Three Essays on Labor and Credit Markets

by

Benjamin J. Keys

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2009

Doctoral Committee:

Professor Charles C. Brown, Chair
Professor Sheldon E. Danziger
Professor John E. DiNardo
Assistant Professor Michael W. L. Elsby

To Meghan

ACKNOWLEDGEMENTS

To borrow a phrase, it took a village to complete this dissertation. I am indebted to my committee, Charlie Brown, John DiNardo, Sheldon Danziger, and Michael Elsby, for their support and guidance. Chapter II is joint work with Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, while Chapter III is co-authored with Brian C. Cadena. I thank Soren Anderson, Michael Barr, Rebecca Blank, John Bound, Kerwin Charles, Adam Cole, Paul Courant, Jane Dokko, David Gunter, Samara Gunter, Miles Kimball, Patrick Kline, Brian Kovak, Jens Ludwig, Karen Pence, Alexandra Resch, James Sallee, Dan Silverman, Jeff Smith, anonymous referees, and numerous seminar participants for their helpful comments and suggestions. I am grateful to the Department of Education's Jacob K. Javits Fellowship program, the University of Michigan's Rackham Pre-Doctoral Fellowship Program, the Population Studies Center, and the National Poverty Center for generous financial support. Finally, I owe a debt of gratitude to my parents, extended family, friends, and most of all my fiancee Meghan for moral support throughout. I can only hope to repay their love and kindness in full...with interest.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER I

Introduction

The first paper explores the effect of job displacement on credit market outcomes. I develop a dynamic model where the relationship between employment shocks and the decision to file for bankruptcy is explicit. The model shows that the bankruptcy decision depends on permanent income shocks, debt levels, and expectations about future income. The magnitude of a negative income shock - and its persistence - determines whether the household files for bankruptcy immediately or delays filing and increases indebtedness in anticipation of filing in a subsequent period.

The model's insights on the timing of the bankruptcy decision motivate an empirical investigation using two independent sources of data. First, using the NLSY, I perform an event-study analysis to identify the lead-lag relationship between job displacement and bankruptcy. In contrast to previous research, I find that job displacement leads to a heightened probability of filing for bankruptcy in the subsequent three years, consistent with the model's predictions. Using county-level data to address aggregate trends, I find that 1000 job losses are associated with 8-11 bankruptcies, that the effects also last 2-3 years, and that the loss of a manufacturing job is three times more likely to lead to bankruptcy than the loss of a non-manufacturing job. The individual- and county-level results suggest that job displacements associated with persistent earnings losses influence the timing of bankruptcies and increase their frequency.

The second paper, co-authored with Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, explores the question of whether securitization influenced screening standards in the subprime mortgage market. Theories of financial intermediation suggest that securitization, the act of converting illiquid loans into liquid securities, could reduce the incentives of financial intermediaries to screen borrowers. We empirically examine this question using a unique dataset on securitized subprime mortgage loan contracts in the United States. We exploit a specific rule of thumb in the lending market, the FICO score of 620, to generate exogenous variation in the ease of securitization and compare the composition and performance of lenders' portfolios around the ad-hoc threshold.

Using a sample of more than two million home purchase loans during the period 2001-2006, we empirically confirm that the number of loans securitized varies systematically around the credit cutoffs. In particular, for loans with a potential for significant "soft" (unobservable) information - "low" documentation loans - we find that there are twice as many loans securitized just above the credit threshold than just below it. Conditional on being securitized, the portfolio that is more likely to be securitized has a 20% higher default rate than a similar risk profile group with a lower probability of securitization. Crucially, these two portfolios have similar observable risk characteristics and loan terms. We use variation across lenders, state foreclosure laws, and the timing of the passage of anti-predatory laws to rule out alternative explanations. Our results suggest that securitization does adversely affect the screening incentives of lenders.

The third paper, co-authored with Brian C. Cadena, seeks an explanation for why college students avoid "free money." Nearly 20 percent of undergraduate students who are offered subsidized Stafford interest-free loans turn them down. An agent with time-consistent preferences would not reject these loans because doing so requires foregoing a significant expansion of consumption possibilities. In contrast, we demonstrate that students with time-inconsistent preferences may optimally choose to turn down subsidized loans to avoid excessive consumption during school. Thus, debt-averse behavior can occur even among students who have no direct distaste for debt.

We use the National Postsecondary Student Aid Study to test a prediction supported by the self-control model: Students will make different borrowing decisions based on the magnitude of the liquidity increase induced by accepting the loan. Some students who live off-campus receive a portion of their loan as a cash refund, which they must manage over the course of the semester. Similar students living on-campus, however, have their loans applied directly to their room and board expenses, mitigating the effect of the loan on liquidity. We exploit this variation to run difference-in-differences models of loan take-up, using students who would not receive a refund in either housing situation to control for unobserved factors correlated with housing location that may affect borrowing behavior.

We find that students who would receive a refund check are seven percentage points less likely to accept their loans than similar students living on-campus. As an additional robustness check, we use the dormitory capacity of the institution as an instrument for a student's housing location. These specifications also show that students living off-campus are more likely to reject the loan, with the IV results somewhat larger in magnitude. This finding is inconsistent with a model that explains the lack of perfect take-up through a simple distaste for debt. It is consistent, however, with a behavioral model in which consumers choose to limit their liquidity as a forward-looking optimal response to anticipated temptation.

CHAPTER II

The Credit Market Consequences of Job Displacement

2.1 Introduction

More than one million households file for bankruptcy each year. The system is designed to help households that are unable to repay their debts regain control of their finances. By limiting the risk associated with borrowing, however, bankruptcy laws create an incentive for individuals to increase their debt. This tension between the desire to give households a "fresh start" and the moral hazard therein has been a central point of conflict in the politics of bankruptcy reform and in the present academic research on bankruptcy. On the one hand, two-thirds of bankruptcy filers cite the loss of a job or other source of income as the main reasons for filing, by far the most commonly provided motive (Sullivan, Warren, and Westbrook 1999, Warren and Tyagi 2003). These findings form the basis for the claim that unanticipated "trigger events" such as job loss, divorce, or health crises cause bankruptcy. On the other hand, some researchers counter that "strategic" behavior drives the decision to file for bankruptcy, as households continue to borrow and wait until the benefit from filing is at a maximum before discharging their debts. In their influential paper, Fay, Hurst, and White (2002) analyze filing patterns in the PSID and argue that "discharge of debt is the dominant consideration in households' decisions to file" (pg. 716).

This paper shows that these two perspectives, rather than being mutually exclusive, are both essential to understanding the personal bankruptcy decision. I develop a dynamic, forward-looking model of household behavior where the relationship between income shocks and the decision to file for bankruptcy is explicit. The model implies that strategic agents respond to adverse events optimally, both in their borrowing patterns and in the likelihood and timing of bankruptcy. Intuitively, the decision to file for bankruptcy is irreversible and costly, and as such, there is an option value to delaying (White 1998). Unanticipated shocks lead to asset positions where filing is financially beneficial, while expectations about future earnings play an important role in both the decision to file and the timing of when to file. The model provides two key predictions: First, job separations and other income shocks can lead to lagged responses of bankruptcy filing. Second, the bankruptcy decision crucially depends on both the magnitude and the expected persistence of the income shock.

I test these predictions using individual-level data from the National Longitudinal Survey of Youth (NLSY) and county aggregate data collected from the U.S. Courts. The effect of job loss on bankruptcy is estimated in the NLSY using an event-study framework that carefully controls for the timing of income shocks. Unlike previous research on bankruptcy, the event-study methodology explicitly addresses the source of exogenous variation and allows for estimation of pre-shock differences in bankruptcy likelihoods. Using this approach, I find that households are four times as likely to file for bankruptcy in the year immediately following a job displacement. Bankruptcy risk then declines in magnitude but persists for two to three years. The persistence of a higher bankruptcy risk after displacement is consistent with the model, which formalizes the option value to delaying filing.

To explore further the implications of the model and to test additional hypotheses raised by the "adverse events" empirical literature, I investigate the impact of divorce and health crises on the household bankruptcy decision. In contrast to previous research, I find that divorce is not a "proximate cause" of bankruptcy, as the likelihood of filing for bankruptcy rises prior to divorce. I also find that the timing of health shocks are highly related to the timing of bankruptcy. Overall, the evidence suggests that plausibly exogenous job displacement and negative health shocks can play a role in predicting future bankruptcies among those at-risk.

Although the NLSY is the best available panel data to study bankruptcy, its small sample size does not yield the statistical power necessary to distinguish the effects of job loss based on the severity of the displacement or the demographics of the displaced. To examine these issues, I use county-level data from the last three decades to estimate the aggregate relationship between bankruptcy and job loss. This independent analysis, using different data and a different empirical specification, yields similar results. I find that 1000 additional job losses are associated with 8-11 bankruptcies and that the effects of job loss persist for two to three years, consistent with the model and corroborating the individual-level results using the NLSY.

To examine the model's prediction that more permanent income shocks are more likely to lead to bankruptcy, I separate the county-level job losses into manufacturing and non-manufacturing jobs. Manufacturing jobs are generally associated with longer tenure relationships and greater firm-specific human capital.[1] Losing a manufacturing job often leads to deeper and more persistent earnings shortfalls (Carrington 1993). Consistent with the model's predictions, I find that the loss of a manufacturing job is three times more likely to lead to bankruptcy than the loss of a non-manufacturing job. This is the first empirical evidence that the structural shift away from the manufacturing sector has contributed to increases in bankruptcy, and confirms that the micro foundations of the dynamic model are supported by the macro patterns in the data.

Separating the effects by county demographics and macroeconomic conditions provides greater insight into the consequences of job loss. I find that job losses are more likely to lead to bankruptcies in counties that are more educated, wealthier, and have a larger fraction of working-age individuals. These results suggest that job loss may be more painful in these types of counties, with losses anticipated to be more permanent, or representing greater destruction of tenure and firm-specific human capital. Similarly, during high-unemployment periods when unemployment durations are expected to be significantly longer, the loss of 1000 job leads to 40 more bankruptcies, while during low-unemployment periods the relationship is small and statistically insignificant. These results provide robustness to the main findings and offer an explanation for the cyclical patterns of bankruptcy observed in the aggregate data.

These two complementary empirical analyses at the micro and aggregate levels contribute to the literature on job loss by providing strong evidence that the consequences of displacement extend into the credit market. In a similar context, Sullivan (2008) finds that households increase their unsecured borrowing via credit cards in response to a short-term earnings shock. Though unemployment spells are usually brief, these short-term shocks can have larger long-term consequences on a worker's well-being. Recent research has documented decreased long-term earnings and consumption, greater marital discord, and even heightened mortality resulting from job losses.[2]

The costs of bankruptcy are steep for the bankruptcy courts, which review more than one million cases per year, for all borrowers, who pay higher interest rates to compensate for the cost of discharged debts, and for the households in jeopardy of default. Timely intervention on the part of policymakers or the private sector potentially could reduce the costs of bankruptcy. In 2005, a new provision to personal bankruptcy law was enacted which requires all debtors to undergo credit counseling prior to filing for a discharge of their debts. However, this feature of the new bankruptcy code has not been successful in deterring filings, as clients receive counseling only after contacting

[1] See, for instance, Brown (1989), Anderson and Meyer (1994), and Topel (1990).

[2] See, for instance, Jacobson, LaLonde and Sullivan (1993) on earnings; Stephens (2001) and Browning and Crossley (2003; 2008) on consumption; Charles and Stephens (2004) on divorce; Sullivan and von Wachter (2007) on mortality.

a bankruptcy lawyer.[3] Because the likelihood of filing for bankruptcy is heightened in the years following a layoff, providing credit counseling at the time of job displacement, or when an individual exhausts Unemployment Insurance benefits, might help some households avoid bankruptcy. It is not clear what form a successful intervention would take, whether it would require targeted extensions of credit, greater repayment flexibility, or forcing households to declare bankruptcy sooner and thus avoid accumulating additional unsecured debt. Designing feasible policy initiatives based on these results is an important direction for future research.

The next section describes the institutional details of filing for bankruptcy and discusses the costs and benefits to doing so. In section 2.3, I develop a model which formally connects job losses, earnings shocks and bankruptcy and provides intuitive predictions about which households are most likely to file. Section 2.4 describes the data, the NLSY, and the event study methods used to identify a strong relationship between the timing of job loss and the timing of bankruptcy. An analysis of aggregate trends in bankruptcy using county-level data is presented in section 2.5. Section 2.6 concludes with policy implications and directions for future research.

2.2 The Costs and Benefits of Filing for Bankruptcy

This paper focuses on the time period prior to the passage of the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005 (BAPCPA), which reformed United States bankruptcy law.[4] Between 1980 and 2004, when the bankruptcy code was largely unchanged and the "insurance" value of personal bankruptcy in the U.S. was considered one of the most generous in the world, filing rates increased from 2.0 per thousand working-age adults in 1980 to 8.5 per thousand in 2004. Households were able to choose between two different options for resolving outstanding debts in bankruptcy court.[5] The first, known as Chapter 7, permitted full discharge of allowable debts after deducting non-exempt assets. Back taxes, alimony, child support, and student loans are not dischargeable liabilities, but all other unsecured debts are discharged under Chapter 7 rules.[6]

In theory, under Chapter 7 any non-exempt assets are forfeited to pay off these debts. In

[3] A recent GAO report (GAO-07-203), appropriately titled "Value of Credit Counseling Requirement is Not Clear," supports the view that counseling would be more effective if individuals with severe credit risks were identified at an earlier date: "Anecdotal evidence suggests that by the time most consumers receive the pre-filing counseling, their financial situations are dire, leaving them with no viable alternative to bankruptcy."

[4] See Ashcraft et al. (2007) for more on the reforms of BAPCPA. The new bankruptcy rules have not altered the fundamental choices made by households regarding the timing of the filing decision.

[5] Outside of the legal system, households can simply cease making payments, thereby forcing creditors to garnish wages or attach liens to property. See Dawsey and Ausubel (2004) for more details on this "informal bankruptcy" option.

[6] Prior to 1998, government-guaranteed student loans were eligible for discharge if they were in repayment for more than seven years. Private student loans made by for-profit lenders were dischargeable in bankruptcy until the 2005 BAPCPA.

practice, however, non-exempt assets usually amount to less than 5% of all debts recovered by creditors (Livshits, MacGee, and Tertilt 2007a). Exemption rules vary by state (Gropp et al. 1997), but generally protect retirement plans such as IRAs and 401(k)s, provide a homestead exemption up to a dollar amount (unlimited in a few states), and grant additional exemptions for automobiles and personal belongings such as clothing.

The alternative to filing for discharge is a reorganization of debts, Chapter 13. Under Chapter 13, households agree to a repayment plan of a portion of their debts, worked out through the bankruptcy court with their creditors. These repayment plans are usually scheduled for 3-5 years, however most Chapter 13 filers fall behind and many re-file in Chapter 7. An AOUSC report found that between 1980 and 1988, only 36% of Chapter 13 filers completed their repayment plan (GAO 1999). Individuals are not allowed to file again for seven years if filing Chapter 7, but can re-file sooner if filing Chapter 13.[7] For the purposes of the micro-level analysis, I pool Chapter 7 and Chapter 13 non-business filings together and treat both filing options as having the same costs and benefits. This choice is standard in the literature, often because of small sample sizes for a given chapter and because the decision of which chapter to file is often made after consulting with a bankruptcy attorney. In the aggregate analysis, however, I distinguish between Chapter 7 and Chapter 13 filings.

In addition to the discharge of eligible debts, another benefit to filing is the suspension of all garnishment and other debt collection techniques. Aggressive collection tactics, such as harassing phone calls, as well as wage garnishment and repossession efforts are often cited as a "last straw" in leading households to file (Luckett 2002). The tangible costs of filing are the fees to file the paperwork and pay a bankruptcy lawyer which are on the order of \$500 - \$1500. The most costly aspect of bankruptcy is the flag placed on one's credit score, which is present for up to ten years and has a strong impact on both access to credit and the price of credit (Musto 1999; Fisher et al. 2004).

An often-discussed intangible cost of filing is the role of "stigma," the emotional punishment inflicted by oneself or one's peers for filing for bankruptcy. While there have been claims that declining stigma can explain some of the growth in bankruptcy filing (Gross and Souleles 2002, Fay et al. 2002), subjective survey research indicates that individuals' distaste for bankruptcy has been relatively constant over time (NORC, as cited in Sullivan et al. 2006). In fact, Sullivan et al. (2006) point out that the increased transparency of searchable online bankruptcy databases has, if anything, likely led to *more* stigmatizing experiences from bankruptcy.

[7] This time between filings has been extended to nine years by the passage of BAPCPA in 2005.

Thus households with negative assets must weigh the benefits (debt discharge) and costs (access and price of future credit, forfeiture of assets) against the alternative of not filing and repaying their outstanding debts. Intuitively, households which experience a "large enough" negative deviation from average lifetime expected earnings such that debt repayment is more painful than the costs of filing should optimally file for bankruptcy. The decision is heavily influenced both by the amount of outstanding debt and the magnitude and persistence of the income shock. This intuition is formalized in the model presented in the next section.

2.3 Model

In this section I present a dynamic model where the relationship between unemployment and bankruptcy is explicit in order to highlight the role of shocks on both the incidence and timing of bankruptcy. The model yields two key predictions: First, job separations and other income shocks can lead to lagged responses of bankruptcy filing. Second, the bankruptcy decision crucially depends on both the magnitude and the expected persistence of the income shock. The model implies that strategic agents respond to adverse events optimally, both in their borrowing patterns and in the likelihood and timing of bankruptcy. Thus, the model shows that "strategic borrowing" and "adverse events" frameworks used to understand bankruptcy filing in the previous empirical literature are not mutually exclusive. This dynamic perspective clarifies the policy implications of increased bankruptcy filing and the potential role for intervention, as discussed in section VI.

