from the Survey of Consumer Finances,” *Federal Reserve Bulletin*, 2020, 106 (5).
- Blanchet, Thomas, Emmanuel Saez, and Gabriel Zucman**, “Real-time inequality,” Technical Report, National Bureau of Economic Research 2022.
- Blumenberg, Adam**, “Avoid These Common Holiday Health Emergencies,” 2024. <https://www.columbiadoctors.org/news/avoid-these-common-holiday-health-emergencies> [Accessed: March 2026].
- Bohm, Michele K, Yong Liu, Marissa B Esser, Jessica B Mesnick, Hua Lu, Yi Pan, and Kurt J Greenlund**, “Binge drinking among adults, by select characteristics and state - United States, 2018,” *MMWR Morb. Mortal. Wkly. Rep.*, October 2021, 70 (41), 1441–1446.
- Brown, Don**, “Are Recast Loans Skewing Agency Speeds?,” 2023. Accessed: 2025-06-04.
- Caetano, Gregorio, Miguel Palacios, and Harry A. Patrinos**, “Measuring Aversion to Debt: An Experiment among Student Loan Candidates,” *Journal of Family and Economic Issues*, 2019, 40, 117–131.
- Case, Anne and Angus Deaton**, “Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century,” *Proceedings of the National Academy of Sciences of the United States of America*, Dec 2015, 112 (49), 15078–15083.
- Cattaneo, Matias D., Richard K. Crump, Max H. Farrell, and Yingjie Feng**, “On Binscatter,” *American Economic Review*, May 2024, 114 (5), 1488–1514.

- Centers for Disease Control and Prevention**, “Understanding the Opioid Overdose Epidemic,” June 2025.
- Chartier, Karen and Raul Caetano**, “Ethnicity and Health Disparities in Alcohol Research,” *Alcohol Research & Health*, 2010, 33 (1–2), 152–160.
- Chetty, Raj, Adam Looney, and Kory Kroft**, “Salience and Taxation: Theory and Evidence,” *American Economic Review*, 2009, 99 (4), 1145–77.
- Choi, James J.**, “Popular Personal Financial Advice versus the Professors,” *Journal of Economic Perspectives*, 2022, 36 (4), 167–192.
- Cicero, Theodore J. and Matthew S. Ellis**, “Abuse-Deterrent Formulations and the Prescription Opioid Abuse Epidemic in the United States: Lessons learned from OxyContin,” *JAMA Psychiatry*, May 2015, 72 (5), 424–430.
- Cocco, Joao F, Francisco J Gomes, and Pascal J Maenhout**, “Consumption and portfolio choice over the life cycle,” *The Review of Financial Studies*, 2005, 18 (2), 491–533.
- Collier, Roger**, “Street versions of opioids more potent and dangerous,” *Canadian Medical Association Journal*, Sep 2013, 185 (12), 1027.
- Collins, Susan E**, “Associations between socioeconomic factors and alcohol outcomes,” *Alcohol Res.*, 2016, 38 (1), 83–94.
- Congressional Research Service**, “The Opioid Crisis in the United States: A Brief History,” Technical Report, Congressional Research Service 2022. Available at <https://crsreports.congress.gov/product/pdf/IF/IF12260>.
- Cookson, J. Anthony, Erik Gilje, and Rawley Heimer**, “Shale shocked: Cash windfalls and household debt repayment,” *Journal of Financial Economics*, 2022.
- Cornaggia, Kimberly, John Hund, Giang Nguyen, and Zihan Ye**, “Opioid Crisis Effects on Municipal Finance,” *The Review of Financial Studies*, April 2022, 35 (4), 2019–2066.
- Currie, Janet and Hannes Schwandt**, “The opioid epidemic was not caused by economic distress but by factors that could be more rapidly addressed,” *The ANNALS of the American Academy of Political and Social Science*, 2021, 695, 276–291.
- Custodio, Claudia, Dragana Cvijanovic, and Moritz Wiedemann**, “Opioid Crisis and Real Estate Prices,” 2024. Available at SSRN: <https://ssrn.com/abstract=3712600> or <http://dx.doi.org/10.2139/ssrn.3712600>.
- Cutler, David, Jonathan S. Skinner, Ariel Dora Stern, and David Wennberg**, “Physician Beliefs and Patient Preferences: A New Look at Regional Variation in Health Care Spending,” *American Economic Journal: Economic Policy*, 2019, 11 (1), 192–221.
- Deiana, Claudio and Ludovica Giua**, “The Intended and Unintended Effects of Opioid Policies on Prescription Opioids and Crime,” *The B.E. Journal of Economic Analysis & Policy*, January 2021.