The model of household choice presented here is intended to provide qualitative predictions about the relationship between shock persistence and bankruptcy and to motivate the empirical work based on consumers' choices. The model builds on the work of Lawrence (1995), which presents a two-period model where the consumer default decision is static and therefore cannot address the timing of the decision. In the following model, I extend the Lawrence model to allow for employment shocks and to include multiple periods in order to formalize the dynamic role of income shocks in determining the likelihood and timing of bankruptcy.

A number of recent papers have constructed structural models of bankruptcy, which have improved our understanding of the behavior of borrowers and banks in a lending environment that includes the possibility of default.[8] Although these general equilibrium models have at their core the

[8]The models are general equilibrium in the sense that the decisions of both banks and households are endogenous. See the thorough summary by Athreya (2005). Livshits et al. (2007a, 2007b) assess explanations for the increase in bankruptcy rates and find support for an explanation based on a declining cost of filing for bankruptcy. Chatterjee et al. (2007) use their model to examine the welfare implications of a counterfactual policy experiment where means-testing is imposed as in the post-BAPCPA regime, and find large welfare benefits attributed to a decreased interest rate on unsecured debt.

dynamic relationship between income shocks and the bankruptcy decision, they do not emphasize the potential of delayed filing in response to shocks or make an explicit link between unemployment and bankruptcy. I incorporate aspects of the more complex structural models (income shock processes, wage garnishment) to the better capture the costs of bankruptcy and to highlight the timing choices in the bankruptcy decision.

2.3.1 Setup of the Bankruptcy Model

Consider a multi-period model where household income, $y_1, y_2, ..., y_T$, is random in all periods, and when employed, log earnings follow an AR(1) process: $\ln(y_t) = \rho \ln(y_{t-1}) + \epsilon_t$. Households face a risk of unemployment, in which case they receive unemployment insurance benefits, z.[9] The risk of unemployment follows a Markov process, where the probability of staying employed, given employment in the previous period, is $0 \leq \eta_1 \leq 1$ and the probability of staying unemployed, given unemployment in the previous period, is $0 \leq \eta_2 \leq 1$. In other words, $(1 - \eta_1)$ is the separation probability and $(1 - \eta_2)$ is the job finding probability. The values of η_1 and η_2 shape households' expectations of the length their employment and unemployment relationships.[10] In the first period all households begin in the employed state.

As most household debt is shared between spouses, and most bankruptcy petitions are jointly filed, bankruptcy is treated here as a household-level decision. Bankruptcy is not allowed in period 1 but is allowed in all subsequent periods. If a household chooses to file for bankruptcy, they face three punishments in the model. First, they are constrained from the credit market in the period they file and in subsequent periods, able neither to borrow nor to save.[11] This assumption is broadly consistent with the bankruptcy flag which appears on the filer's credit report for up to 10 years, and the potentially prohibitively high cost of obtaining credit

Second, the household pays a portion of their earnings, ϕ, to the bankruptcy court in the year in which they file. This garnishment is intended to represent the inability of households to hide their nonexempt assets from the bankruptcy courts. Finally, the third cost of filing is to repay a portion of the debt even in filing (S for debt service), which will be shown to be necessary for interior optimal borrowing behavior, i.e. not borrowing up to the credit limit in all periods. These costs are built into the model to best fit the real-world punishments from bankruptcy, and are adapted

[9]In the simulations which follow, mean log earnings is chosen to be 4, so mean earnings are around 54 and range from 30 to 100, while unemployment benefits are set to $z = 20$. The qualitative results are not sensitive to the choice of these values, within reason.

[10]When households are unemployed, they receive a "shadow" draw from the distribution of earnings to provide a basis for future earnings expectations if they exit unemployment.

[11]The restriction on saving is included so that households are unable to preserve any liquid assets.

from previous models (see, e.g., Livshits et al. 2007a). The model does not directly incorporate the bankruptcy "stigma" as an additional cost. If stigma was hypothesized to be proportional to household earnings, then a portion of ϕ could be interpreted as such. Similarly, if stigma was considered proportional to the amount of debt discharged, then S would reflect the cost of stigma.

Let $V_t(x_{t-1}, y_t, b_{t-1})$ be the value function for a given debt ($x > 0$) or asset ($x < 0$) level in period t, where b_{t-1} is an indicator for whether the household had previously filed for bankruptcy. The maximized value function for filing for bankruptcy is given by V^B and not filing given by V^N. If households receive a positive income shock then they save, $x < 0$, and earn interest r. If they experience a negative income shock, either due to a low draw from the wage distribution or from an unanticipated unemployment spell, households accumulate debt, $x > 0$, with exogenously determined interest rate $R > r$ charged by the bank to offset write-offs from bankruptcies. Households are assumed to be borrowing constrained up to a fraction of current income.[12]

The model can be solved by backwards induction. The essential features of the multi-period model are described most easily in a three-period setting. In period 3, the final period, the household chooses whether or not to file for bankruptcy, giving the value function in the last period:

$$V_3(x_2, Y_3, b_2) = \max\{V_3^N(x_2, Y_3, b_2), V_3^B(x_2, Y_3, b_2)\}$$

where the household chooses to file only when optimal to do so, $V_3^B > V_3^N$. The payoff to not filing, V_3^N, depends on behavior in the second period and the assets or debts brought forward to the final period. If the household did not file in period 2, $b_2 = 0$, then it consumes its period 3 labor income minus interest payments on borrowing (or interest income from saving):

$$V_3^N(x_2, Y_3, 0) = \begin{cases} u(Y_3 - rx_2) & \text{if } x_2 < 0 \text{ (saving)} \\ u(Y_3 - Rx_2) & \text{if } x_2 > 0 \text{ (borrowing)} \end{cases}$$

If the household did file for bankruptcy in period 2, it simply consumes its period 3 labor income:

$$V_3^N(x_2, Y_3, 1) = u(Y_3)$$

The payoff to filing, V_3^B, is period 3 wages net of garnishment minus the portion of debt which is

[12] This assumption is required so that households do not borrow an infinite amount and then attempt to file for bankruptcy. Even if the interest rate was a function of the amount borrowed, some households might borrow as much as they could until the interest rate were infinite with the full intention of defaulting in the subsequent period.

not forgiven:

$$V_3^B(x_2, Y_3, 0) = u((1-\phi)Y_3 - Sx_2)$$

In making the bankruptcy decision in period 3, the model's final period, the household does not need to consider the lack of access to the credit market in future periods. The household chooses bankruptcy in period 3 when the punishment mechanisms, garnishment ϕ and debt service S, are less painful than repaying the debt accrued in period 2: $(1-\phi)Y_3 - Sx_2 > Y_3 - Rx_2$. If the household saved in period 2, $x_2 < 0$, then there is no benefit to filing for bankruptcy, and the household consumes all of its income and savings, $Y_3 - rx_2$.

In period 2 (and any additional "mid-life" periods in a multiple-period setting), the decision rule is more complicated; the household chooses the amount to consume, c_2, or equivalently the amount to borrow or save, x_2, as well as whether to file for bankruptcy:

$$V_2(x_1, Y_2) = \max\{V_2^N(x_1, Y_2), V_2^B(x_1, Y_2)\}$$

The payoff from not filing, V_2^N, is determined by the household's income draw, Y_2, amount of borrowing or saving in the previous period, x_1, and the expected payoff in period 3, represented by the integral term, which is determined by expectations about the distribution of future income, F:

$$V_2^N(x_1, Y_2) = \begin{cases} \max\limits_{x_2} u(Y_2 + x_2 - rx_1) + \beta \int V_3(x_2, Y_3, 0) dF(Y_3|Y_2) & \text{if } x_1 < 0 \\ \max\limits_{x_2} u(Y_2 + x_2 - Rx_1) + \beta \int V_3(x_2, Y_3, 0) dF(Y_3|Y_2) & \text{if } x_1 > 0 \end{cases}$$

The payoff to filing for bankruptcy, V_2^B, also depends on expectations about future earnings. The bankrupt household consumes period 2 income net of garnishment, minus the portion of debt which is not forgiven:

$$V_2^B(x_1, Y_2) = u((1-\phi)Y_2 - Sx_1) + \beta \int V_3(0, Y_3, 1) dF(Y_3|Y_2)$$

In period 1, the preliminary period, the household chooses how much to borrow or save, x_1, based on their income draw, but cannot file for bankruptcy:

$$V_1(Y_1) = \max_{x_1} u(Y_1 + x_1) + \beta \int V_2(x_1, Y_2) dF(Y_2|Y_1)$$

Despite the simplicity of this three-period model, the optimal x_1^*, x_2^*, and b_2^* do not have clear analytical solutions. Thus it is necessary to select parameter values, functional forms, and simulate households' responses.[13]

The model is simulated to provide an understanding of the household response to income and employment shocks. The qualitative insights that I highlight are captured by the optimal Bellman equation in period 2. Periods 1 and 3 in this setting are discussed in less detail, as the choices made in these periods are designed to capture the dynamic aspects of the household's period 2 decision.[14] The predictions from the model motivate the empirical methodology used in sections IV and V.

2.3.2 Optimal Borrowing and Bankruptcy Decisions

Figure 2.1 shows the period 1 decision of how much to borrow or save, x_1, depending on the values of the income draw, Y_1. The solid line represents the choices of households when bankruptcy is an available option in later periods. Households with low income borrow, but most borrow relatively small amounts. Households with the lowest income borrow such that the borrowing constraint binds, which leads to the flat portion of the borrowing curve. Because of the uncertainty of income and possible job separation in the subsequent periods, households with high income save a significant fraction of their income.

The decision whether to borrow or save in the first period is based not only on the income draw but also on expected future draws and the value of filing for bankruptcy (even though households cannot file in this period). As a counterfactual, the dashed line in Figure 2.1 presents the optimal choices of households when there is no option to file for bankruptcy in later periods. Low-income households borrow much less than when they possess a default option in the future, and even wealthy households save more in a bankruptcy-free world than in a world which allows for a "fresh start." The difference between the solid and the dashed lines highlights the strategic aspect of the optimal bankruptcy decision: borrowing (dis-saving) is everywhere higher when bankruptcy is available, and especially so for households in debt.

The amount borrowed or saved in period 1 is brought into period 2, where the most interesting decisions occur. The household chooses whether to file for bankruptcy or decide to wait and see the income realization in period 3 before filing. Thus delaying filing has an option value. If the household files in period 2, it cannot borrow or save to smooth consumption in the next period,

[13] Household utility is assumed to exhibit constant relative risk aversion (CRRA): $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$, where $1/\sigma$ is the degree of intertemporal elasticity of substitution. The parameters used in this simulation are: $\beta = 0.85$, $\phi = 0.4$, $S = 0.1$, $r = 1.05$, $R = 1.1$, $\sigma = 2$, $\rho = 0.15$, $\sigma_\epsilon = 0.15$, $\eta_1 = 0.95$, $\eta_2 = 0.4$, and an exogenous borrowing limit of 1.5 times current income. The main qualitative results of the model are not sensitive to the choice of parameters within reason.

[14] Understanding banks' optimal lending rules under incomplete information regarding the income and employment processes is an important extension of this model and is left for future research.

and it must pay both the garnishment and the debt service penalties. On the other hand, if the household chooses to borrow, it has the option of filing in the subsequent period, so it will borrow more than if it was required to repay all debt in period 3.

Figure 2.2 shows the borrowing and saving decision in period 2, based on the income draws in periods 1 and 2, Y_1 and Y_2, among the employed. Among employed households, as second period income increases (moving to the right in the graph), households first borrow more in anticipation of bankruptcy in the subsequent period, then begin saving. Those individuals with the worst draws are borrowing constrained and cannot borrow as much as they would like, leading to the flat part of the graph on the left hand side. The individuals with high income in the first period but low in the second choose to draw down their savings but not borrow (the flat part at "0" for the 80th percentile, the "high" income line). The individuals with relatively good income realizations in period 1 and bad income draws in period 2 borrow heavily in anticipation of bankruptcy in the next period. Given that the likelihood of employment in the next period is high for those who are currently employed, $\eta_1 = .95$, none of the employed file for bankruptcy in period 2 in this simulation. If plotted for all percentiles of the Y_1 distribution, the graph would not be symmetric because the expected distribution of future income draws, $dF(Y_3|Y_2)$, depends only on Y_2 and not on Y_1.

The only households who choose to file for bankruptcy in period 2 are those with bad income realizations while employed in period 1 and subsequently unemployed in period 2, who borrowed in period 1 in anticipation of a better outcome and now wish to default on their debts. Given the high value of η_1, unemployment was a relatively low-probability event. Figure 2.3 shows the bankruptcy decision in period 2, depending on the expected persistence of the unemployment shock, η_2. Each point on the graph is from a separate simulation, and represents the maximum value of income in the first period, Y_1, for which a household who is unemployed in period 2 would file for bankruptcy.

For low values of unemployment persistence, the household expects to return to the labor force quickly, so only those households with very low values of Y_1 (and thus very high values of period 1 debt) file for bankruptcy in period 2. However, as the persistence of unemployment increases, more and more households file immediately in response to the unemployment shock. When $\eta_2 = 1$, and the unemployment spell is expected to be permanent, 50% more households file for bankruptcy. This figure highlights the important role which the persistence of shocks and the formation of expectations, only relevant in a dynamic context, play in shaping the household's bankruptcy decision.

2.3.3 Implications of the Model

Based on the income and employment draws in period 2, the household either files immediately, borrows while waiting to see the next period's draw, or implicitly plans to never file by saving. These responses are consistent with both the "strategic" and "non-strategic" explanations for filing in the empirical literature, as those who file immediately respond to the adverse unemployment shock, while those who accumulate debt intend to maximize the value of filing in period 3. In contrast to existing empirical work which analyzes a one-year horizon (e.g. Fay et al. 2002), these results imply that longer windows of observation around a job separation or other income shock are necessary to appropriately identify its dynamic impact.

Changes in expectations can also have a serious impact on the bankruptcy decision. The proxies that have been used for "stigma" have generally been previous filing rates in the same state (e.g. Fay et al. 2002), but these filing rates take into consideration other households' expectations about future earnings. Similarly, Gross and Souleles (2002) argue that a change in the fit of a model could represent a shift in the underlying distaste for debt. However, the decrease in the goodness of fit also may be indicative of changing expectations about future earnings. Thus the proxies for stigma are potentially identified off of the omitted variation in expectations about income trajectories. In addition, differences in expected earnings can allow for two households who appear similar in terms of indebtedness to behave very differently in making bankruptcy filing decisions.

2.4 Microdata Analysis

The central prediction of the model is that the timing and likelihood of bankruptcy are determined by the magnitude and persistence of income and employment shocks. Thus negative household shocks can have delayed effects on the bankruptcy decision. The model provides insight into the inherently dynamic nature of the bankruptcy decision and the role of shocks and expectations in household decision-making. The next two sections investigate the predictions of this model empirically, using both individual and aggregate data as independent tests of the relationship between job separations and bankruptcy.

2.4.1 Data

The NLSY

The National Longitudinal Survey of Youth initiated a panel study of young people aged 14-21 in 1979. The survey was conducted annually until 1994, and has been biennial thereafter. Questions about education, employment, family formation and dissolution, and respondents' health have been asked in every wave of the survey. When the respondents had all reached the legal age of adulthood in 1985, they were asked about their assets and debts (independent of their parents' resources).[15] These questions have expanded as the respondents have aged and accumulated diverse assets such as 401(k)s and stock portfolios.[16]

The core set of questions on assets and debts can be used to estimate the overall net worth of each respondent. I obtained the restricted-license NLSY data in order to identify the respondents' state of residence which determines the relevant bankruptcy exemptions, as discussed above.[17] Thus the benefit of filing for bankruptcy in any given year can be estimated by deducting a respondent's exempted assets from her net worth.

In the wave of the survey conducted in 2004, respondents were asked if they had ever filed for personal bankruptcy and if so, in what year.[18] Respondents also provided the chapter of filing and whether the filing was due to a business failure.[19] In the analysis that follows, I combine Chapter 7 and Chapter 13 filings, and I focus exclusively on non-business filings by omitting any filings which were classified as a Chapter 11 business reorganization or where the respondent reported that filing was due to the failure of a business.

The lack of representative micro-level data on bankruptcies has made understanding the household decision particularly challenging. Evidence based only on surveys of filers cannot identify any timing relationships between shocks and bankruptcy. These surveys lack a control group of individuals who have experienced shocks but not filed for bankruptcy.[20] Unlike previous research which has relied on the PSID, this paper uses the NLSY to investigate the timing and financial determinants of bankruptcy.[21] The NLSY is the best available panel dataset to study the determinants

[15]For more information on the wealth questions in NLSY79, see Zagorsky (1999).

[16]Asset and debt variables have been top-coded for confidentiality purposes and I apply the lowest consistent top-code to all wealth variables. This affects many but not all of the questions regarding asset and debt variables. Unfortunately the uncensored wealth responses are not available, even with the restricted license dataset. All dollar value variables are adjusted by the CPI-U to real values with the year 2000 as the base year.

[17]The restricted license application can be obtained through the BLS website. I thank the BLS staff for their assistance.

[18]Because of the timing of this question, my sample consists of respondents who answered the NLSY survey in 2004.

[19]Two-thirds of filers said they filed for Chapter 7, with the remainder filing Chapter 13.

[20]See Livshits et al. (2007b) for a summary of surveys of bankruptcy filers (Appendix B).

[21]Zagorsky (2007) looks at the correlation between bankruptcy filing and IQ in the NLSY and finds a hump-shaped

and timing of personal bankruptcy, primarily because the asset and debt variables are available at an annual level in the NLSY. Although the wealth questions in the PSID have more detail than those in the NLSY, respondents answer them only every five years (and now biennially). Estimating wealth between PSID supplements would require interpolating wealth data across five-year periods where a bankruptcy may have occurred in the interim.

Evaluating the Quality of Retrospective Data

The two main critiques of retrospective survey data on bankruptcy are that bankruptcies may be underreported and that individuals may not remember the precise timing of their filing date. First, it is possible that respondents do not report events which may have a negative "stigma" attached to them. In their analysis of PSID data, Fay, Hurst, and White (2002) find that the bankruptcy rate in the sample is roughly one-half of the national rate. In other words, there is potentially 50% under-reporting of bankruptcy experiences. In Figure 2.4, I compare the national bankruptcy rate to the filing rate in the NLSY, the PSID, and the Survey of Consumer Finances (SCF), which asked a similar retrospective question of a nationally representative sample in 2004.[22] The national rates are calculated by dividing the total number of non-business bankruptcies by the Census Bureau's estimate of the total number of US households.

Although the NLSY only follows one cohort over time, the level and trend of the filing rate is consistent with aggregate patterns, albeit slightly below the national rate and slightly above that of the PSID. Turning to the overall rate of ever having filed, the NLSY cohort's rate is 13.3%, whereas the rate in the SCF is 12.2% for respondents of the same age range (aged 39-48 in 2004). That the reported annual filing rates are lower in a retrospective survey such as the SCF and the NLSY is not surprising, as some individuals from the cohort interviewed would not have been old enough to file for bankruptcy in 1979, and some individuals who filed multiple times would only be counted as filing once. In the NLSY, 9% of filers say they have filed more than once, yet respondents were given the opportunity to report only one date of bankruptcy filing.

Alternatively, respondents may not remember the timing of their bankruptcy filing, which would lead to measurement error (and potentially inconsistent estimates) in all subsequent analysis. Without administrative confirmation, there is no way to exhaustively assess the magnitude of this problem. One approach is to compare the respondents' reported retrospective date of bankruptcy with

relationship across the IQ distribution, and Zagorsky and Lupica (2008) analyze respondents' post-bankruptcy wealth outcomes.

[22] For confidentiality purposes the SCF assigned responses into two-year periods, which explains why the filing rates are the same in two-year intervals in the figure.

their debt and asset levels which were reported in each survey year. If the bankruptcy information provided in 2004 can predict a break in the asset and debt data provided in each survey, then the timing of the bankruptcy is sensible.[23]

In Figures 2.5-2.7, I confirm that respondents accurately remember the year in which they filed for bankruptcy. The numbers from the figures are reported in Table 2.1. Figure 2.5 shows the total debt reported by bankruptcy filers, plotted against the relative years before or after bankruptcy, relative to the debts of respondents who never filed for bankruptcy.[24] Year 0 is the year of filing for bankruptcy, and the years to the left are years prior to bankruptcy; to the right are years since filing. The plotted points are relative to those who have never filed for bankruptcy to control for time effects (as described below in more detail), so the "0" on the y-axis is equivalent to the mean value of debts for non-filers, $36,961. The figure shows that total debts fall by $15,000 upon discharge, with a large drop in the year reported as the bankruptcy year. The figure also suggests that debts re-accumulate after bankruptcy, and almost as rapidly as prior to bankruptcy.

The increase in total debts in the years following bankruptcy is a surprising pattern in Figure 2.5 given the damage which bankruptcy does to one's credit score. Other questions in the NLSY provide clear evidence that filing for bankruptcy has a large negative impact on post-bankruptcy credit access: over half of filers who applied for credit were rejected or received less than they asked for, compared to only 20% of non-filers who did not receive the loan they desired. This difference remains nearly thirty percentage points even after controlling for income, age, gender, race, marital status, family size, and education. Furthermore, 32% of bankruptcy filers were dissuaded from applying for credit because they anticipated rejection compared to only 13% of non-filers. Although some debts are re-accumulating well before the removal of the bankruptcy flag on the credit report (ten years), these are likely at a high cost of credit.[25]

Figure 2.6 presents the relative average amounts of "other" debts, as classified by the NLSY, which importantly includes credit card debt, around the time of bankruptcy filing. Again, these coefficients are relative to those who never filed for bankruptcy, so the "0" on the y-axis is equivalent to the mean value of debts for non-filers, $2,669. The amount of these "other" debts peaks in the two years prior to bankruptcy, and then falls by nearly $5,000. This component of debt does not re-accumulate in the six years following bankruptcy but begins to increase in years 8-10. In Figure 2.7, the homeownership rate of bankruptcy filers is plotted in a similar fashion, relative to those

[23] Also, if we believe that there is a significant stigma to bankruptcy, then it should be easy for respondents to recall the year in which the filing occurred.

[24] 95% confidence intervals are plotted in dashes, based on standard errors clustered at the individual level.

[25] Some lenders may eagerly lend to these poor-credit households because after filing they have no means of immediate discharge of their debts.

who never filed. As this is a young cohort, the mean homeownership rate for never-filers is only 33% (which should be interpreted as the "0" value on the y-axis in the figure). The fraction of bankruptcy filers owning a home falls by ten percentage points around the timing of bankruptcy. Homeownership does rebound in the years following bankruptcy, which likely contributes to the increase in total debts shown in Figure 2.5 (which includes mortgage debt). These graphs document the challenges to post-bankruptcy credit access and show that the dates reported retrospectively by respondents in 2004 accurately identify the inflection points in debt reported in earlier years.

For completeness, Table 2.2 presents the summary statistics of the NLSY data. The top portion of the table shows the mean values of standard demographic characteristics such as age, race, gender, education and parents' education. The bottom portion of the table provides a summary of what events respondents have experienced by the time of the 2004 survey. Using my definition of non-business bankruptcy, 12% of respondents have filed at some point in their lives. Many more households have experienced a displacement (a head or a spouse on UI), a health problem, or a divorce, with the proportion of the sample for each ranging from 30-48%. These shocks form the basis of the tests of the model, whether households respond to shocks in the timing and likelihood of filing for bankruptcy.

2.4.2 An Event Study Approach to the Bankruptcy Decision

To carefully identify the timing of filing for bankruptcy around plausibly exogenous shocks, I follow the event study framework of Jacobson, LaLonde, and Sullivan (1993), which has been used in many contexts related to job loss (see, e.g., Sullivan and von Wachter 2007). In this framework, the regressions take the form:

$$Y_{it} = \sum_{j=-s}^{s} \alpha_j * 1(RelTime = j) + \beta X_{it} + \gamma_t + \epsilon_{it},$$

where Y_{it} is an indicator for whether or not the respondent filed for bankruptcy in year t, γ_t are year fixed effects, the vector X_{it} is a set of individual-level characteristics. The vector of α_j are relative time dummies which reflect the time pattern of the response to the shock. During the observation window $(-s, s)$, each α_j represents the effect on bankruptcy j years before or after the shock.

I estimate this model using logistic regression, and report odds ratios in the figures and tables which follow, so the relevant null hypotheses are whether the odds ratios are significantly different from one.[26] The individual-level characteristics, X_{it}, are a full set of age dummies, race, and edu-

[26] Using a duration framework and estimating Cox proportional hazards models yields similar results. However,

cation. The core results presented below are qualitatively similar when additional control variables were included. Most notably, specifications which included pre-shock values of the wealth and marital status of the respondent are nearly identical. Pre-shock values were included because of their potentially endogeneity to the income shock (see, e.g. Charles and Stephens 2004). In addition, specifications which include state of residence fixed effects also yield similar results. Standard errors are clustered to allow for arbitrary heteroscedasticity and correlation of errors over time for individuals.[27]

Does Job Displacement Cause Bankruptcy?

Figure 2.8 shows the pattern of the relative year coefficients for male job losses, defined as the respondent's first time on Unemployment Insurance (UI), with the coefficients reported in Column 1 of Table 2.3.[28] The sample consists of all male respondents who worked full-time in the year previous to a UI spell, with the control group of full-time male workers who never experienced a UI spell (worked at least 45 weeks). I use two-year bins because the biennial NLSY surveys do not allow me to identify responses within the two year window. The coefficients are odds ratios relative to the group of respondents who have never received UI benefits (never experienced job loss while covered by Unemployment Insurance). There is a spike in the odds ratio in the year in which the job loss is experienced: households which experience a job displacement are over four times as likely to file for bankruptcy than those which have not lost a job. The heightened likelihood of bankruptcy filing then falls to 2.12 times the likelihood of the never-unemployed, though still significantly different, in the 2-3 years after bankruptcy. Subsequent years are no longer significantly different from the group who never experienced a job loss.

In addition, prior to job separation there is no difference in the likelihood of filing for bankruptcy between future job losers and the never-unemployed. The estimated coefficients of the pre-displacement relative time indicators are not significantly different from one, the baseline hazard. Furthermore, the coefficient for the period immediately prior to job loss (-1 to -2 years) is statistically different from the time of job loss (year 0 and 1), with the null ($\alpha_0 = \alpha_{-1}$) rejected at the 8% level by an F-test (reported in the last row of Table 2.3). Thus the bankruptcy hazard in the year of job separation is not only significantly different from the non-separation control group, but is also different

there is not a natural "spell" to use in this context. See Fisher (2002) and Gross and Souleles (2002) for examples of duration analysis of personal bankruptcy.

[27]Because of the clustered sampling structure of NLSY, it may be desirable to allow for unspecified correlation at the level of the sampling stratum. Estimates of the main results using standard errors clustered by the sampling strata yield almost identical confidence intervals and are available from the author upon request.

[28]I focus on the first displacement because of the potential endogeneity of subsequent displacements. See Stevens (1997) for a careful analysis of the role of additional displacements on earnings and wage losses.

from the years prior to job separation.

Figure 2.9 shows the relative job loss year coefficients for female job separations. These coefficients exhibit a similar pattern to those of men, with 3.1 times higher likelihood of filing in the year of job displacement and the year following displacement. However, the test of pre-separation coefficients is weaker for women, as the test of the null of equality of the pre-separation coefficient and the year of separation ($\alpha_0 = \alpha_{-1}$) cannot be rejected (p=0.18). While the statistical relationship is weaker, the pattern of coefficients is consistent with the timing of bankruptcy depending on the timing of job loss.

The results for job separations confirm that the timing of bankruptcy is strongly related to the timing of job displacement, particularly for male respondents. Unlike the previous empirical methodology on bankruptcy, the event study framework allows for estimation of pre-shock differences in bankruptcy likelihoods, and to test the coefficients across years. The results support the dynamic forward-looking model presented above, which predicted that households which suffer large negative shocks would file immediately upon receiving information about future employment and permanent income, while other households will delay filing as its effect on permanent income may not be immediately known. These findings highlight the role of "strategic" timing, even in response to adverse events, reconciling the arguments of the previous empirical literature in a dynamic context.

Do Divorce or Health Shocks Cause Bankruptcy?

Proponents of the adverse events hypothesis also suggest that divorce and health problems lead directly to bankruptcy. I test these additional claims in Figures 2.10 and 2.11 (with the coefficients reported in Table 2.3). Figure 2.10 presents the odds ratios by relative time of divorce, defined here as being married in year (t-1) and unmarried in year t. The control group is those individuals who have been married and never divorced. Although significantly different from those who never filed for divorce, the bankruptcy likelihood begins increasing one to two years prior to divorce. Further, we cannot reject the test of equality of the pre-divorce and year of divorce coefficients ($\alpha_0 = \alpha_{-1}$). Thus while bankruptcy is correlated with marital separation, the results presented here suggest that divorce is also related to money problems on their own.

There is less power in the NLSY to detect an impact on bankruptcy from a negative health shock, as the NLSY cohort is (for the most part) young and healthy. Nonetheless, 36% of the sample has experienced a health limitation which has reduced their ability to work at some point.

Following Burkhauser and Daly (1996), I define a health shock as the first time a respondent reports being healthy for one period and then limited for two consecutive periods.[29] The results based on this negative health shock, shown in Figure 2.11, show an increased likelihood of bankruptcy at the time of the shock, with unhealthy individuals 2.6 times more likely to file for bankruptcy than those who never experienced this pattern of health. An F-test for the equality of the pre-shock and year-of-shock coefficients ($\alpha_0 = \alpha_{-1}$) can be rejected at the 6% level. These results, which are presented in the fourth column of Table 2.3, suggest that health shocks appear to act as a direct "trigger" for bankruptcy independent of income and wealth.

2.5 Aggregate Analysis

While the results above suggest a strong relationship between job loss and bankruptcy in the cross-section, the NLSY follows only one cohort over time, and the sample is not large enough to detect differences in the effects of job loss based on the severity of the displacement or the demographics of the displaced. I thus turn to an aggregate analysis to investigate the relationship between job losses and bankruptcy using county-level data. This independent empirical approach corroborates the evidence provided using the individual-level panel data, and reinforces the importance of the model in interpreting the employment-bankruptcy relationship.

The bankruptcy data are collected from the Administrative Office of the U.S. Courts. The dataset contains the number of business and non-business filings, by chapter, for each county for each year from 1980-2004, but no information is collected on the causes of bankruptcy or the characteristics of the filer at this level.[30] Data on employment at the county level is collected from County Business Patterns (CBP) from 1980-2004. As described earlier, the bankruptcy code was essentially unchanged from 1980-2004.[31] I construct county-level measures of manufacturing, non-manufacturing, and total employment for 1980-2004.[32] The model suggests that the bankruptcy decision should be made on the basis of new information. As such, in the specifications which follow the change in the number of jobs is the independent variable of interest, rather than the stock of jobs at a given time.

[29] This stricter definition of health shocks only applies to 10% of respondents.

[30] To the best of my knowledge, these data have only been used in the recent literature to address the consequences of expanded access to casino gambling (Evans and Topoleski 2002; Barron, Staten, and Wilshusen 2002).

[31] Amendments have modified some exemption rules and changes were made in 1984 intended to limit write-offs from debts incurred immediately prior to bankruptcy, so-called "bad faith" debts.

[32] These measures use the appropriate NAICS and SIC codes (2-digit classifications), which have changed over time. Some values are coded as a range for confidentiality purposes. I impute using the midpoints of the ranges provided. Data limitations prevent the use of finer sub-classifications, as well as alternative precise indicators of the severity of job loss.

I regress the total number of new personal bankruptcies in county i in year t, Y_{it}, on the change in the number of jobs, $\Delta Jobs_{it}$, in the same county:

$$Y_{it} = \beta \Delta Jobs_{it} + \gamma_i + \mu_t + \epsilon_{it}$$

To control for time and location differences, I include both year dummies (μ_t) and fixed effects for all 3135 counties (γ_i).[33] The year dummies remove the trends in bankruptcy filing at the national level, as well as any cyclical aggregate variation. The county-specific fixed effects partial out the time-invariant characteristics of counties and account for the fact that some counties would have many more bankruptcies due solely to population size, even in the absence of employment shocks. As the independent variable of interest, $\Delta Jobs_{it}$, is expressed as a deviation from the previous year rather than levels, this fixed effects specification directly addresses the concern that some counties may have persistent job destruction and large numbers of bankruptcies for unobserved reasons.

The results from the aggregate analysis are presented in Table 2.4. All else equal, counties which experience more job losses have a greater number of bankruptcies. The top panel of Table 2.4, Column 1 shows that 1000 additional jobs lost in a county lead to 11 more bankruptcies, even after accounting for the fixed attributes of the county and the macroeconomic conditions in the year of observation. Columns 2 and 3 add in lagged changes in jobs, and the point estimates are similar for the lags as well as the change in the current year, which suggests that job losses have lasting effects on the local bankruptcy rate. The effects are significant for two to three years, but additional lag terms are not statistically significant (see Column 4). This finding is highly consistent with the timing of the effects of unemployment shocks in the NLSY. That the aggregate results support the findings from the microdata further establishes the importance of the dynamic aspect of the household bankruptcy decision.

Manufacturing jobs are generally more likely to be unionized, have longer tenures, and provide better health care and pensions than non-manufacturing jobs. These are also jobs where the accumulation of specific human capital may be particularly important in determining the costs of job separation (Topel 1990, Carrington 1993). The lower panel of Table 2.4 separates the county-level changes in jobs by manufacturing and non-manufacturing job changes. I find that manufacturing job losses are three times more likely to lead to bankruptcy than non-manufacturing jobs. For every 1000 manufacturing jobs lost in a county there are more than 30 additional bankruptcies.

[33] A small number of counties have changed boundaries over this period, and I construct consistent county definitions across all 25 years where necessary.

These results suggest that the changing structure of employment, towards shorter-tenure jobs and away from manufacturing industries, which provided steadier employment, has been a contributing factor to the growth in consumer bankruptcies. In terms of the model presented in Section 2.3, the manufacturing losses have been both more severe (in terms of dollar magnitude) and more persistent (in terms of future earnings) than non-manufacturing job displacements.

As unemployment durations during recessions are much longer on average than during booms (see, e.g. Valletta 2005, CBO 2007), as an additional test, I interact the job change coefficient with whether the aggregate unemployment rate is above the median for the sample period. This interaction specification focuses attention on job losses in those counties which were more affected by high unemployment rates than others. The results are presented in the first column of Table 2.5. For 1977-2004, the median annual national unemployment rate was 6%. During high-unemployment periods, when unemployment durations are expected to be long, the loss of 1000 additional jobs leads to 48 more bankruptcies, while during low-unemployment periods the relationship is small and statistically insignificant. Thus the cyclical component of unemployment has a meaningful impact on the relationship between job displacement and bankruptcy, and provides an explanation for the cyclical pattern in the national bankruptcy filing rate observed in Figure 2.4.

What types of counties are driving these effects? Table 2.5 explores this question by including the interaction of demographic characteristics of counties in 1980 with the number of job losses in the county. The main effect of the demographic characteristic in the county is removed by the county fixed effect, as the measure is not time-varying. Column 2 shows that job losses in more educated counties are more likely to lead to bankruptcy. The coefficient on "Change in Total Jobs" provides the estimated relationship in less educated counties (defined as below the median in fraction with a high school degree), while the coefficient on the interaction of the change in total jobs and the education variable gives the *difference* between above-median education and below-median education counties. Thus the effect of job loss on bankruptcy in highly educated counties is the sum of the coefficients, $-0.001 + -0.010 = -0.011$.