- Dobkin, Carlos and Steven L. Puller**, “The Effects of Government Transfers on Monthly Cycles in Drug Abuse, Hospitalization and Mortality,” *Journal of Public Economics*, 2007, 91 (11–12), 2137–2157.
- Doleac, Jennifer L. and Anita Mukherjee**, “The Effects of Naloxone Access Laws on Opioid Abuse, Mortality, and Crime,” *The Journal of Law and Economics*, 2022, 65 (2), 211–238.
- Doran, Kelly M., Neloufar Rahai, Ryan P. McCormack, Jacqueline Milian, Donna Shelley, John Rotrosen, and Lillian Gelberg**, “Substance use and homelessness among emergency department patients,” *Drug and Alcohol Dependence*, 2018, 188, 328–333.
- Drug Enforcement Administration (DEA)**, “DEA launches new initiative to combat drug-related violence and overdoses,” February 7 2022.
- Duffy, John and Andreas Orland**, “Liquidity constraints and buffer stock savings: Theory and experimental evidence,” *SSRN Electronic Journal*. Cited on, 2020, p. 5.
- Duncan, Greg J, Katherine Magnuson, Alicia S Kunin-Batson, Hirokazu Yoshikawa, Nathan A Fox, Sarah Halpern-Meehin, Nicholas J Ainsworth, Sarah R Black, Jennifer Mize Nelson, Timothy D Nelson et al.**, “Cash transfers and their effect on maternal and young children’s health: a randomized clinical trial,” *JAMA pediatrics*, 2025, 179 (8), 867–875.
- Eckel, Catherine C., Cathleen Johnson, Claude Montmarquette, and Christian Rojas**, “Debt aversion and the demand for loans for postsecondary education,” *Public Finance Review*, 2007, 35 (2), 233–262.
- Evans, Douglas N and Jeremy R Porter**, “Criminal history and landlord rental decisions: a New York quasi-experimental study,” *J. Exp. Criminol.*, March 2015, 11 (1), 21–42.
- Evans, William N. and Timothy J. Moore**, “The short-term mortality consequences of income receipt,” *Journal of Public Economics*, 2011, 95 (11), 1410–1424. Special Issue: International Seminar for Public Economics on Normative Tax Theory.
- , **Ethan M. J. Lieber, and Patrick Power**, “How the Reformulation of OxyContin Ignited the Heroin Epidemic,” *The Review of Economics and Statistics*, 03 2019, 101 (1), 1–15.
- Farrell, Diana, Fiona Greig, , and Amar Hamoudi**, “Deferred Care: How Tax Refunds Enable Healthcare Spending,” Technical Report, JPMorgan Chase Institute 2018.
- Federal Bureau of Investigation**, “Uniform Crime Reporting (UCR) Program: Crime in the United States, 2010,” Accessed: 2024-06-11 2010. <https://ucr.fbi.gov/crime-in-the-u.s/2010/crime-in-the-u.s.-2010/aboutucrmain>.
- Field, Erica**, “Educational Debt Burden and Career Choice: Evidence from a Financial Aid Experiment at NYU Law School,” *American Economic Journal: Applied Economics*, 2009, 1 (1), 1–21.

- Finkelstein, Anna Chorniy Amy and Matthew J Notowidigdo**, “Paternalistic Social Assistance: Evidence and Implications from Cash vs. In-Kind Transfers,” 2025.
- Freedman, Matthew and Emily G. Owens**, “Your Friends and Neighbors: Localized Economic Development and Criminal Activity,” *The Review of Economics and Statistics*, 05 2016, 98 (2), 233–253.
- Friedman, Milton**, “The permanent income hypothesis,” in “A theory of the consumption function,” Princeton University Press, 1957, pp. 20–37.
- Friesen, Erik Loewen, Jacob Bailey, Sarah Hyett, Sina Sedighi, Mitchell Lennox de Snoo, Kenneth Williams, Rebecca Barry, Anders Erickson, Farid Foroutan, Peter Selby, Laura Rosella, and Paul Kurdyak**, “Hazardous alcohol use and alcohol-related harm in rural and remote communities: a scoping review,” *Lancet Public Health*, February 2022, 7 (2), e177–e187.
- Fryer, Roland G., Paul S. Heaton, Steven D. Levitt, and Kevin M. Murphy**, “Measuring Crack Cocain and Its Impact,” *Economic Inquiry*, 2013, 51 (3), 1651–1681.
- Fu, Q.**, *Mortgage Curtailment and Its Roles in Mortgage Pricing: Theory and Estimates*, University of Wisconsin–Madison, 1998.
- Fulford, Scott L.**, “How important is variability in consumer credit limits?,” *Journal of Monetary Economics*, 2015, 72, 42–63.
- Gale, William, Swati Joshi, Christopher Pulliam, and John Sabelhaus**, “Simulating Income Tax Liabilities in the Survey of Consumer Finances,” *Tax Policy Center*, February 2022. Urban Institute & Brookings Institution.
- Ganong, Peter and Pascal J. Noel**, “Liquidity versus wealth in household debt obligations: Evidence from housing policy in the great recession,” *American Economic Review*, 2020, 110 (10), 3100–3138.
- Genesove, David and Christopher Mayer**, “Loss Aversion and Seller Behavior: Evidence from the Housing Market,” *Quarterly Journal of Economics*, 2001, 116 (4), 1233–1260.
- Gibson, Marcia, Wendy Hearty, and Peter Craig**, “The public health effects of interventions similar to basic income: a scoping review,” *The Lancet Public Health*, 2020, 5 (3), e165–e176.
- Goldstein, Paul J**, “The drugs/violence nexus: A tripartite conceptual framework,” *J. Drug Issues*, October 1985, 15 (4), 493–506.
- Gopalan, Radhakrishnan, Barton H. Hamilton, Jorge Sabat, and David Sovich**, “Aversion to Student Debt? Evidence From Low-Wage Workers,” Working Paper 3843961, SSRN 2021.
- Gross, David B. and Nicholas S. Souleles**, “Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data,” *The Quarterly Journal of Economics*, 2002, 117 (1), 149–185.