Columns 3 through 7 of Table 2.5 display the results from similar specifications, with different county demographic characteristics interacted with the county's change in jobs. Separating the effects in this manner shows that job losses are more likely to lead to bankruptcies in counties that are more educated, wealthier, and younger, as well as those counties with a smaller fraction of African-American households. These results suggest that job loss may be more painful in these types of counties, with losses anticipated to be more permanent, or greater destruction of tenure and firm-specific human capital. An additional explanation may simply be that individuals living in wealthier counties were more able to receive credit from lenders and thus amass larger amounts of debt. Alternatively, households in these counties may be more willing to respond to a job loss by filing for bankruptcy for other reasons, such as a greater willingness to contact a lawyer or a lower community-wide stigma from filing. The model predicts increased likelihood of bankruptcy filing when losses are more severe or more persistent, and the results in Table 2.5 appear consistent with that prediction.

Separate specifications for bankruptcies by type of filing suggest that job loss has a bigger impact on Chapter 7 filings than on Chapter 13 filings. Based on these specifications, 1000 additional lost jobs lead to 7.4 new Chapter 7 filings and 3.2 new Chapter 13 filings, which adds up to the total filings coefficient of 10.6 in Column 1 of the top panel of Table 2.4. The effect of losing a manufacturing job is seven times larger than losing a non-manufacturing job on Chapter 7 filings, whereas the effect of each type is roughly equivalent for Chapter 13 filings. That manufacturing job losses are so strongly related to Chapter 7 filings fits the model's prediction as Chapter 7 filings could be taken as evidence of a lower likelihood of even partial repayment (as in a Chapter 13 reorganization).

Although the bankruptcy and job loss variables are estimated as deviations from county means, larger counties would likely have more bankruptcies than smaller counties even in deviations. When I include measures of county size, such as the number of households, to address this potentially confounding relationship, the coefficient on job changes is not significantly different: An additional 1000 jobs lost in a county now leads to 10 more bankruptcies rather than 11. A specification which normalizes both the measures of bankruptcies and and job changes by the size of the county population yields a similar result.[34] These specifications suggest that differences in county size are not driving factors in the observed relationship between job loss and bankruptcy.

To address the potentially confounding role of localized growth patterns in bankruptcy, county-specific linear and quadratic time trends were added to the fixed-effects regressions (results not shown). The results are similar, as 1000 lost jobs are now associated with 6-8 more bankruptcies, the coefficients remain statistically significant, and manufacturing job losses continue to drive this result. However, the impact of lagged changes in jobs are smaller and generally less statistically significant. These checks verify that the results are not merely an artifact of correlated trends in job losses and bankruptcy at the local level, as they are robust to controlling for county-specific patterns of growth in filings.

[34] Results are available upon request.

Because the people in the county who file for bankruptcy are not necessarily the same people who suffered the loss of a job, job losses could have both direct and indirect effects on bankruptcy in this aggregate analysis. In the direct case, the household which loses the job also files for bankruptcy. If the job losses have an indirect effect on other members of the regional economy (through the service sector or the housing market, for instance) then these general equilibrium effects would also be included in the reduced-form estimates. Note that if the indirect effects take time to develop, they could give the illusion that lagged job losses matter in predicting bankruptcy, when in fact the household-level direct response could be instantaneous.

The results from the county-level data confirm the predictions of the dynamic model and corroborate the findings using the individual-level panel data of the NLSY. Job displacement increases the likelihood of bankruptcy and has persistent effects over two to three years after separation, and the effects are strongest when the job losses are likely the source of permanent negative income shocks.

2.6 Discussion and Policy Implications

Filing for bankruptcy has become so common that over 13% of all U.S. households aged 39-48 have experienced it at some point in their lives. And yet economists know very little about the determinants of bankruptcy, due in large part to the lack of representative micro-level data with information on bankruptcies. Evidence based only on surveys of filers cannot identify any timing relationships between shocks and bankruptcy, while previous longitudinal studies have used the PSID, which has long intervals between wealth surveys. Unlike previous research, this paper takes advantage of a new retrospective bankruptcy question and annual measures of wealth in the NLSY to identify the role which job loss plays in the timing and likelihood of filing for bankruptcy.[35]

By employing an event-study methodology, I find that households are over four times more likely to file for bankruptcy immediately following a job displacement, and that the effect diminishes but persists for 2-3 years. This result is robust to the inclusion of additional individual controls such as pre-displacement wealth. Using aggregate data, I find that job losses at the county level also have persistent effects on the number of bankruptcies in the county, on the order of 11 per year for three years for an additional 1000 lost jobs. The effects are much larger for the loss of manufacturing jobs, jobs lost during recessions, and jobs lost in counties which are highly educated. Stronger

[35] These two datasets, the PSID and the NLSY, are the only longitudinal surveys available for research on bankruptcy. However, a much more detailed analysis could be accomplished by matching administrative records on employment (e.g. the LEHD) to administrative bankruptcy records.

effects in more educated counties and in times when the expected unemployment duration is long support the model's predictions about the consequences of permanent income shocks and losses of firm-specific human capital on bankruptcy risk.

The results for both the NLSY and the county-level analysis suggest a pattern of bankruptcy filing in response to negative labor market shocks which is consistent with the model of the household bankruptcy decision presented in section 2.3. The filing and borrowing behavior in the model reconciles some aspects of the "strategic" and "non-strategic" motivations for filing discussed in the empirical literature, and clarifies the relationship between income shocks and bankruptcy in existing structural models. The results reinforce the importance of dynamic micro-foundations in interpreting both household decision-making and the aggregate patterns in unemployment and bankruptcy. In addition, the effects of job displacement in both the NLSY and the county-level data last for roughly two to three years, an empirical regularity across independent datasets and different estimation methodologies.

The ambiguous theoretical distinction between "adverse events" and "strategic behavior" has confused the policy implications of bankruptcy reform. Even the 2005 bankruptcy reform act suffers from this false dichotomy, as it addresses both "Abuse Prevention" and "Consumer Protection" to cater to both sides of the debate. Instead, the focus of policymakers should be to identify timely policy interventions. The results of this paper suggest an intervention of credit counseling at the time of job loss could be beneficial to at-risk households. Just as "rapid response" efforts are made during mass layoffs, the UI system could partner with credit counselors to assess household's spending patterns, current financial outlook, and credit market opportunities. Whether individuals should be extended more or less credit at the time of job loss is an open question and one that is deserving of further investigation.

Notably, the duration of the unemployment shocks analyzed in this paper are brief, only 8-10 weeks on average. That a relatively brief layoff can lead to a serious credit crisis such as bankruptcy points to a puzzle: Why are households unable to insure against or smooth consumption around these shocks? Easing credit constraints should theoretically improve households' ability to smooth consumption; however, there has been a marked increase in consumption volatility over the last 25 years (Gorbachev 2007; Keys 2008). For some households, credit expansion clearly has not kept pace with the growth in earnings volatility as documented by Moffitt and Gottschalk (2002) and Shin and Solon (2008). Reconciling these patterns is a topic for future consideration.

Overall, the results suggest that labor market shocks are crucial to understanding the timing and likelihood of bankruptcy. The bankruptcy decision relies on not only current income and wealth, but also expectations about future employment and earnings possibilities. The role of expectations has been ignored in many discussions of “stigma” in the empirical literature, and has been overlooked in its influence on bankruptcy “triggers.” Incorporating a dynamic model of bankruptcy into the empirical literature is a central goal of this paper, but further work is needed to better understand individuals’ formation of expectations about the labor market and its consequences.

Figure 2.1: The borrowing and saving decision in period 1

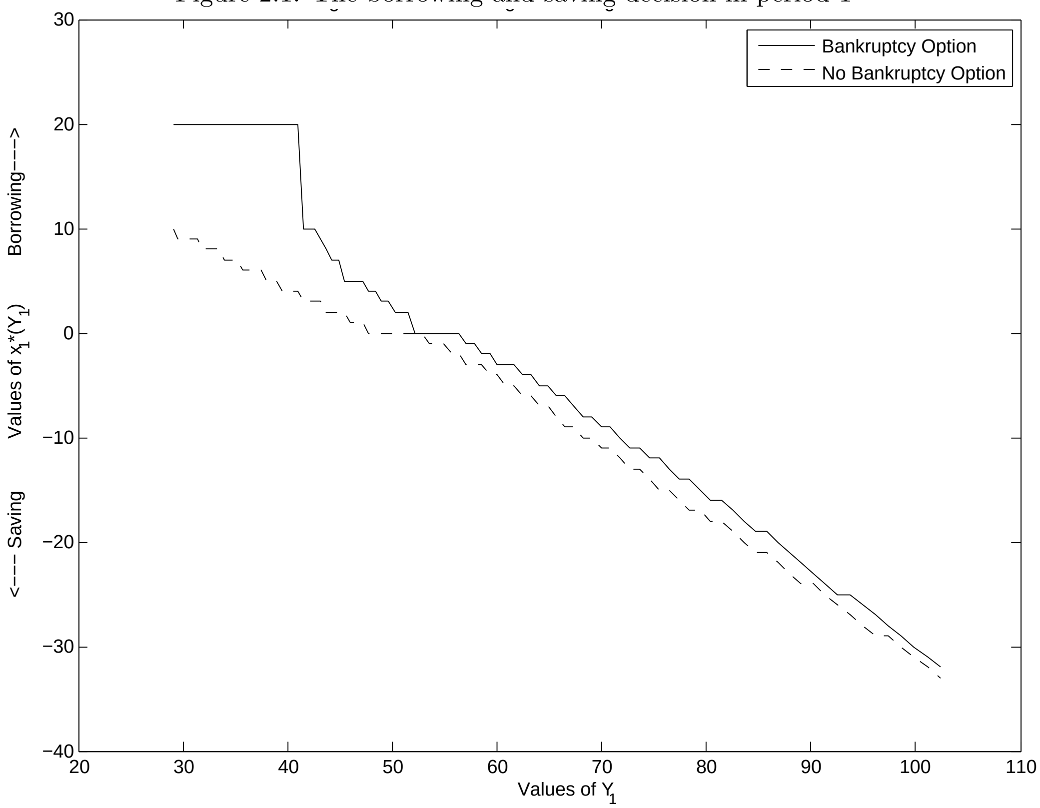

Figure 2.2: The borrowing and saving decision in period 2 - Employed

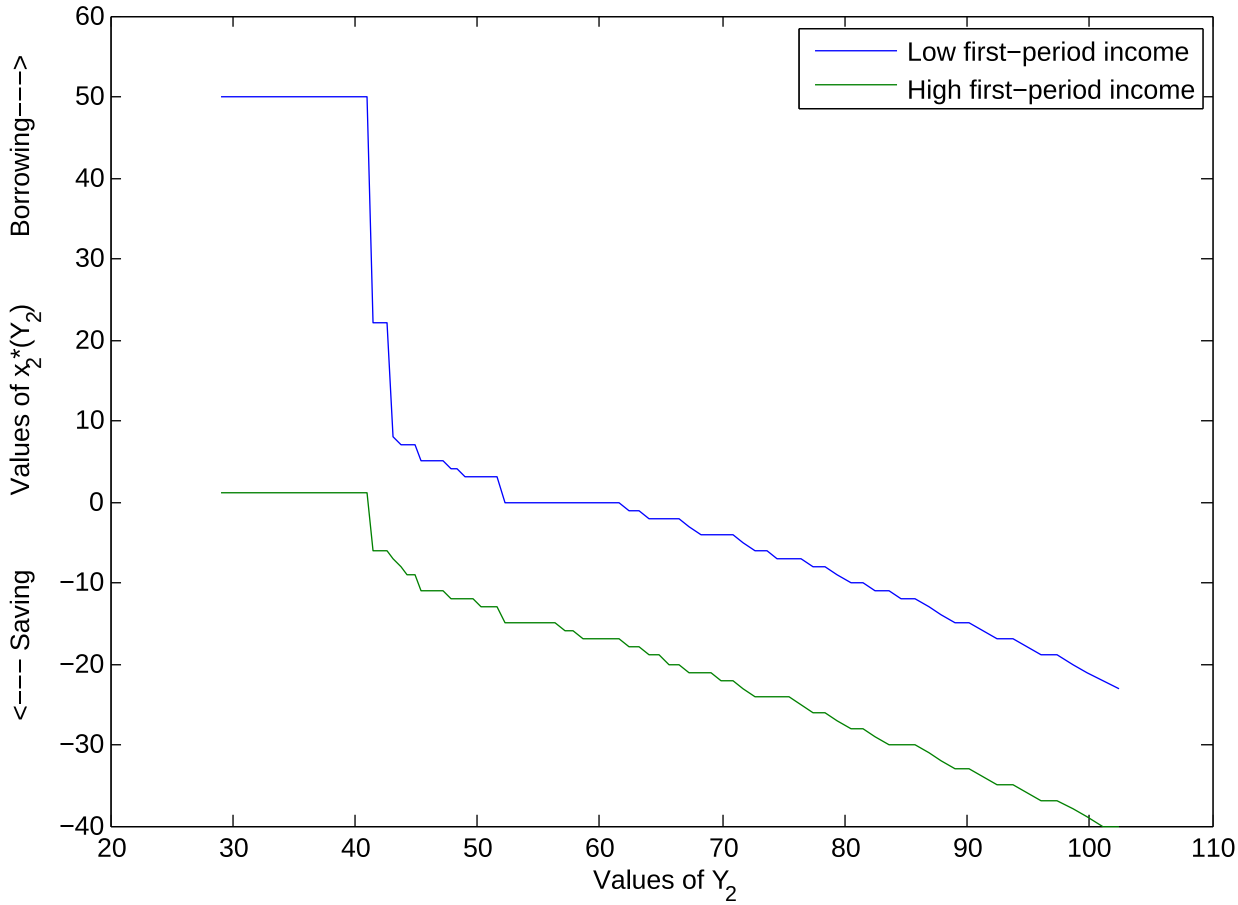

Figure 2.3: The bankruptcy decision in period 2 – Maximum value of Y_1 which results in bankruptcy, by persistence of unemployment

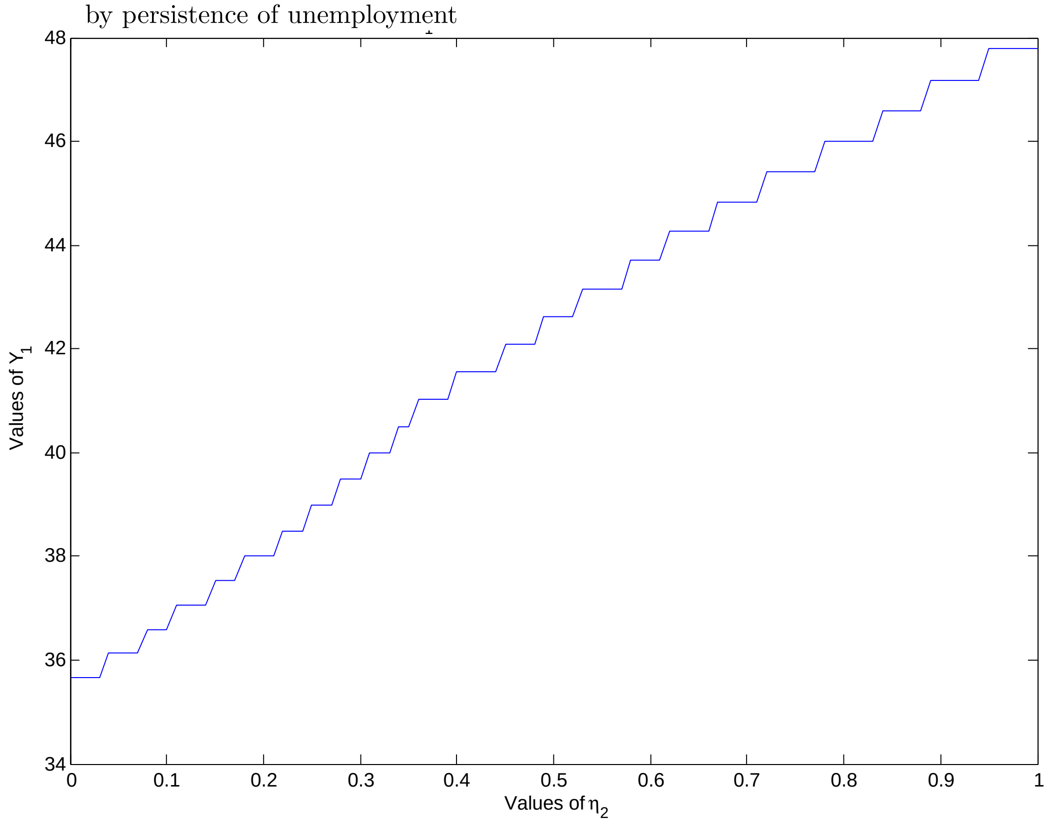

Figure 2.4: Comparison of NLSY, PSID, and SCF filing rates to national filing rate, 1979-2002

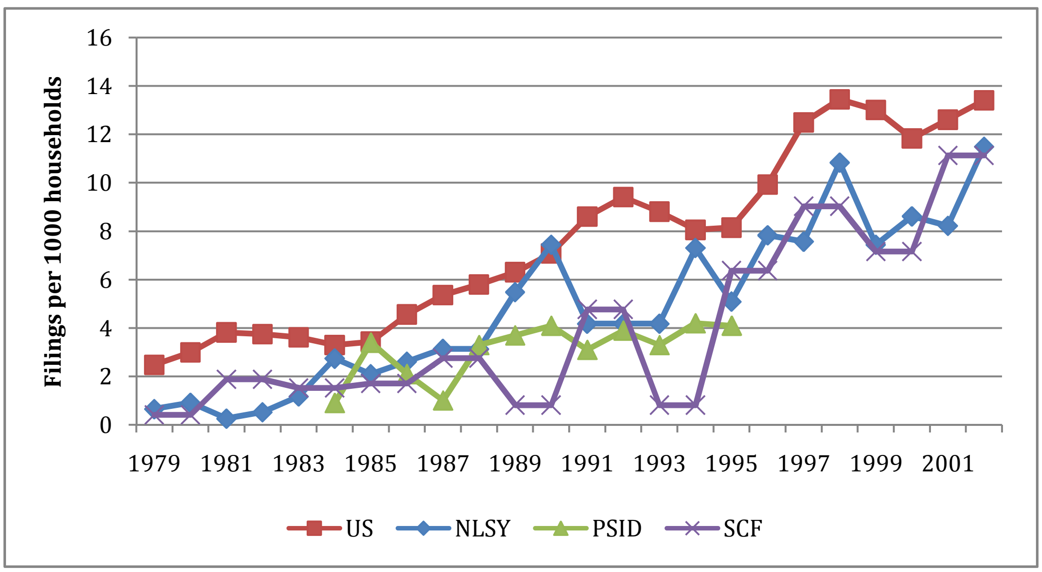

Figure 2.5: Total debts of bankruptcy filers, relative to non-filers, by time of bankruptcy shock

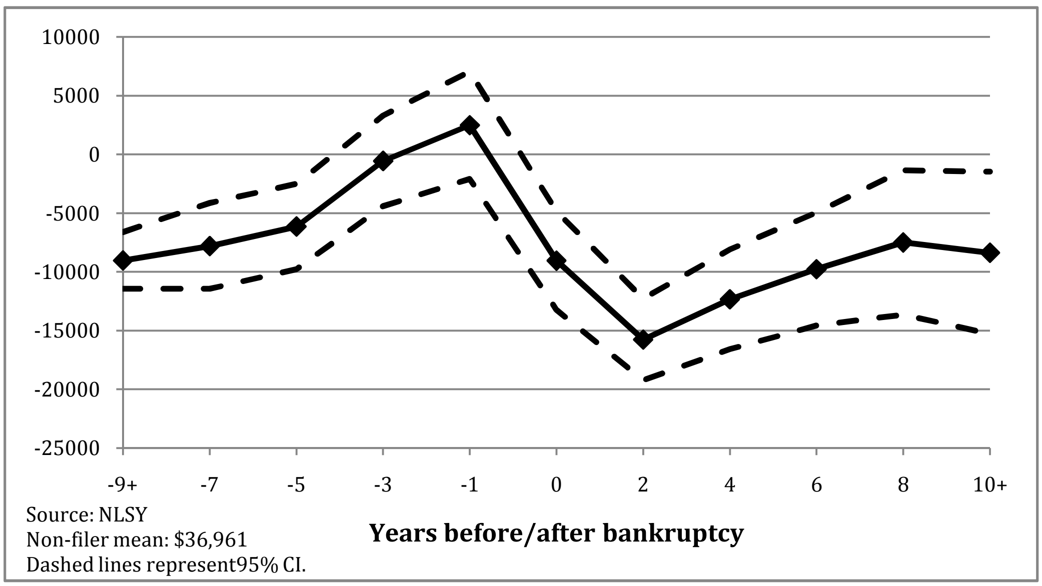

Figure 2.6: “Other” debts of bankruptcy filers, relative to non-filers, by time of bankruptcy shock

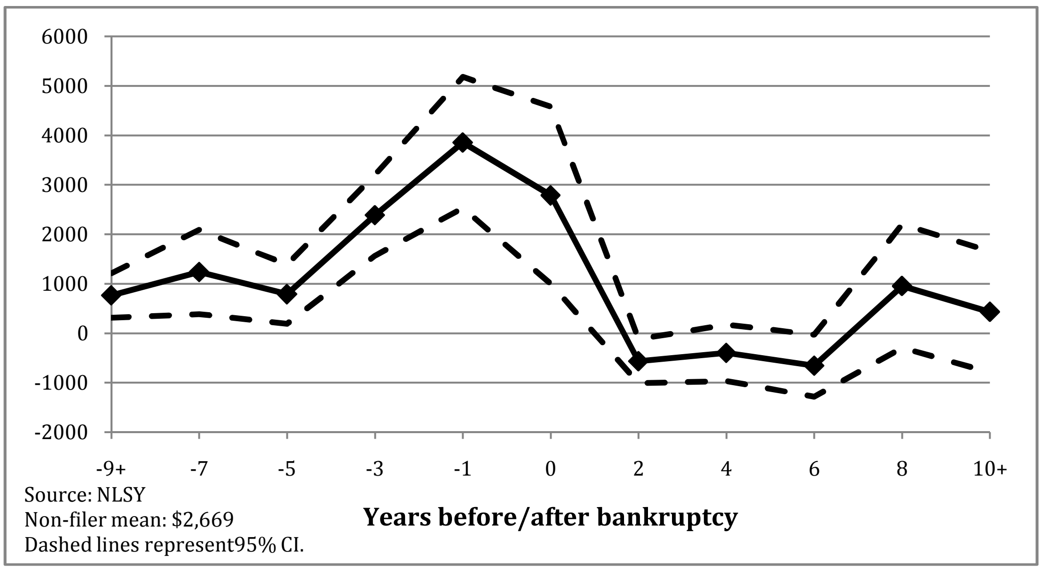

Figure 2.7: Homeownership rates of bankruptcy filers, relative to non-filers, by time of bankruptcy shock

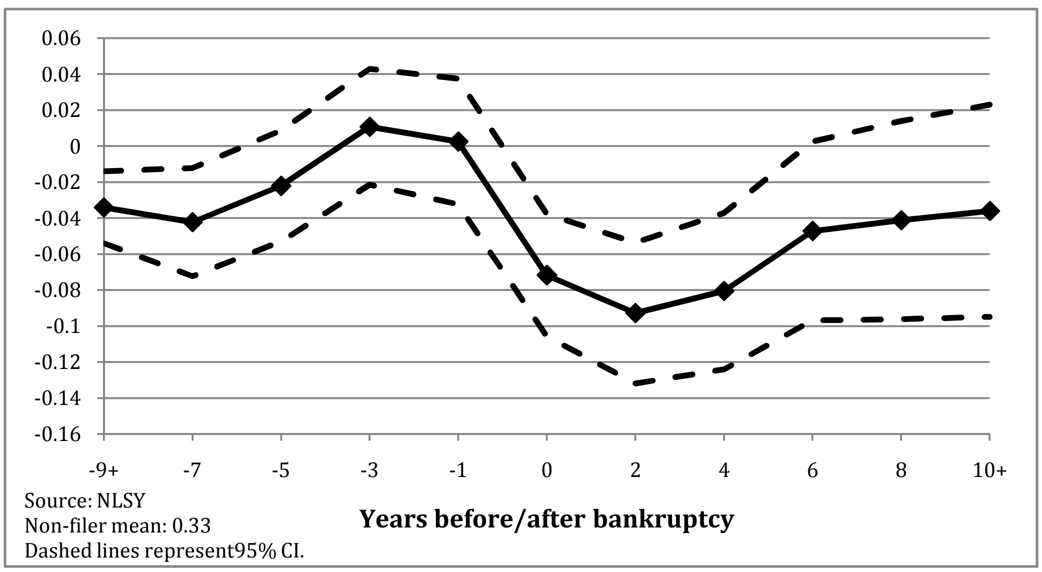

Figure 2.8: Odds ratios of bankruptcy filing, by relative time from UI shock - men

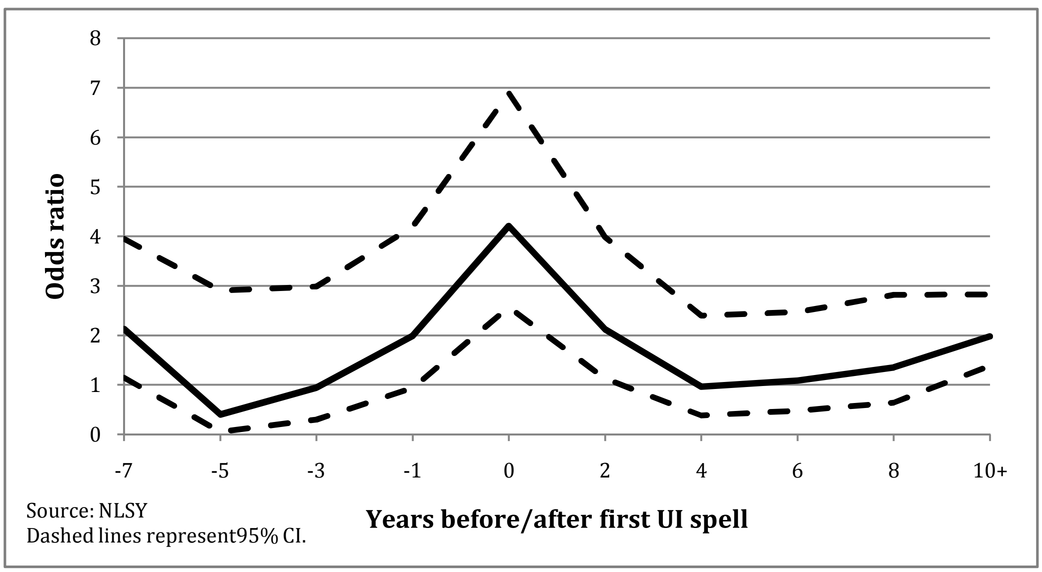

Figure 2.9: Odds ratios of bankruptcy filing, by relative time from UI shock - women

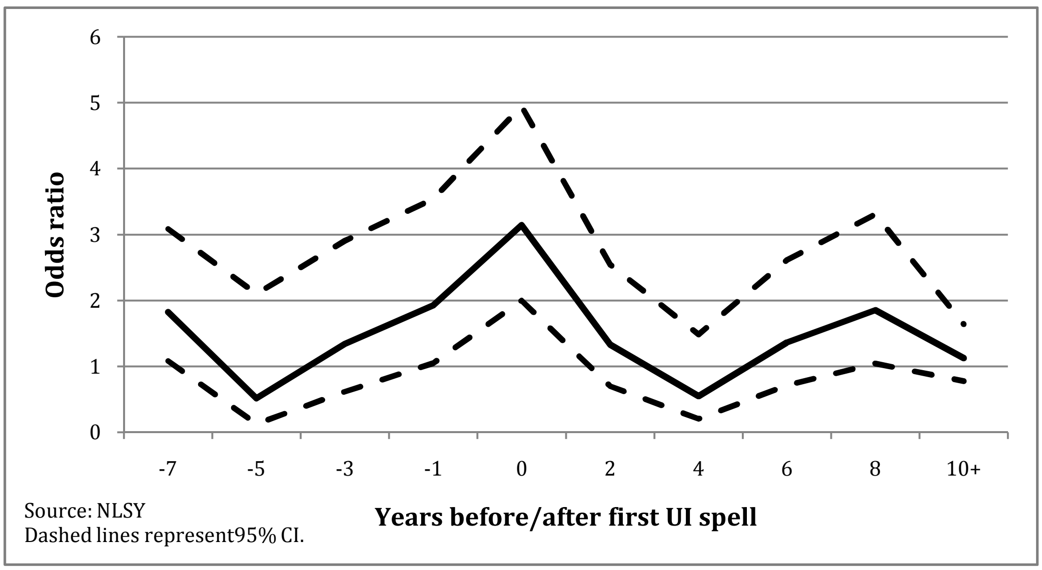

Figure 2.10: Odds ratios of bankruptcy filing, by relative time from divorce

Figure 2.11: Odds ratios of bankruptcy filing, by relative time from negative health shock

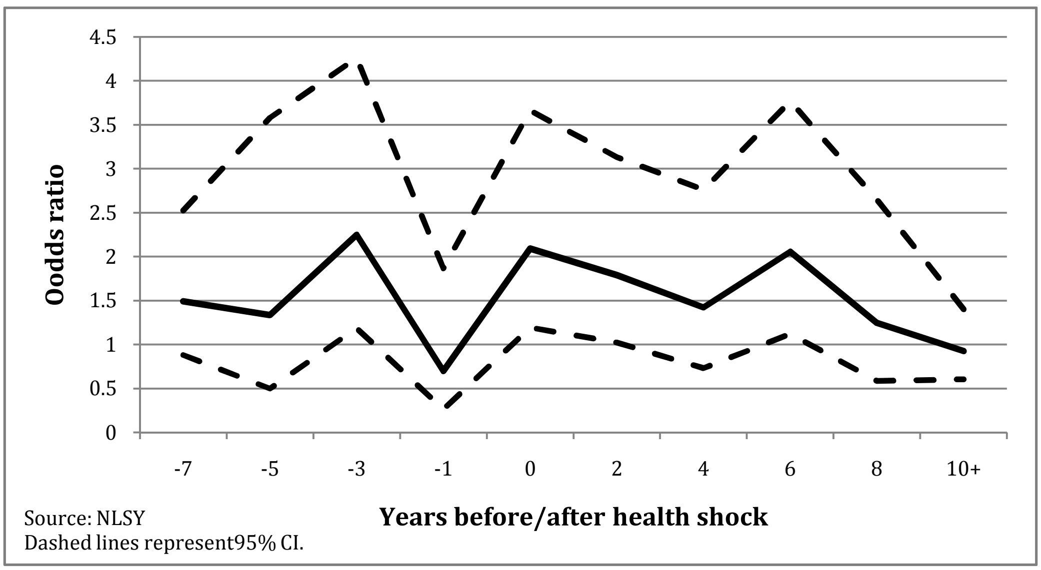

Table 2.1: Debts, assets, and the timing of bankruptcy

By type of asset or debt						
	Total Debt		"Other" Debts		Own Home?	
Relative Time Coefficients	Coef.	SE	Coef.	SE	Coef.	SE
9 or more years before	-9026	1226	765	230	-0.03	0.01
7-8 years before	-7800	1860	1240	436	-0.04	0.02
5-6 years before	-6136	1851	790	305	-0.02	0.02
3-4 years before	-556	1968	2388	420	0.01	0.02
1-2 years before	2480	2334	3856	675	0.00	0.02
year of bankruptcy + 1 year after	-9042	2133	2788	914	-0.07	0.02
2-3 years after	-15755	1768	-565	228	-0.09	0.02
4-5 years after	-12333	2166	-397	292	-0.08	0.02
6-7 years after	-9771	2451	-655	321	-0.05	0.03
8-9 years after	-7491	3131	956	637	-0.04	0.03
10 or more years after	-8368	3522	433	626	-0.04	0.03
Individuals	7661		7659		7661	
Observations	96354		87735		129198	
Non-filer mean	$36,961		$2,669		0.33	

Source: Author's calculations using NLSY79, 1979-2004.
Estimates derived from fixed-effects model, see text for details.
Standard errors clustered by individuals.
The results correspond to Figures 2.5, 2.6, and 2.7, respectively.
"Other" debts include credit card debt, medical and legal bills
and other outstanding debts.

Table 2.2: Summary statistics, NLSY in 2004

Variable	Obs	Mean	Std. Dev.
Less than High School	7661	8.2%	0.27
High School	7661	42.0%	0.49
Some College	7661	23.2%	0.42
College and Up	7661	26.6%	0.44
Age	7661	43.3	2.32
Mother's highest grade completed	7188	11.6	2.78
Father's highest grade completed	6534	11.8	3.60
Male	7661	50.9%	0.50
African-American	7661	14.3%	0.35
Ever filed for bankruptcy	7661	11.1%	0.31
Ever on UI - male	7661	40.5%	0.49
Ever on UI - female	7661	28.0%	0.45
Ever had health problem	7661	8.8%	0.28
Ever divorced	7661	45.4%	0.50

Source: NLSY79, 1979-2004. Observations weighted using sample weights.

Table 2.3: Odds ratios of the probability of filing for bankruptcy

By type of event	UI - Men		UI - Women		Divorce		Health Shock	
Relative Time Coefficients	Odds Ratio	z-stat	Odds Ratio	z-stat	Odds Ratio	z-stat	Odds Ratio	z-stat
7 or more years before	2.13	2.39	1.83	2.25	1.13	0.79	1.49	1.48
5-6 years before	0.40	-0.91	0.52	-0.92	1.34	1.28	1.33	0.57
3-4 years before	0.95	-0.09	1.34	0.74	1.24	1.03	2.25	2.49
1-2 years before	1.99	1.80	1.93	2.11	1.61	2.76	0.70	-0.72
year of event + 1 year after	4.21	5.71	3.14	4.95	2.09	5.04	2.09	2.58
2-3 years after	2.12	2.33	1.33	0.87	2.04	4.69	1.79	2.04
4-5 years after	0.96	-0.08	0.55	-1.18	1.81	3.7	1.42	1.04
6-7 years after	1.09	0.20	1.36	0.93	1.09	0.43	2.06	2.33
8-9 years after	1.35	0.80	1.85	2.09	1.17	0.79	1.24	0.56
10 or more years after	1.98	3.75	1.13	0.63	1.75	5.14	0.92	-0.39
Additional Controls	Yes		Yes		Yes		Yes	
Individuals	2661		2900		6541		7615	
Observations	35671		31435		125872		146621	
p-value for test of shock year = previous year	0.08		0.18		0.21		0.06	

Source: Author's calculations using NLSY, 1979-2004.
Estimates derived from logit model, see text for definition of spells.
Z-statistics calculated using standard errors clustered by individuals.
Additional controls are education (highest grade completed) and race.
Respondents' gender is included as a covariate in the last two specifications.