- Gross, Tal and Jeremy Tobacman**, “Dangerous Liquidity and the Demand for Health Care: Evidence from the 2008 Stimulus Payments,” *Journal of Human Resources*, 2014, 49 (2), 424–445.
- Guettabi, Mouhcine and Allison Witman**, “Universal Cash Transfers and Prescription Utilization: Evidence from the Alaska Permanent Fund Dividend,” *Journal of Health Economics*, 2023, 90, 102758.
- Han, Benjamin H, Kelly M Doran, and Noa Krawczyk**, “National trends in substance use treatment admissions for opioid use disorder among adults experiencing homelessness,” *J. Subst. Abuse Treat.*, January 2022, 132 (108504), 108504.
- Harris, Matthew C., Lawrence M. Kessler, Matthew N. Murray, and Beth Glenn**, “Prescription Opioids and Labor Market Pains: The Effect of Schedule II Opioids on Labor Force Participation and Unemployment,” *Journal of Human Resources*, October 2020, 55 (4), 1319–1364.
- Hsieh, Chang-Tai**, “Do consumers react to anticipated income changes? Evidence from the Alaska permanent fund,” *American Economic Review*, 2003, 93 (1), 397–405.
- Johnson, Guy and Chris Chamberlain**, “An Evaluation of the Effect of the OxyContin Reformulation on Unintentional Fatal and Nonfatal Overdose,” *Australian Social Work*, 2008, 61 (4), 342–356.
- Jones, Christopher M, Beth Han, Grant T Baldwin, Emily B Einstein, and Wilson M Compton**, “Use of Medication for Opioid Use Disorder Among Adults With Past-Year Opioid Use Disorder in the US, 2021,” *JAMA Network Open*, 2023, 6 (8), e2327488.
- Jones, Damon and Ioana Marinescu**, “The labor market impacts of universal and permanent cash transfers: Evidence from the Alaska Permanent Fund,” *American Economic Journal: Economic Policy*, 2022, 14 (2), 315–340.
- Jurewicz, Amanda, Deborah K Padgett, Ziwei Ran, Donna G Castelblanco, Ryan P McCormack, Lillian Gelberg, Donna Shelley, and Kelly M Doran**, “Social relationships, homelessness, and substance use among emergency department patients,” *Subst. Abus.*, 2022, 43 (1), 573–580.
- Kaestner, Robert and Engy Ziedan**, “Effects of prescription opioids on employment, earnings, marriage, disability and mortality: Evidence from state opioid control policies,” *Labour Economics*, June 2023, 82, 102358.
- Kaplan, Jacob**, *Uniform Crime Reporting (UCR) Program Data: A Practitioner’s Guide* 2021.
- Karimli, Tural**, “Opioid Epidemic and Mortgage Default,” 2022. Available at SSRN: <https://ssrn.com/abstract=4177492> or <http://dx.doi.org/10.2139/ssrn.4177492>.
- Keyes, Katherine M. and Richard Miech**, “Age, period, and cohort effects in heavy episodic drinking in the US from 1985 to 2009,” *Drug and Alcohol Dependence*, 2013, 132 (1), 140–148.