Table 2.4: County-level estimates of job loss and bankruptcy

Dependent variable = number of non-business bankruptcies in a county				
	(1)	(2)	(3)	(4)
Change in Total Jobs	-0.011***	-0.009***	-0.009***	-0.011***
	(0.0017)	(0.0015)	(0.0015)	(0.0023)
1st lag		-.013***	-0.011***	-0.006***
		(0.003)	(0.002)	(0.0018)
2nd lag			-0.011***	-0.006***
			(0.004)	(0.001)
3rd lag				-0.010
				(0.010)
Change in Manufacturing jobs	-0.032**	-0.030***	-0.032***	-0.036**
	(0.012)	(0.011)	(0.011)	(0.015)
1st lag		-0.037**	-0.032**	-0.011
		(0.016)	(0.013)	(0.008)
2nd lag			-0.036***	-0.017**
			(0.011)	(0.008)
3rd lag				-0.053***
				(0.018)
Change in Non-manufacturing jobs	-0.007***	-0.005**	-0.005**	-0.007***
	(0.002)	(0.003)	(0.003)	(0.002)
1st lag		-0.009***	-0.007***	-0.006**
		(0.001)	(0.002)	(0.003)
2nd lag			-0.007***	-0.004**
			(0.002)	(0.002)
3rd lag				0.003
				(0.008)
County fixed effects?	Yes	Yes	Yes	Yes
Year fixed effects?	Yes	Yes	Yes	Yes
Number of counties	3135	3135	3135	3135
Number of observations	78375	78375	78375	75240

Standard errors (in parentheses) clustered at the county level.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$
Note: The bottom panel presents 3 separate regressions (where the changes in manufacturing and non-manufacturing jobs enter the equation separately).
Sources: Bankruptcy data: AOUSC 1980-2004, Employment data: CBP 1977-2004.

Table 2.5: County-level estimates of job loss and bankruptcy with additional interactions

Dependent variable = number of non-business bankruptcies in a county

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Above Median in 1980:					
Interacted Covariate	High Unemp. Rate	Fraction w/HS diploma	Fraction w/college degree	% African-American	Median Age	County Population	Median Household Income
Change in Total Jobs	0.002	-0.001	-0.003	-0.019***	-0.014***	-0.008***	-0.005**
	(0.007)	(0.004)	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)
Change in Total Jobs	-0.048***	-0.010**	-0.008***	0.008**	0.006**	-0.003	-0.006**
x Covariate	(0.014)	(0.004)	(0.002)	(0.003)	(0.003)	(0.002)	(0.002)
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of counties	3135	3135	3135	3135	3135	3135	3135
Number of obs	78735	78375	78375	78375	78375	78375	78375

Standard errors (in parentheses) clustered at the county level. *** p<0.01, ** p<0.05, * p<0.1
In Column 1, high unemployment = 1 if national unemployment>6%, the median for 1977-2004.
In Columns 2-7, each covariate = 1 if the 1980 value is greater than the median for all 3135 counties.
Sources: Bankruptcy data: AOUSC 1980-2004, Employment data: CBP 1977-2004
Unemployment Rate: BLS, Demographic data: 1980 Census

CHAPTER III

Did Securitization Lead to Lax Screening? Evidence from Subprime Loans

3.1 Introduction

Securitization, converting illiquid assets into liquid securities, has grown tremendously in recent years, with the securitized universe of mortgage loans reaching $3.6 trillion in 2006. The option to sell loans to investors has transformed the traditional role of financial intermediaries in the mortgage market from "buying and holding" to "buying and selling." The perceived benefits of this financial innovation, such as improving risk sharing and reducing banks' cost of capital, are widely cited (e.g. Pennacchi 1988). However, delinquencies in the heavily securitized subprime housing market increased by 50% from 2005 to 2007, forcing many mortgage lenders out of business and setting off a wave of financial crises which spread worldwide. In light of the central role of the subprime mortgage market in the current crisis, critiques of the securitization process have gained increased prominence (Blinder 2007; Stiglitz 2007).

The rationale for concern over the "originate-to-distribute" model during the crisis derives from theories of financial intermediation. Delegating monitoring to a single lender avoids the problems of duplication, coordination failure, and free-rider problems associated with multiple lenders (Diamond 1984). However, in order for a lender to screen and monitor, it must be given appropriate incentives (Holmstrom and Tirole 1997) and this is provided by the illiquid loans on their balance sheet (Diamond and Rajan 2003). By creating distance between a loan's originator and the bearer of the loan's default risk, securitization may have potentially reduced lenders' incentives to carefully screen and monitor borrowers (Petersen and Rajan 2002). On the other hand, proponents of securitization argue reputation concerns, regulatory oversight, or sufficient balance sheet risk may have prevented moral hazard on the part of lenders. What the effects of existing securitization practices on screening were, thus, remains an empirical question.

This paper investigates the relationship between securitization and screening standards in the context of subprime mortgage loans. The challenge in making a causal claim is the difficulty in isolating differences in loan outcomes independent of contract and borrower characteristics. First, in any cross-section of loans, those which are securitized may differ on observable and unobservable risk characteristics from loans which are kept on the balance sheet (not securitized). Second, in a time-series framework, simply documenting a correlation between securitization rates and defaults may be insufficient. This inference relies on establishing the optimal level of defaults at any given point in time. Moreover, this approach ignores macroeconomic factors and policy initiatives which may be independent of lax screening and yet may induce compositional differences in mortgage borrowers over time. For instance, house price appreciation and the changing role of Government-Sponsored Enterprises (GSEs) in the subprime market may also have accelerated the trend toward originating mortgages to riskier borrowers in exchange for higher payments.

We overcome these challenges by exploiting a specific *rule of thumb* in the lending market which induces exogenous variation in the ease of securitization of a loan compared to a loan with similar characteristics. This *rule of thumb* is based on the summary measure of borrower credit quality known as the FICO score. Since the mid-1990s, the FICO score has become the credit indicator most widely used by lenders, rating agencies, and investors. Underwriting guidelines established by the GSEs, Fannie Mae and Freddie Mac, standardized purchases of lenders' mortgage loans. These guidelines cautioned against lending to risky borrowers, the most prominent rule of thumb being not lending to borrowers with FICO scores below 620 (Avery et al. 1996; Loesch 1996; Calomiris and Mason 1999; Capone 2002; Freddie Mac 2001, 2007).[1] While the GSEs actively securitized loans when the nascent subprime market was relatively small, since 2000 this role has shifted entirely to investment banks and hedge funds (the non-agency sector). We argue that persistent adherence to this ad-hoc cutoff by investors who purchase securitized pools from non-agencies generates a differential increase in the ease of securitization for loans. That is, loans made to borrowers which fall just above the 620 credit cutoff have a higher unconditional likelihood of being securitized and are therefore more liquid relative to loans below this cutoff.

To evaluate the effect of securitization on screening decisions, we examine the performance of loans originated by lenders around this threshold. As an example of our design, consider two

[1]We discuss the 620 rule of thumb in more detail in Section 3.3 and in reference to other cutoffs in the lending market in Section 3.4.7.

borrowers, one with a FICO score of 621 (620^{+}) while the other has a FICO score of 619 (620^{-}), who approach the lender for a loan. In order to evaluate the quality of the loan applicant, screening involves collecting both "hard" information, such as the credit score, and "soft" information, such as a measure of future income stability of the borrower. Hard information by definition is something that is easy to contract upon (and transmit), while the lender has to exert an unobservable effort to collect soft information (Stein 2002). We argue that the lender has a weaker incentive to base origination decisions on both hard and soft information, less carefully screening the borrower, at 620^{+} where there is a higher likelihood that this loan will be eventually securitized. In other words, because investors purchase securitized loans based on hard information, the cost of collecting soft information is internalized by lenders to a lesser extent when screening borrowers at 620^{+} than at 620^{-}. Therefore, by comparing the portfolio of loans on either side of the credit score threshold, we can assess whether differential access to securitization led to changes in the behavior of lenders who offered these loans to consumers with nearly identical risk profiles.

Using a sample of more than one million home purchase loans during the period 2001-2006, we empirically confirm that the number of loans securitized varies systematically around the 620 FICO cutoff. For loans with a potential for significant soft information – *low documentation* loans – we find that there are more than twice as many loans securitized above the credit threshold at 620^{+} vs. below the threshold at 620^{-}. Since the FICO score distribution in the population is smooth (constructed from a logistic function; see Figure 1), the underlying creditworthiness and demand for mortgage loans (at a given price) is the same for prospective buyers with a credit score of either 620^{-} or 620^{+}. Therefore, these differences in the number of loans confirm that the unconditional probability of securitization is higher above the FICO threshold, i.e., it is easier to securitize 620^{+} loans.

Strikingly, we find that while 620^{+} loans should be of slightly better credit quality than those at 620^{-}, low documentation loans that are originated above the credit threshold tend to default within two years of origination at a rate 10-25% higher than the mean default rate of 5% (which amounts to roughly a 0.5-1% increase in delinquencies). As this result is conditional on observable loan and borrower characteristics, the only remaining difference between the loans around the threshold is the increased ease of securitization. Therefore, the greater default probability of loans above the credit threshold must be due to a reduction in screening by lenders.

Since our results are conditional on securitization, we conduct additional analyses to address selection on the part of borrowers, lenders, or investors as explanations for the differences in the performance of loans around the credit threshold. First, we rule out borrower selection on observables, as the loan terms and borrower characteristics are smooth through the FICO score threshold. Next, selection of loans by investors is mitigated because the decisions of investors (Special Purpose Vehicles, SPVs) are based on the same (smooth through the threshold) loan and borrower variables as in our data (Kornfeld 2007).

Finally, strategic adverse selection on the part of lenders may also be a concern. However, lenders offer the entire pool of loans to investors, and, conditional on observables, SPVs largely follow a randomized selection rule to create bundles of loans out of these pools, suggesting securitized loans would look similar to those that remain on the balance sheet (Gorton and Souleles 2005; *Comptroller's Handbook* 1997). Furthermore, if at all present, this selection will tend to be more severe below the threshold, thereby biasing the results against us finding any screening effect. We also constrain our analysis to a subset of lenders who are not susceptible to strategic securitization of loans. The results for these lenders are qualitatively similar to the findings using the full sample, highlighting that screening is the driving force behind our results.

Could the 620 threshold be set by lenders as an optimal cutoff for screening that is unrelated to differential securitization? We investigate further using a natural experiment in the passage and subsequent repeal of anti-predatory laws in New Jersey (2002) and Georgia (2003) that varied the ease of securitization around the threshold. If lenders use 620 as an optimal cutoff for screening unrelated to securitization, we expect the passage of these laws to have no effect on the differential screening standards around the threshold. However, if these laws affected the differential ease of securitization around the threshold, our hypothesis would predict an impact on the screening standards. Our results confirm that the discontinuity in the number of loans around the threshold diminished during a period of strict enforcement of anti-predatory lending laws. In addition, there was a rapid return of a discontinuity after the law was revoked. Importantly, our performance results follow the same pattern, i.e., screening differentials attenuated only during the period of enforcement. Taken together, this evidence suggests that our results are indeed related to differential securitization at the credit threshold and that lenders did not follow the rule of thumb in all instances. Importantly, the natural experiment also suggests that prime-influenced selection is not at play.

Once we have confirmed that lenders are screening more rigorously at 620^{-} than 620^{+}, we assess whether borrowers were aware of the differential screening around the threshold. Although there is no difference in contract terms around the cutoff, borrowers may have an incentive to manipulate their credit scores in order to take advantage of differential screening around the threshold (consistent with our central claim). Aside from outright fraud, it is difficult to strategically manipulate one's FICO score in a targeted manner and any actions to improve one's score take relatively long periods of time, on the order of three to six months (Fair Isaac). Nonetheless, we investigate further using the same natural experiment evaluating the performance effects over a relatively short time horizon. The results reveal a rapid return of a discontinuity in loan performance around the 620 threshold which suggests that rather than manipulation, our results are largely driven by differential screening on the part of lenders.

As a test of the role of soft information on screening incentives of lenders, we investigate the *full documentation* loan lending market. These loans have potentially significant hard information because complete background information about the borrower's ability to repay is provided. In this market, we identify another credit cutoff, a FICO score of 600, based on the advice of the three credit repositories. We find that twice as many full documentation loans are securitized above the credit threshold at 600^+ vs. below the threshold at 600^-. Interestingly, however, we find no significant difference in default rates of full documentation loans originated around this credit threshold. This result suggests that despite a difference in ease of securitization around the threshold, differences in the returns to screening are attenuated due to the presence of more hard information. Our findings for full documentation loans suggest that the role of soft information is crucial to understanding what worked and what did not in the existing securitized subprime loan market. We discuss this issue in more detail in Section 3.6.

This paper connects several strands of literature. Our evidence sheds new light on the subprime housing crisis, as discussed in the contemporaneous work of Benmelech and Dlugosz (2008), Doms, Furlong, and Krainer (2007), Dell'Ariccia, Igan and Laeven (2008), Demyanyk and Van Hemert (2008), Gerardi, Shapiro and Willen (2007), Mayer, Piskorski, and Tchistyi (2008), Mian and Sufi (2008) and Rajan, Seru and Vig (2008).[2] This paper also speaks to the literature which discusses the benefits (Kashyap and Stein 2000 and Loutskina and Strahan 2007), and the costs (Parlour and Plantin 2007 and Morrison 2005) of securitization. In a related line of research, Drucker and Mayer (2008) document how underwriters exploit inside information to their advantage in secondary mortgage markets, while Gorton and Pennacchi (1995), Drucker and Puri (2007) and Sufi (2006) investigate how contract terms are structured to mitigate some of these agency conflicts.[3]

[2]For thorough summaries of the subprime mortgage crisis and the research which has sought to explain it, see Mayer and Pence (2008) and Mayer, Pence, and Sherlund (2008).

[3]Our paper also sheds light on the classic liquidity/incentives trade-off that is at the core of the financial contracting literature (see Coffee 1991, Diamond and Rajan 2003, Aghion et al. 2004, DeMarzo and Urosevic 2006).

The rest of the paper is organized as follows. Section 3.2 provides a brief overview of lending in the subprime market and describes the data and sample construction. Section 3.3 discusses the framework and empirical methodology used in the paper, while Sections 3.4 and 3.5 present the empirical results in the paper. Section 3.6 concludes.

3.2 Lending in Subprime Market

3.2.1 Background

Approximately 60% of outstanding U.S. mortgage debt is traded in mortgage-backed securities (MBS), making the U.S. secondary mortgage market the largest fixed-income market in the world (Chomsisengphet and Pennington-Cross 2006). The bulk of this securitized universe ($3.6 trillion outstanding as of January 2006) is comprised of agency pass-through pools – those issued by Freddie Mac, Fannie Mae and Ginnie Mae. The remainder, approximately, $2.1 trillion as of January 2006 has been securitized in non-agency securities. While the non-agency MBS market is relatively small as a percentage of all U.S. mortgage debt, it is nevertheless large on an absolute dollar basis. The two markets are separated based on the eligibility criteria of loans that the GSEs have established. Broadly, agency eligibility is established on the basis of loan size, credit score, and underwriting standards.

Unlike the agency market, the non-agency (referred to as "subprime" in the paper) market was not always this size. This market gained momentum in the mid- to late-1990s. Inside B&C Lending – a publication which covers subprime mortgage lending extensively – reports that total subprime lending (B&C originations) has grown from $65 billion in 1995 to $500 billion in 2005. Growth in mortgage-backed securities led to an increase in securitization rates (the ratio of the dollar-value of loans securitized divided by the dollar-value of loans originated) from less than 30 percent in 1995 to over 80 percent in 2006.

From the borrower's perspective, the primary distinguishing feature between prime and subprime loans is that the up-front and continuing costs are higher for subprime loans.[4] The subprime mortgage market actively prices loans based on the risk associated with the borrower. Specifically, the interest rate on the loan depends on credit scores, debt-to-income ratios and the documentation level of the borrower. In addition, the exact pricing may depend on loan-to-value ratios (the amount of equity of the borrower), the length of the loan, the flexibility of the interest rate (adjustable,

[4]Up-front costs include application fees, appraisal fees, and other fees associated with originating a mortgage. The continuing costs include mortgage insurance payments, principle and interest payments, late fees for delinquent payments, and fees levied by a locality (such as property taxes and special assessments).

fixed, or hybrid), the lien position, the property type and whether stipulations are made for any prepayment penalties.[5]

For investors who hold the eventual mortgage-backed security, credit risk in the agency sector is mitigated by an implicit or explicit government guarantee, but subprime securities have no such guarantee. Instead, credit enhancement for non-agency deals is in most cases provided internally by means of a deal structure which bundles loans into "tranches," or segments of the overall portfolio (Lucas, Goodman and Fabozzi 2006).

3.2.2 Data

Our primary data contain individual loan data leased from LoanPerformance. The database is the only source which provides a detailed perspective on the non-agency securities market. The data includes information on issuers, broker dealers/deal underwriters, servicers, master servicers, bond and trust administrators, trustees, and other third parties. As of December 2006, more than 8,000 home equity and nonprime loan pools (over 7,000 active) that include 16.5 million loans (more than seven million active) with over $1.6 trillion in outstanding balances were included. LoanPerformance estimates that as of 2006, the data covers over 90% of the subprime loans that are securitized.[6] The dataset includes all standard loan application variables such as the loan amount, term, LTV ratio, credit score, and interest rate type – *all* data elements that are disclosed and form the basis of contracts in non-agency securitized mortgage pools. We now describe some of these variables in more detail.

For our purpose, the most important piece of information about a particular loan is the creditworthiness of the borrower. The borrower's credit quality is captured by a summary measure called the FICO score. FICO scores are calculated using various measures of credit history, such as types of credit in use and amount of outstanding debt, but do *not* include any information about a borrower's income or assets (Fishelson-Holstein 2005). The software used to generate the score from individual credit reports is licensed by the Fair Isaac Corporation to the three major credit repositories – TransUnion, Experian, and Equifax. These repositories, in turn, sell FICO scores and credit reports to lenders and consumers. FICO scores provide a ranking of potential borrowers

[5] For example, the rate and underwriting matrix of Countrywide Home Loans Inc., a leading lender of prime and subprime loans, shows how the credit score of the borrower and the loan-to-value ratio are used to determine the rate at which different documentation-level loans are made (www.countrywide.com).