- Keys, Benjamin J., Devin G Pope, and Jaren C Pope**, “Failure to refinance,” *Journal of Financial Economics*, 2016, 122 (3), 482–499.
- Kirby, Kris N, Nancy M Petry, and Warren K Bickel**, “Heroin Addicts Have Higher Discount Rates for Delayed Rewards than Non-Drug-Using Controls,” *Journal of Experimental Psychology*, 1999, 128, 78.
- Kueng, Lorenz**, “Excess sensitivity of high-income consumers,” *The Quarterly Journal of Economics*, 2018, 133 (4), 1693–1751.
- Kuhn, Randall and Dennis P Culhane**, “Applying cluster analysis to test a typology of homelessness by pattern of shelter utilization: Results from the analysis of administrative data,” *American journal of community psychology*, 1998, 26 (2), 207–232.
- Lerner, Michele**, “Recasting’ is another way, besides refinancing, to save money on your mortgage,” 2021. Accessed: 2025-06-04.
- Loewenstein, G and D Prelec**, “Anomalies in Intertemporal Choice: Evidence and an Interpretation,” *The Quarterly Journal of Economics*, 1992, 107 (2), 573–597.
- Macleane, Johanna Catherine, Justine Mallatt, Christopher J. Ruhm, and Kosali I. Simon**, “The Opioid Crisis, Health, Healthcare, and Crime: A Review Of Quasi-Experimental Economic Studies,” Technical Report 29983, National Bureau of Economic Research April 2022.
- Madden, Gregory J, Nancy M Petry, Gary J Badger, and Warren K Bickel**, “Impulsive and Self-Control Choices in Opioid-Dependent Patients and Non-Drug-Using Control Patients: Drug and Monetary Rewards,” *Experimental and Clinical Psychopharmacology*, 1997, 5, 256.
- Mallatt, Justine**, “Policy-Induced Substitution to Illicit Drugs and Implications for Law Enforcement Activity,” *American Journal of Health Economics*, 2022, 8 (1), 30–64.
- Martínez-Marquina, Alejandro and Mike Shi**, “The Opportunity Cost of Debt Aversion,” *American Economic Review*, 2024, 114 (4), 1140–1172.
- McClelland, Rober, Livia Mucciolo, and Safia Sayed**, “New Evidence on the effect of the TCJA on the housing market,” Research Report, Tax Policy Center 2022.
- McCollum, Meagan N., Hong Lee, and R. Kelley Pace**, “Deleveraging and mortgage curtailment,” *Journal of Banking and Finance*, 2015, 60, 60–75.
- McDermott, KW and L Liang**, “Overview of major ambulatory surgeries performed in hospital-owned facilities, 2019. Healthcare Cost and Utilization Project. Published December 2021,” 2023.
- McKetta, Sarah C. and Katherine M. Keyes**, “Trends in U.S. women’s binge drinking in middle adulthood by socioeconomic status, 2006–2018,” *Drug and Alcohol Dependence*, 2020, 212, 108026.
- Meissner, Thomas**, “Intertemporal consumption and debt aversion: an experimental study,” *Experimental Economics*, 2016, 19, 281–298.

- **and David Albrecht**, “Debt Aversion: Theory and Measurement,” *arXiv preprint arXiv:2207.07538*, 2022.
- **and Steffen Altmann**, “Debt Aversion: Theory and Evidence,” *Journal of Monetary Economics*, 2022, 128, 80–97.
- Miller, Norman S., Mark S. Gold, and John C. Mahler**, “Violent behaviors associated with cocaine use: possible pharmacological mechanisms,” *International Journal of the Addictions*, October 1991, 26 (10), 1077–1088.
- Moore, Timothy J. and Kevin T. Schnepel**, “Opioid Use, Mortality Risks and Crime: Insights from a Rapid Reduction in Heroin Supply,” *The Review of Economics and Statistics*, 2024.
- National Drug Intelligence Center**, “Background on OxyContin,” 2001. Accessed: 2024-06-04.
- National Institute on Drug Abuse (NIDA)**, “What are the immediate (short-term) effects of heroin use?,” April 13 2021.
- Nolen-Hoeksema, Susan**, “Gender differences in risk factors and consequences for alcohol use and problems,” *Clinical Psychology Review*, 2004, 24 (8), 981–1010.
- of Alaska Department of Revenue, Permanent Fund Dividend Division State**, “Annual Reports – Permanent Fund Dividend Division,” <https://pfd.alaska.gov/Division-Info/annual-reports>. Accessed: 2025-10-22.
- Olvera, Johabed G., Felipe Lozano-Rojas, Julio A. Ramos Pastrana, and Sumedha Gupta**, “Opioids prescribing restrictions and homelessness: Evidence from hydrocodone rescheduling,” *Journal of Housing Economics*, 2024, 66, 102010.
- Park, Sujeong and David Powell**, “Is the rise in illicit opioids affecting labor supply and disability claiming rates?,” *Journal of Health Economics*, 2021, 76, 102430.
- Patrick, Megan, Patrick Wightman, Robert Schoeni, and John Schulenberg**, “Socioeconomic Status and Substance Use Among Young Adults: A Comparison Across Constructs and Drugs,” *Journal of Studies on Alcohol and Drugs*, 09 2012, 73, 772–782.
- Powell, David and Rosalie Liccardo Pacula**, “The Evolving Consequences of OxyContin Re-formulation on Drug Overdoses,” *American Journal of Health Economics*, 2021, 7 (1), 41–67.
- Prelec, Drazen and George Loewenstein**, “The red and the black: Mental accounting of savings and debt,” *Marketing science*, 1998, 17 (1), 4–28.
- Raimondi, Meredith**, “NCUIH releases new resource highlighting Medicaid’s crucial role for American Indian and Alaska native communities,” Mar 2025.
- Ramsey, Dave**, *The Financial Peace Planner: A Step-by-Step Guide to Restoring Your Family’s Financial Health*, Penguin, 1998.

- Reis, Ricardo**, “Inattentive consumers,” *Journal of monetary Economics*, 2006, 53 (8), 1761–1800.
- Rivera, Olga J. Santiago, Jennifer R. Havens, Maria A. Parker, and James C. Anthony**, “Risk of Heroin Dependence in Newly Incident Heroin Users,” *JAMA Psychiatry*, 08 2018, 75 (8), 863–864.
- Rountree, Janey, Nathan Hess, and Austin Lyke**, “Health conditions among unsheltered adults in the US,” 2019.
- Sadoff, Sally and Anya Samek**, “Can Interventions Affect Commitment Demand? A Field Experiment on Food Choice,” *Journal of Economic Behavior & Organization*, 2019, 158, 90–109.
- Silver, David and Jonathan Zhang**, “Invisible Wounds: How Mental Disability Benefits Shape Veteran Well-Being,” *American Economic Journal: Economic Policy*, 2026, 18 (1), 282–317.
- Sim, Yongbo**, “The effect of opioids on crime: Evidence from the introduction of OxyContin,” *International Review of Law and Economics*, June 2023, 74, 106136.
- Sims, Christopher A**, “Implications of rational inattention,” *Journal of monetary Economics*, 2003, 50 (3), 665–690.
- Stanford Basic Income Lab**, “Global Map of Basic Income Experiments,” 2026.
- State of Alaska, Department of Revenue**, “Permanent Fund Dividend: Eligibility Requirements,” <https://pfd.alaska.gov/eligibility/eligibility-requirements> 2024. Accessed: 2026-03-13.
- Statista**, “Total federal drug control spending in the U.S. from 2008 to 2023,” 2024.
- Substance Abuse and Mental Health Services Administration (SAMHSA)**, “Full-Time College Students: Heavy Alcohol Use and Illicit Drug Use,” 2014. Accessed: 2025-02-22.
- Tan, Bowen**, “The Effects of OxyContin Reformulation on Homicides,” 2024.
- Telyukova, Irina A. and Randall Wright**, “A Model of Money and Credit, with Application to the Credit Card Debt Puzzle,” *The Review of Economic Studies*, 2008, 75 (2), 629–647.
- Thaler, Richard H.**, “Savings, Fungibility, and Mental Accounts,” *Journal of Economic Perspectives*, 1990, 4 (1), 193–205.
- The White House**, “Ending Crime and Disorder on America’s Streets,” <https://www.whitehouse.gov/presidential-actions/2025/07/ending-crime-and-disorder-on-americas-streets/> July 2025. Executive Order.
- U.S. Department of Housing and Urban Development**, “Notice CPD-16-13: Prioritizing Persons Experiencing Chronic Homelessness and Other Vulnerable Homeless Persons in Permanent Supportive Housing,” 2016. Accessed: 2024-11-04.

— , “CoC Homeless Populations and Subpopulations Reports,” 2024. Accessed: 2024-11-04.

Vanguard, “How America Saves 2024,” Technical Report, Vanguard Institutional 2024.

Watson, Brett, Mouhcine Guettabi, and Matthew Reimer, “Universal Cash and Crime,” *The Review of Economics and Statistics*, 10 2020, 102 (4), 678–689.

Woodward, Susan E. and Robert E. Hall, “Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence,” *American Economic Review*, 2012, 102 (7), 3249–3276.