[6] Note that only loans that are securitized are reported in the LoanPerformance database. Communication with the database provider suggests that the roughly 10% of loans that are not reported are for privacy concerns from lenders. Importantly for our purpose, the exclusion is not based on any selection criteria that the vendor follows (e.g., loan characteristics or borrower characteristics). Moreover, based on estimates provided by LoanPerformance, the total number of non-agency loans securitized relative to all loans originated has increased from about 65% in early 2000 to over 92% since 2004.

by the probability of having some negative credit event in the next *two years*. Probabilities are rescaled into a range of 400-900, though nearly all scores are between 500 and 800, with a higher score implying a lower probability of a negative event. The negative credit events foreshadowed by the FICO score can be as small as one missed payment or as large as bankruptcy. Borrowers with lower scores are proportionally more likely to have all types of negative credit events than are borrowers with higher scores.

FICO scores have been found to be accurate even for low-income and minority populations (see Fair Isaac website www.myfico.com; also see Chomsisengphet and Pennington-Cross 2006). More importantly, the applicability of scores available at loan origination extends reliably up to two years. By design, FICO measures the probability of a negative credit event over a two-year horizon. Mortgage lenders, on the other hand, are interested in credit risk over a much longer period of time. The continued acceptance of FICO scores in automated underwriting systems indicates that there is a level of comfort with their value in determining lifetime default probability differences.[7] Keeping this as a backdrop, most of our tests of borrower default will examine the default rates up to 24 months from the time the loan is originated.

Borrower quality can also be gauged by the level of documentation collected by the lender when taking the loan. The documents collected provide historical and current information about the income and assets of the borrower. Documentation in the market (and reported in the database) is categorized as full, limited or no documentation. Borrowers with full documentation provide verification of income as well as assets. Borrowers with limited documentation provide no information about their income but do provide some information about their assets. “No-documentation” borrowers provide no information about income or assets, which is a very rare degree of screening lenience on the part of lenders. In our analysis, we combine limited and no-documentation borrowers and call them low documentation borrowers. Our results are unchanged if we remove the very small portion of loans which are no documentation.

Finally, there is also information about the property being financed by the borrower, and the purpose of the loan. Specifically, we have information on the type of mortgage loan (fixed rate, adjustable rate, balloon or hybrid), and the loan-to-value (LTV) ratio of the loan, which measures the amount of the loan expressed as a percentage of the value of the home. Typically loans are classified as either for purchase or refinance, though for convenience we focus exclusively on loans

[7] An econometric study by Freddie Mac researchers showed that the predictive power of FICO scores drops by about 25 percent once one moves to a three-to-five year performance window (Holloway, MacDonald and Straka 1993). FICO scores are still predictive, but do not contribute as much to the default rate probability equation after the first two years.

for home purchases.[8] Information about the geography where the dwelling is located (zipcode) is also available in the database.[9]

Most of the loans in our sample are for the owner-occupied single-family residences, townhouses, or condominiums (single unit loans account for more than 90% of the loans in our sample). Therefore, to ensure reasonable comparisons we restrict the loans in our sample to these groups. We also drop non-conventional properties, such as those that are FHA or VA insured or pledged properties, and also exclude buy down mortgages. We also exclude Alt-A loans, since the coverage for these loans in the database is limited. Only those loans with valid FICO scores are used in our sample. We conduct our analysis for the period January 2001 to December 2006, since the securitization market in the subprime market grew to a meaningful size post-2000 (Gramlich 2007).

3.3 Framework and Methodology

When a borrower approaches a lender for a mortgage loan, the lender asks the borrower to fill out a credit application. In addition, the lender obtains the borrower's credit report from the three credit bureaus. Part of the background information on the application and report could be considered "hard" information (e.g., the FICO score of the borrower), while the rest is "soft" (e.g., a measure of future income stability of the borrower, how many years of documentation were provided by the borrower, joint income status) in the sense that it is less easy to summarize on a legal contract. The lender expends effort to process the soft and hard information about the borrower and, based on this assessment, offers a menu of contracts to the borrower. Subsequently, borrowers decide to accept or decline the loan contract offered by the lender.

Once a loan contract has been accepted, the loan can be sold as part of a securitized pool to investors. Notably, only the hard information about the borrower (FICO score) and the contractual terms (e.g., LTV ratio, interest rate) are used by investors when buying these loans as a part of securitized pool.[10] In fact, the variables about the borrowers and the loan terms in the LoanPerformance database are identical to those used by investors and rating agencies to rate tranches of the securitized pool. Therefore, while lenders are compensated for the hard information about the borrower, the incentive for lenders to process soft information critically depends on whether they have to bear the risk of loans they originate (Gorton and Pennacchi 1995; Parlour and Plantin 2007;

[8]We find similar rules of thumb and default outcomes in the refinance market.

[9]See Keys et al. (2009) for a discussion of the interaction of securitization and variation in regulation, driven by the geography of loans and the type of lender.

[10]See Testimony of Warren Kornfeld, Managing Director of Moodys Investors Service before the subcommittee on Financial Institutions and Consumer Credit U.S. House of Representatives May 8, 2007.

Rajan et al. 2008). The central claim in this paper is that lenders are less likely to expend effort to process soft information as the ease of securitization increases.

We exploit a specific *rule of thumb* at the FICO score of 620 which makes securitization of loans more likely if a certain FICO score threshold is attained. Historically, this score was established as a minimum threshold in the mid-1990's by Fannie Mae and Freddie Mac in their guidelines on loan eligibility (Avery et al. 1996 and Capone 2002). Guidelines by Freddie Mac suggest that FICO scores below 620 are placed in the *Cautious Review Category*, and Freddie Mac considers a score below 620 "as a strong indication that the borrower's credit reputation is not acceptable." (Freddie Mac 2001, 2007).[11] This is also reflected in Fair Isaac's statement, "...those agencies [Fannie Mae and Freddie Mac], which buy mortgages from banks and resell them to investors, have indicated to lenders that any consumer with a FICO score above 620 is good, while consumers below 620 should result in further inquiry from the lender....". While the GSEs actively securitized loans when the nascent subprime market was relatively small, this role shifted entirely to investment banks and hedge funds (the non-agency sector) in recent times (Gramlich, 2007).

We argue that adherence to this cutoff by subprime MBS investors, following the advice of GSEs, generates an increase in demand for securitized loans which are just above the credit cutoff relative to loans below this cutoff. There is widespread evidence that is consistent with 620 being a rule of thumb in the securitized subprime lending market. For instance, rating agencies (Fitch and Standard and Poor's) used this cutoff to determine default probabilities of loans when rating mortgage backed securities with subprime collateral (Loesch 1996; Temkin, Johnson and Levy 2002). Similarly, Calomiris and Mason (1999) survey the high risk mortgage loan market and find 620 as a rule of thumb for subprime loans. We also confirmed this view by conducting a survey of origination matrices used by several of the top 50 originators in the subprime market (a list obtained from *Inside B&C Lending*; these lenders amount to about 70% of loan volume). The credit threshold of 620 was used by nearly all the lenders.

Since investors purchase securitized loans based on hard information, our assertion is that the cost of collecting soft information are internalized by lenders to a greater extent when screening borrowers at 620^- than at 620^+. There is widespread anecdotal evidence that lenders in the subprime market review both soft and hard information more carefully for borrowers with credit scores below 620. For instance, the website of Advantage Mortgage, a subprime securitized loan originator, claims that "...all loans with credit scores below 620 require a second level review....There

[11]These guidelines appeared at least as far back as 1995 in a letter by the Executive Vice President of Freddie Mac (Michael K. Stamper) to the CEOs and Credit Officers of all Freddie Mac Sellers and Servicers (see internet appendix Exhibit 1).

are no exceptions, regardless of the strengths of the collateral or capacity components of the loan."[12] By focusing on the lender as a unit of observation we attempt to learn about the differential impact ease of securitization had on behavior of lenders around the cutoff.

To begin with, our tests empirically identify a statistical discontinuity in the distribution of loans securitized around the credit threshold of 620. In order to do so, we show that the number of loans securitized dramatically increases when we move along the FICO distribution from 620^- to 620^+. We argue that this is equivalent to showing that the unconditional probability of securitization increases as one moves from 620^- to 620^+. To see this, denote $N_s^{620^+}$ and $N_s^{620^-}$ as the number of loans securitized at 620^+ and 620^- respectively. Showing that $N_s^{620^+} > N_s^{620^-}$ is equivalent to showing $\frac{N_s^{620^+}}{N_p} > \frac{N_s^{620^-}}{N_p}$, where N_p is the number of prospective borrowers at 620^+ or 620^-. If we assume that the number of prospective borrowers at 620^+ or 620^- are similar, i.e., $N_p^{620^+} \approx N_p^{620^+} = N_p$ (a reasonable assumption as discussed below), then the unconditional probability of securitization is higher at 620^+. We refer to the difference in these unconditional probabilities as the differential *ease of securitization* around the threshold. Notably, our assertion of differential screening by lenders does not rely on knowledge of the proportion of prospective borrowers that applied, were rejected, or were held on the lenders' balance sheet. We simply require that lenders' are aware that a prospective borrower at 620^+ has a higher likelihood of eventual securitization.

We measure the extent of the jump by using techniques which are commonly used in the literature on regression discontinuity (e.g., see DiNardo and Lee 2004; Card et al. 2007). Specifically, we collapse the data on each FICO score (500-800) i, and estimate equations of the form:

$$Y_i = \alpha + \beta T_i + \theta f(FICO(i)) + \delta T_i * f(FICO(i)) + \epsilon_i \tag{3.1}$$

where Y_i is the number of loans at FICO score i, T_i is an indicator which takes a value of 1 at FICO $\geq$ 620 and a value of 0 if FICO $<$ 620 and ϵ_i is a mean-zero error term. $f(FICO)$ and $T * f(FICO)$ are flexible seventh-order polynomials, with the goal of these functions being to fit the smoothed curves on either side of the cutoff as closely to the data presented in the figures as possible.[13] $f(FICO)$ is estimated from 620^- to the left, and $T * f(FICO)$ is estimated from 620^+ to the right. The magnitude of the discontinuity, β, is estimated by the difference in these two smoothed functions evaluated at the cutoff. The data are *re-centered* such that $FICO = 620$

[12] This position for loans below 620 is reflected in lending guidelines of numerous other subprime lenders.

[13] We have also estimated these functions of the FICO score using 3rd order and 5th order polynomials in FICO, as well as relaxing parametric assumptions and estimating using local linear regression. The estimates throughout are not sensitive to the specification of these functions. In Section 3.4, we also examine the size and power of the test using the seventh-order polynomial specification following the approach of Card et al. (2007).

corresponds to "0," thus at the cutoff the polynomials are evaluated at 0 and drop out of the calculation, which allows β to be interpreted as the magnitude of the discontinuity at the FICO threshold. This coefficient should be interpreted locally in the immediate vicinity of the credit score threshold.

After documenting a large jump at the ad-hoc credit thresholds, we focus on the performance of the loans around these thresholds. We evaluate the performance of the loans by examining the default probability of loans – i.e., whether or not the loan defaulted t months after it was originated. If lenders screen similarly for the loan of credit quality 620^+ and the loan of 620^- credit quality, there should not be any discernible differences in default rates of these loans. Our maintained claim is that any differences in default rates on either side of the cutoff, after controlling for hard information, should be only due to the impact that securitization has on lenders' screening standards.

This claim relies on several identification assumptions. First, as we approach the cutoff from either side, any differences in the characteristics of prospective borrowers are assumed to be random. This implies that the underlying creditworthiness and the demand for mortgage loans (at a given price) is the same for prospective buyers with a credit score of 620^- or 620^+. This seems reasonable as it amounts to saying that the calculation Fair Isaac performs (using a logistic function) to generate credit scores has a random error component around any specific score. Figure 3.1 shows the FICO distribution in the U.S. population in 2004. This data is from an anonymous credit which assures us that the data exhibits similar patterns during the other years of our sample. Note that the FICO distribution across the population is smooth, so the number of prospective borrowers around a given credit score is similar (in the example above, $N_p^{620^+} \approx N_p^{620^+} = N_p$).

Second, we assume that screening is costly for the lender. The collection of information – hard systematic data (e.g., FICO score) as well as soft information (e.g., joint income status) about the creditworthiness of the borrower – requires time and effort by loan officers. If lenders did not have to expend resources to collect information, it would be difficult to argue that the differences in performance we estimate are a result of ease of securitization around the credit threshold affecting banks incentives to screen and monitor. Again, this seems to be a reasonable assumption (see Gorton and Pennacchi 1995).

Note that our discussion thus far has assumed that there is no explicit manipulation of FICO scores by the lenders or borrowers. However, the borrower may have incentives to do so if loan contracts or screening differ around the threshold. Our analysis in Section 3.4.6 focuses on a natural experiment and shows that the effects of securitization on performance are not being driven by strategic manipulation.

3.4 Main Empirical Results

3.4.1 Descriptive Statistics

As noted earlier, the non-agency market differs from the agency market on three dimensions: FICO scores, loan-to-value ratios and the amount of documentation asked of the borrower. We next look at the descriptive statistics of our sample with special emphasis on these dimensions. Our analysis uses more than one million loans across the period 2001 to 2006. As mentioned earlier, the non-agency securitization market has grown dramatically since 2000, which is apparent in Panel A of Table 3.1, which shows the number of subprime loans securitized across years. These patterns are similar to those described in Demyanyk and Van Hemert (2007) and Gramlich (2007). The market has witnessed an increase in the number of loans with reduced hard information in the form of limited or no documentation. Note that while limited documentation provides no information about income but does provide some information about assets, a no-documentation loan provides information about neither income nor assets. In our analysis we combine both types of limited-documentation loans and denote them as *low* documentation loans. The full documentation market grew by 445% from 2001 to 2005, while the number of low documentation loans grew by 972%.

We find similar trends for loan-to-value ratios and FICO scores in the two documentation groups. LTV ratios have gone up over time, as borrowers have put in less and less equity into their homes when financing loans. This increase is consistent with a better appetite of market participants to absorb risk. In fact, this is often considered the bright side of securitization – borrowers are able to borrow at better credit terms since risk is being borne by investors who can bear more risk than individual banks. Panel A also shows that average FICO scores of individuals who access the subprime market has been increasing over time. The mean FICO score among low documentation borrowers increased from 630 in 2001 to 655 in 2006. This increase in average FICO scores is consistent with the rule of thumb leading to a larger expansion of the market above the 620 threshold. Average LTV ratios are lower and FICO scores higher for low documentation as compared to the full documentation sample. This possibly reflects the additional uncertainty lenders have about the quality of low documentation borrowers.

Panel B compares the low and full documentation segments of the subprime market on a number of the explanatory variables used in the analysis. Low documentation loans are on average larger and given to borrowers with higher credit scores than loans where full information on income and assets are provided. However, the two groups of loans have similar contract terms such as interest rate, loan-to-value, prepayment penalties, and whether the interest rate is adjustable or not. Our analysis below focuses first on the low documentation segment of the market, and we explore the full documentation market in Section 3.5.

3.4.2 Establishing the Rule of Thumb

We first present results that show that large differences exist in the number of low documentation loans that are securitized around the credit threshold we described earlier. We then examine whether this jump in securitization has any consequences on the subsequent performance of the loans above and below this credit threshold.

As mentioned in Section 3.3, the rule of thumb in the lending market impacts the ease of securitization around the credit score of 620. We therefore expect to see a substantial increase in the number of loans just above this credit threshold as compared to number of loans just below this threshold. In order to examine this, we start by plotting the number of loans at each FICO score in the two documentation categories around the credit cutoff of 620 across years starting with 2001 and ending in 2006. As can be seen from Figure 3.2, there is a marked increase in number of low documentation loans around the credit score of 620 – that is, at 620^+ relative to number of loans at 620^-. We do not find any such jump for full documentation loans at FICO of 620.[14] Given this evidence, we focus on the 620 credit threshold for low documentation loans.

From Figure 3.2, it is clear that the number of loans see roughly a 100% jump in 2004 for low documentation loans around the credit score of 620 – i.e., there are twice as many loans securitized at 620^+ as compared to loans securitized at 620^-. Clearly, this is consistent with the hypothesis that the ease of securitization is higher at 620^+ than at scores just below this credit cutoff.

To estimate the jumps in the number of loans, we use the methods described above in Section 3.3 using the specification provided in equation (3.1). As reported in Table 3.2, we find that low documentation loans see a dramatic increase above the credit threshold of 620. In particular, the coefficient estimate (β) is significant at the 1% level and is on average around 110% (from 73 to 193%) higher for 620^+ as compared to 620^- for loans during the sample period. For instance, in 2001, the estimated discontinuity in Panel A is 85. The mean average number of low documentation loans at a FICO score for 2001 is 117. The ratio is around 73%. These jumps are plainly visible from the yearly graphs in Figure 3.2.

[14] We will elaborate more on full documentation loans in Section 3.5.

In addition, we conduct permutation tests (or "randomization" tests), where we varied the location of the discontinuity (T_i) across the range of all possible FICO scores and re-estimated equation (1). The test treats every value of the FICO distribution as a potential discontinuity, and estimates the magnitude of the observed discontinuity at each point, forming a counterfactual distribution of discontinuity estimates. This is equivalent to a bootstrapping procedure which varies the cutoff but does not re-sample the order of the points in the distribution (Johnston and DiNardo 1996). We then compare the value of the estimated discontinuity at 620 to the counterfactual distribution and construct a test statistic based on the asymptotic normality of the counterfactual distribution and report the p-value from this test. The null hypothesis is that the estimated discontinuity at a FICO score of 620 is that of the mean of the 300 possible discontinuities.[15]

The precision of the permutation test is limited by the number of observations used at each FICO score. As a result, regressions which pool across years provide the greatest power for statistical testing. While constructing the counterfactuals, we therefore use pooled specifications with year fixed effects removed to account for differences in vintage. The result of this test is shown in Table 3.2 and shows that the estimate at 620 for low documentation loans is a strong outlier relative to the estimated jumps at other locations in the distribution. The estimated discontinuity when the years are pooled together is 780 loans with a permutation test p-value of 0.003. In summary, if the underlying creditworthiness and the demand for mortgage loans is the same for potential buyers with a credit score of 620^- or 620^+, this result confirms that it is easier to securitize loans above the FICO threshold.