Xu, Yingqi, “Essays on Mortgage Curtailment.” PhD dissertation, Georgetown University 2023.

Yakovlev, Evgeny, “Demand for alcohol consumption in Russia and its implication for mortality,” *American Economic Journal: Applied Economics*, 2018, 10 (1), 106–149.

Zee, Art Van, “The promotion and marketing of OxyContin: commercial triumph, public health tragedy,” *American Journal of Public Health*, 2009, 99 (2), 221–227. Epub 2008 Sep 17.

Appendix A

Chapter 1 Appendix

A.1 Detailed Data Construction

A.1.1 Identifying early payments

Our main sample is the Fannie Mae Single-Family Loan Performance dataset, a monthly panel of conventional loans purchased by Fannie Mae. We observe the mortgage's interest rate, term, balance at origination (rounded to the nearest thousand) and current unpaid principal balance (the exact dollar amount). We do not observe a loan's monthly minimum payment in the data, but we are able to calculate an estimated payment using the monthly interest rate, the original loan principal, and the term of the loan. Specifically, for a given principal P , term (in months) N , and monthly interest rate r_m , we calculate monthly payment MP as

$$MP = \frac{P \times r_m}{1 - (1 + r_m)^{-N}} \quad (\text{A.1})$$

The share of each monthly payment that goes toward interest in a given month is equal to the monthly interest rate r_m times the remaining unpaid principal; the share that is applied toward principal is simply the monthly payment minus that share.

We observe the outstanding principal balance each month for a given mortgage. Using the outstanding balance in month t , we calculate the *expected* outstanding balance in month $t + 1$ – i.e., the balance remaining if the household paid precisely its mortgage payment – by subtracting the principal share of the monthly payment from the outstanding balance in month t .

If the difference between the observed balance and the expected balance is negative, the household has paid in excess of its monthly minimum. We define prepayment to be excess payments greater than \$20 to account for measurement error.

It is not always the case that a household paying in excess of its monthly minimum is prepay-

ment. Some households pay extra, but do not specify that the excess cash should be put toward the remaining principal. Instead, the excess payment is used a future payment. This behavior does not reduce the overall interest burden nor does it shorten the effective maturity date; we do not consider this prepayment, and so we remove the flag for these observations. To do so, we start by flagging all households who pay in excess of \$20 or more in a given month. Then, we remove the flag if we observe that the household “underpaid” by a near-symmetric amount in the next three months. Consider a month in which we observe a household with a \$500 monthly payment, \$450 of which goes toward principal. Suppose this household is flagged as paying an excess of \$50. In the following month, we observe that the household has underpaid by \$50, meaning the difference between expected and actual remaining principal is -50. This means the household did not originally prepay, but rather earmarked the extra cash to go to a future payment.

Lastly, we do not consider any excess payment in the final month of a loan’s term to be a prepayment. This is because we are unable to distinguish if the loan has matured or has been refinanced. Because of this, the estimated probability of prepaying for very small (less than \$2,000 or so) remaining balances is mechanically attenuated. This is not because households are less likely to prepay with small remaining balances – the exact opposite is true (see Figure 1.9). Figure ?? shows the probability of prepaying against remaining balance using the Fannie Mae data. We see the exact pattern as in the credit registry data, but for very small balances, the estimated probabilities are artificially low.

A.1.2 Tax data

We use the 2010-2019 Surveys of Consumer Finances (SCF) and the NBER TAXSIM program to calculate marginal tax rates, using the approach established by Gale et al. (2022). Our approach is as follows.

First, since the public use SCF does not include respondents’ geographic information, we randomly assign households to states. Specifically, we make the probability of being assigned to a state proportional to each state’s population; for example, since California’s population is about 1/9 of the U.S. population as a whole, each household has a 1/9 probability of being assigned to California.

Second, we restrict the set of SCF respondents to closely resemble the households we observe in the Fannie Mae dataset. That is, we remove ARMs, mortgages originated before 1999, etc. Using this subset, we calculate the federal, state, and mortgage interest marginal tax rates for each respondent in the SCF. Using TAXSIM, we then calculate the effective federal and state marginal tax rates with respect to taxpayer earnings for each respondent in the SCF. We also calculate the implicit marginal tax rate with respect to mortgage interest by adding \$100 to each household’s mortgage interest expense; the difference between the baseline tax liability and the new tax liability after adding the \$100 is the implicit marginal tax rate on mortgage interest.

Lastly, since the SCF is a triennial survey, we apply the results for each year to the next two

subsequent years. That is, the rates calculated for 2010 will be the same in 2011 and 2012. Tax rates calculated using the 2019 SCF are used for 2019-2023. Within each state-year-origination balance bin¹, we randomly match SCF respondents' tax rates to loans in the Fannie Mae dataset.

For our benchmark assumptions, we assign federal and state marginal tax rates to 0 for respondents that indicate they have some sort of retirement account.

A.2 Inertia

One possibility is that households begin prepaying in a low interest rate environment when prepayment was beneficial; when rates rise, they fail to reoptimize because prepaying was a sound investment previously, so inattention leads them to believe prepaying remains financially wise. This theory is motivated by the fact that the probability of a household prepaying in month t given prepaying in month $t - 1$ is about 80%. To study this, we identify households whose first prepayment occurred in a low-rate environment (January 2011 - February 2022) and those whose first prepayment occurred in a high-rate environment (March 2022 - November 2024). We remove likely "autopayers" – loans that pay the exact same amount each month for 12 consecutive months – to capture *active* decisions to prepay. We then calculate the conditional probability of prepaying at month $t + k$ given the household has prepaid in the previous k periods for each household type. In other words, we first calculate the unconditional probability of prepaying, then the probability of prepaying conditional on doing so last month, then the probability of prepaying conditional on doing so the last two months, and so on.

Results are shown in Figure 1.11b. The blue line represents loans that first prepaid in a high rate environment, and the red line represents loans that first prepaid in a low rate environment. While the unconditional probabilities differ substantially ($\sim 20\%$ vs $\sim 35\%$) across the two groups, by just two consecutive months later, both groups' conditional probabilities hover asymptotically at about 90%. These results show that households do exhibit persistence in their decisions to partially prepay, but this persistence does not correlate with the interest rate environment when they first began prepaying.

¹Mortgage origination balances are rounded to the nearest \$250,000 and capped at \$1 million.

Figure A.1: Data dictionary

Fannie Mae variables

Name	Description
Mortgage interest rate (r_m)	Pre-tax mortgage interest rate
Origination principal balance	Principal balance of the loan at origination (rounded to the nearest \$1000)
Current principal balance	Remaining principal balance in month t
Term	Length (in months) of the loan
Months to legal maturity	Number of months remaining until the loan is due to be paid in full as defined by the contract
Months to maturity	Number of months remaining until the outstanding principal balance of the mortgage amortizes to a zero balance after considering early payments


TAXSIM variables

Name	Description
Federal marginal tax rate	The estimated federal marginal tax rate on taxpayer earnings
State marginal tax rate	The estimated state marginal tax rate on taxpayer earnings

Constructed variables


Name	Description
Monthly payment (MP)	The estimated monthly payment of the mortgage
Principal share of MP	The amount of the estimated monthly payment that is applied toward remaining principal each month
Expected principal balance	How much of the outstanding principal balance should remain in month t if the household pays exactly the monthly payment in month $t-1$
Curtailement amount	Expected principal balance - current (observed) principal balance
Curtailement flag	Indicator equal to 1 if the curtailement amount > \$20

Figure A.2: Credit Registry vs. Loan Performance Data



Note: Figure plots the probability of prepaying by quarter from 2011Q3 - 2023Q1. The red lines represent Fannie Mae Single-Family Loan Performance data. The solid red line uses the benchmark prepayment flag and the dashed red line uses the benchmark flag + clerical prepayment. The blue line represents the UC-CCP data.


Figure A.3: Share of households prepaying same amount each month




Note: Figure shows the share of households prepaying the exact amount each month by calendar year. Prepayment amounts are rounded to the nearest \$5. Authors' calculations based on the Fannie Mae Single-Family Loan Performance dataset.

Figure A.4: Effects of interest rate spreads on prepayment: binscatter approach

(a) 2011-2024




(b) 2022-2024




Note: Both panels show binscatter estimates of Equation 1.5 controlling for loan fixed effects using Cattaneo et al. (2024) estimator. Panel A uses observations from 2011-2024, and Panel B only uses observations from 2022-2024. In both specifications, we estimate a third-degree polynomial with three smoothness constraints. Authors' calculations based on the Fannie Mae Single-Family Loan Performance dataset.

Figure A.5: Comparing Prepayment – Income Gradients: Positive vs. Negative Rate Spreads



Note: Panel A shows a binned scatter plot of a prepayment indicator against imputed income. Panel B shows a binned scatter plot of prepayment amount against imputed income. Prepayment amounts greater than the 99th percentile are excluded. The blue series is for loans with an interest rate spread ≥ 0 . The red series shows loans with a negative interest rate spread. Author's calculations based on Fannie Mae Single-Family Loan Performance dataset.

Figure A.6: Probability of prepaying against LTV at origination



Note: Figure plots the probability of prepaying against loan-to-value (LTV) ratio at origination. Sample includes only observations the first two years after origination. Authors' calculations based on the Fannie Mae Single Family Loan Performance dataset.

Table A.1: Effects of interest rate spread on the probability of prepaying

	(1)	(2)	(3)	(4)	(5)
Spread	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)
1-year lag		-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)
2-year lag			0.00*** (0.00)	0.00*** (0.00)	0.00*** (0.00)
3-year lag				0.00*** (0.00)	0.00*** (0.00)
4-year lag					0.00*** (0.00)
Loan FEs	✓	✓	✓	✓	✓
Number of lags	0	1	2	3	4
Loan-month obs	17,614,564	13,044,558	9,637,158	6,914,454	4,889,247

Note: Table reports estimates from the linear probability model in Equation (1.5). Authors' calculations based on the Fannie Mae Single-Family Loan Performance dataset, FRED, 2010-2019 Surveys of Consumer Finances and NBER Taxsim. Estimates are from Jan 2011 - Nov 2024. Standard errors (in parentheses) are clustered at the loan-level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.2: Credit outcomes by remaining mortgage principal balance

	(1) 1(Curtail)	(2) 1(New auto)	(3) 1(New CC)	(4) 1(Current delinq.)	(5) Revolving credit utilized
≤ 1 %	0.165*** (0.028)	-0.002 (0.002)	0.000 (0.002)	0.003 (0.004)	-1545.091 (826.099)
(1,5] %	0.070*** (0.007)	-0.001 (0.002)	0.003 (0.002)	0.001 (0.004)	157.683 (442.151)
(5,10] %	0.006*** (0.006)	-0.001 (0.002)	0.000 (0.001)	0.001 (0.003)	0.268 (353.887)
(10,20] %	0.060*** (0.005)	-0.001 (0.001)	0.001 (0.001)	-0.001 (0.002)	-125.820 (288.919)
(20,30] %	0.056*** (0.004)	-0.000 (0.001)	0.000 (0.001)	0.002 (0.002)	-331.914 (263.811)
(30,40] %	0.043*** (0.004)	-0.001 (0.001)	0.001 (0.001)	-0.001 (0.002)	-170.921 (235.557)
(40,50] %	0.026*** (0.003)	-0.001 (0.001)	0.001 (0.001)	-0.000 (0.001)	-308.483 (190.570)
Observations	1,943,845	1,943,845	1,943,845	1,943,845	1,757,625
Loan FEs	✓	✓	✓	✓	✓
Quarter FEs	✓	✓	✓	✓	✓
R^2	0.38	0.06	0.12	0.10	0.71

Note: Table plots estimates from Equation (1.8). Column (1) uses a mortgage curtailment indicator for the dependent variable. Column (2) is a new auto loan indicator. Column (3) is a new credit card indicator. Column (4) is a new delinquency indicator. Column (5) is revolving credit utilized. Authors' calculations based on the University of California Consumer Credit Panel (UC-CCP). Standard errors (in parentheses) are clustered at the borrower-level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.3: Effects of interest rate spread on the probability of prepaying: No autopayers

	All			High Credit Score			Low Credit Score		
Spread	-0.01*** (0.00)	-0.01*** (0.00)	-0.02*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.02*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)
Loan FEs	✓			✓			✓		
Month FEs	✓			✓			✓		
$\bar{Y}$	0.28	0.28	0.28	0.30	0.30	0.21	0.26	0.26	0.26

Note: Table reports estimates from the linear probability model in Equation (1.5). "High" and "low" credit score loans are those with above- and below- median credit scores (773), respectively. $\bar{Y}$ represents the baseline probability of prepaying for each group. Interest rate spread is calculated using post-tax mortgage interest rate and post-tax maturity-matched Treasury yields. $N = 8,567,137$ in Columns (1) and (3). $N = 8,497,184$ in Column (2). $N = 4,309,527$ in Columns (4) and (6). $N = 4,278,426$ in Column (5). $N = 4,257,610$ in Columns (7) and (9). $N = 4,218,698$ in Column (8). Authors' calculations based on the Fannie Mae Single-Family Loan Performance dataset, FRED, 2010-2019 Surveys of Consumer Finances and NBER Taxsim. Estimates are from Jan 2011 - Nov 2024. Standard errors (in parentheses) are clustered at the loan-level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Appendix B

Chapter 2 Appendix

Figure B.1: Total alcohol-related Emergency Department visits by gender




Figure B.2: Total alcohol-related Emergency Department visits by geography




Figure B.3: Total alcohol-related Emergency Department visits by primary payer




Figure B.4: Total alcohol-related Emergency Department visits by income quintile