3.4.3 Contract Terms and Borrower Demographics

Before examining the subsequent performance of loans around the credit threshold, we first assess if there are any differences in hard information – either in contract terms or other borrower characteristics – around this threshold. The endogeneity of contractual terms based on the riskiness of borrowers may lead to different contracts and hence, different types of borrowers obtaining loans around the threshold in a systematic way. Though we control for the possible contract differences when we evaluate the performance of loans, it is insightful to examine whether borrower and contract terms also systematically differ around the credit threshold.

We start by examining the contract terms – LTV ratio and interest rates – around the credit

[15]In unreported tests, we also conduct a falsification simulation exercise following Card et al. (2007). In particular, we apply our specification to data generated by a continuous process. We reject the null hypothesis of no effect (using a 2-sided 5% test) in 6.0% of the simulations indicating that the size of our test is a reasonable. A similar test with data generated by a discontinuous process suggests that the power of our test is also reasonable. We reject the null of no effect about 92% of the times (in a 2-sided 5% test) in this case.

threshold. Figures 3.3 and 3.4 show the distribution of interest rates and LTV ratios offered on low documentation loans across the FICO spectrum. As is apparent, we find these loan terms to be very similar – i.e., we find no differences in contract terms for low documentation loans above and below the 620 credit score. We test this formally using an approach equivalent to equation (3.1), replacing the dependent variable Y_i in the regression framework with contract terms (loan-to-value ratios and interest rates) and present the results in the appendix (Table 3.A1). Our results suggest that there is no difference in loan terms around the credit threshold. For instance, for low-documentation loans originated in 2006, the average loan-to-value ratio across the collapsed FICO spectrum is 85%, whereas our estimated discontinuity is only -1.05%, a 1.2% difference. Similarly for the interest rate, for low-documentation loans originated in 2005, the average interest rate is 8.2%, and the difference on either side of the credit score cutoff is only about -0.091%, a 1% difference. Permutation tests reported in Appendix Table 3.A4 confirm that these differences are not outliers relative to the estimated jumps at other locations in the distribution.

Additional contract terms, such as the presence of a prepayment penalty, or whether the not the loan is ARM, FRM or interest only/balloon are also similar around the 620 threshold (results not shown). In addition, if loans have second liens, then a combined LTV (CLTV) ratio is calculated. We find no difference in the CLTV ratios around the threshold for those borrowers with more than one lien on the home. Finally, low documentation loans often do not require that borrowers provide information about their income, so there is only a subset of our sample which provides a debt-to-income (DTI) ratio for the borrowers. Among this subsample, there is no difference in DTI around the 620 threshold in low documentation loans. For brevity, we report only the permutation tests for these contract terms in Appendix Table 3.A4.

Next, we examine whether the characteristics of borrowers differ systematically around the credit threshold. In order to evaluate this, we look at the distribution of the population of borrowers across the FICO spectrum for low documentation loans. The data on borrower demographics comes from Census 2000 and is at the zip code level. As can be seen from Figure 3.5, median household income of the zip codes of borrowers around the credit thresholds look very similar for low documentation loans. We plotted similar distributions for average percent minorities residing in the zip code, and average house value in the zip code across the FICO spectrum (unreported) and again find no differences around the credit threshold.[16]

We use the same specification as equation (3.1), this time with the borrower demographic

[16] Of course, since the census data is at the zip code level, we are to some extent smoothing our distributions. We note, however, that when we conduct our analysis on differences in number of loans (from Section 3.4.2), aggregated at the zip code level, we still find jumps around the credit threshold within each individual zip code.

characteristics as dependent variables and present the results formally in the appendix (Table 3.A2). Consistent with the patterns in the figures, permutation tests (unreported) reveal no differences in borrower demographic characteristics around the credit score threshold. Overall, our results indicate that observable characteristics of loans and borrowers are not different around the credit threshold.

3.4.4 Performance of Loans

We now focus on the performance of the loans that are originated close to the credit score threshold. Note that our analysis in Section 3.4.3 suggests that there is no difference in terms of observable hard information about contract terms or about borrower demographic characteristics around the credit score thresholds. Nevertheless, we will control for these differences when evaluating the subsequent performance of loans in our logit regressions. If there is any remaining difference in the performance of the loans above and below the credit threshold, it can be attributed to differences in unobservable soft information about the loans.

We estimate the differences in default rates on either side of the cutoff using the same framework as equation (3.1), using the dollar-weighted fraction of loans defaulted within 10-15 months of origination as the dependent variable, Y_i. This fraction is calculated as the dollar amount of unpaid loans in default divided by the total dollar amount originated in the same cohort. We classify a loan as under default if any of the conditions is true: (a) payments on the loan are 60+ days late as defined by Office of Thrift Supervision; (b) the loan is in foreclosure; or (c) the loan is real estate owned (REO), i.e. the bank has re-taken possession of the home.[17]

We collapse the data into one-point FICO bins and estimate seventh-order polynomials on either side of the threshold for each year. By estimating the magnitude of β in each year separately, we ensure that no one cohort (or vintage) of loans is driving our results. As shown in Figures **??** to 3.17, the low documentation loans exhibit discontinuities in default rates at the FICO score of 620. A year by year estimate is presented in Panel A of Table 3.3. Contrary to what one might expect, around the credit threshold, we find that loans of higher credit scores actually default *more often* than lower credit loans in the post-2000 period. In particular for loans originated in 2005, the estimate of β is .023 (t-stat=2.10), and the mean delinquency rate is .078, suggesting a 29% increase in defaults to the right of the credit score cutoff. Similarly, in 2006, the estimated size of

[17]While there are two different definitions of delinquency used in the industry (Mortgage Banker's Association (MBA) definition and Office of Thrift Supervision (OTS) definition), we have followed the more stringent OTS definition. While MBA starts counting days a loan has been delinquent from the time a payment is missed, OTS counts days a loan is delinquent one month *after* the first payment is missed.

the jump is .044 (t-stat=2.68), the mean delinquency rate for all FICO bins is .155, which is again a 29% increase in defaults around the FICO score threshold.

Panel B of Table 3.3 presents results of permutation tests, estimated on the residuals obtained after pooling delinquency rates across years and removing year effects. Besides the 60+ late delinquency definition used in Panel A, we also classify a loan in default if it is 90+ late in payments and if it is in foreclosure or REO. Our results yield similar, if not stronger, results. Compared to 620^- loans, 620^+ loans are on average 2.8% more likely to be in arrears of 90+ days, and 2.5% more likely to be in foreclosure or REO. Permutation tests p-values confirm that the jump in defaults at 620 using all the definitions of default are extreme outliers to the rest of the delinquency distribution. For instance, with default defined as foreclosure/REO, the p value for the discontinuity at 620 is 0.004. That we find similar results using different default definitions is consistent with high levels of rollover, whereby loans which are delinquent continue to reach deeper levels of delinquency. As shown in internet appendix Table 1, more than 80% of loans which are 60 days delinquent reach 90+ days delinquent within a year, and 66% of loans which are 90 days delinquent reach foreclosure twelve months after in the low documentation market.

While previous default definitions were dollar-weighted, we also use the raw number of loans in default to estimate the magnitude of the discontinuity in loan performance around the FICO threshold. The unweighted results with 60+ delinquency are also presented Panel B, and continue to exhibit a pattern of higher credit scores leading to higher default rates around the 620 threshold. In fact, the results are statistically stronger than the 60+ weighted results, with a permutation test p-value based on the pooled estimates of 0.004 and the discontinuity estimate being significant in all the years (unreported; see internet appendix Figure 4).

To show how delinquency rates evolve over the age of the loan, in Figure 3.18 we plot the delinquency rates of 620^+ and 620^- for low documentation loans (dollar weighted) by loan age. As discussed earlier, we restrict our analysis to about two years after the loan has been originated. As can be seen from the figure, the differences in the delinquency rates are stark. The differences begin around four months after the loans have been originated and persist up to two years. Differences in default rates also seem quite large in terms of magnitudes. Those with a credit score of 620^- are about 20% less likely to default after a year as compared to loans of credit score 620^+.[18]

[18] Note that Figure 3.18 does not plot cumulative delinquencies. As loans are paid out, say after a foreclosure, the unpaid balance for these loans falls relative to the time when they entered into a 60+ state. This explains the dip in delinquencies in the figure after about 20 months. Our results are similar if we plot cumulative delinquencies, or delinquencies which are calculated using the unweighted number of loans. Also note that the fact that we find no delinquencies early on in the duration of the loan is not surprising, given that originators are required to take back loans on their books if the loans default within three months.

An alternative methodology is to measure the performance of each unweighted loan by tracking whether or not it became delinquent and estimate logit regressions of the following form:

$$Y_{ikt} = \Phi\left(\alpha + \beta T_{it} + \gamma_1 X_{ikt} + +\delta_1 T_{it} * X_{ikt} + \mu_t + \epsilon_{ikt} \right). \tag{3.2}$$

This logistic approach complements the regression discontinuity framework, as we restrict the sample to the 10 FICO points in the immediate vicinity of 620 in order to maintain the same local interpretation of the RD results. Moreover, we are also able to directly control for the possibly endogenous loan terms around the threshold. The dependent variable is an indicator variable (*Delinquency*) for loan i originated in year t that takes a value of 1 if the loan is classified as under default in month k after origination as defined above. We drop the loan from the regression once it is paid out after reaching the REO state. T takes the value 1 if FICO is between 620 and 624, and 0 if it is between 615 and 619 for low documentation loans, thus restricting the analysis to the immediate vicinity of the cutoffs. Controls include FICO scores, the interest rate on the loan, loan-to-value ratio, borrower demographic variables, as well as interaction of these variables with T. We also include a dummy variable for the type of loan (adjustable or fixed rate mortgage). We control for the possible nonlinear effect of age of the loan on defaults by including three dummy variables – that take a value of 1 if the month since origination is between 0-10, 11-20 and more than 20 months respectively. Year of origination fixed effects are included in the estimation and standard errors are clustered at the loan level to account for multiple loan delinquency observations in the data.

As can be seen from the logit coefficients in Panel C of Table 3.3, results from this regression are qualitatively similar to those reported in the figures. In particular, we find that β is positive when we estimate the regressions for low documentation loans. The economic magnitudes are similar to those in the figures as well. For instance, keeping all other variables at their mean level, low documentation loans with credit score of 620^- are about 10-25% less likely to default after a year as compared to low documentation loans of credit score 620^+. These are large magnitudes – for instance, note that the mean delinquency rate for low documentation loans is around 4.45%; the economic magnitude of the effects in Column (2) suggest that the difference in the absolute delinquency rate between loans around the credit threshold is around 0.5-1% for low documentation loans.[19]

[19]Our logistic specification is equivalent to a hazard model if we drop loans as soon as they hit the first indicator of delinquency (60 days in default) and include a full set of duration dummies. Doing so does not change the nature of our results.

To account for the possibility that lax screening might be correlated across different loans within the same vintage, we cluster the loans for each vintage and report the results in Columns (3) and (4). Note that the RD regressions (Panel A) estimated separately by year also alleviates concerns about correlated errors across different loans with the same vintage.

In the mortgage market, the other way for loans to leave the pool is to be repaid in full through refinancing or outright purchase, known as prepayment. This prepayment risk decreases the return to investing in mortgage-backed securities in a similar manner to default risk (see, e.g. Gerardi, Shapiro, and Willen 2007 and Mayer, Piskorski, and Tchistyi 2008). To assess whether there are any differences in actual prepayments around the 620 threshold, we plot the prepayment seasoning curve for all years 2001-2006 in Figure 3.19. As can be observed, prepayments of 620^{+} and 620^{-} borrowers in the low documentation market are similar (also see permutation test in Appendix Table 3.A4). Nevertheless, to formally account for prepayment rates, we also estimate a competing risk model using both prepayment and default as means for exiting the sample. We use the Cox-proportional hazard model based on the econometric specification following Deng, Quigley and Van Order (2000). In unreported tests (internet appendix Table 6), we find results that are similar to our logistic specification.

Finally, the reported specification uses five-point bins of FICO scores around the threshold, but the results are similar (though less precise) if we restrict the bins to fewer FICO scores on either side of 620 (internet appendix Table 2). This issue is also fully addressed by the regression discontinuity results reported in Panels A and B, which use individual FICO score bins as the units of observation. In sum, we find that even after controlling for all observable characteristics of the loan contracts or borrowers, loans made to borrowers with *higher* FICO scores perform *worse* around the credit threshold.

3.4.5 Selection Concerns

Since our results are conditional on securitization, we conduct additional analyses to address selection explanations on account of borrowers, investors and lenders for the differences in the performance of loans around the credit threshold. First, contract terms offered to borrowers above the credit threshold might differ from those below the threshold and attract a riskier pool of borrowers. If this were the case, it would not be surprising if the loans above the credit threshold perform worse than those below it. As shown in Section 3.4.3, loan terms are smooth through the FICO score threshold. We also investigate the loan terms in more detail than in Section 3.4.3 by examining the distribution of interest rates and loan-to-value ratios of contracts offered around 620 for low documentation loans.

Figure **??** depicts the Epanechnikov kernel density of the interest rate on low documentation loans in the year 2004 for two FICO groups – 620^- (615-619) and 620^+ (620-624). The distribution of interest rates observed in the two groups lie directly on top of one another. A Kolmogorov-Smirnov test for equality of distribution functions cannot be rejected at the 1% level. Similarly, Figure **??** depicts density of LTV ratios on low documentation loans in the year 2004 for 620^- and 620^+ groups. Again, a Kolmogorov-Smirnov test for equality of distribution functions cannot be rejected at the 1% level. The fact that we find that the borrowers characteristics are similar around the threshold (Section 3.4.3) also confirms that selection based on observables is unlikely to explain our results.[20]

Second, there might be concerns about selection of loans by investors. In particular, our results could be explained if investors could potentially cherry pick better loans below the threshold. The loan and borrower variables in our data are identical to the data upon which investors base their decisions (Kornfeld 2007). Furthermore, as shown in Section 3.4.3, these variables are smooth through the threshold, mitigating any concerns on selection by investors.[21]

Finally, strategic adverse selection on the part of lenders may also be a concern. Lenders could for instance keep loans of better quality on their balance sheet and offer only loans of worse quality to the investors. This concern is mitigated for several reasons. First, the securitization guidelines suggest that lenders offer the entire pool of loans to investors and that conditional on observables, SPVs largely follow a randomized selection rule to create bundles of loans. This suggests that securitized loans would look similar to those that remain on the balance sheet (Gorton and Souleles 2005; *Comptroller's Handbook* 1997).[22] In addition, this selection if at all present will tend to be

[20]The equality of interest rate distributions also rules out differences in the expected cost of capital around the threshold as an alternative explanation. For instance, lenders could originate riskier loans above the threshold only because the expected cost of capital is lower due to easier securitization. However, in a competitive market, the interest rates charged for these loans should reflect the riskiness of the borrowers. In that case, as mean interest rates above and below the threshold are the same (Section 3.4.3), lenders must have added riskier borrowers above the threshold – resulting in a more dispersed interest rate distribution above the threshold. Our analysis in Figure **??** shows that this is not the case.

[21]An argument might also be made that banks screen similarly around the credit threshold but are able to sell portfolio of loans above and below the threshold to investors with different risk tolerance. If this were the case, it could potentially explain our results in Section 3.4.4. This does not seem likely. Since all the loans in our sample are securitized, our results on performance on loans around the credit threshold are *conditional* on securitization. Moreover, securitized loans are sold to investors in pools which contains a mix of loans from the entire credit score spectrum. As a result, it is difficult to argue that loans of 620^- are purchased by different investors as compared to loans of 620^+.

[22]We confirmed this fact by examining a subset of loans held on the lenders' balance sheets. The alternative dataset covers the top 10 servicers in the subprime market (more than 60% of the market) with details on performance and loan terms of loans that are securitized or held on the lenders' balance sheet. We find no differences in the performance of loans that are securitized relative to those kept by lenders, around the 620 threshold. Results of this analysis are available upon request.

more severe below the credit threshold, thereby biasing us against finding any effect of screening on performance.

We conduct an additional test which also suggests that our results are not driven by selection on the part of lenders. While banks may screen and then strategically hold loans on their balance sheets, independent lenders do not keep a portfolio of loans on their books. These lenders finance their operations entirely out of short-term warehouse lines of credit, have limited equity capital, and no deposit base to absorb losses on loans that they originate (Gramlich 2007). Consequently, they have limited motive for strategically choosing which loans to sell to investors. However, because loans below the threshold are more difficult to securitize and thus are less liquid, these independent lenders still have strong incentives to differentially screen these loans to avoid losses. We focus on these lenders to isolate the effects of screening in our results on defaults (Section 3.4.4).

To test this, we classify the lenders into two categories – banks (banks, subsidiaries, thrifts) and independents – and conduct the performance results only for sample of loans originated by independent lenders. It is difficult to identify all the lenders in the database since many of the lender names are abbreviated. In order to ensure that we are able to cover a majority of our sample, we classify the top 50 lenders (by origination volume) across the years in our sample period, based on a list from the publication 'Inside B&C mortgage'. In unreported results, we confirm that independent lenders also follow the rule of thumb for low documentation loans.[23] Moreover, low documentation loans securitized by independents with credit score of 620^- are about 15% less likely to default after a year as compared to low documentation loans securitized by them with credit score 620^+.[24] Note that the results in the sample of loans originated by lenders without a strategic selling motive are similar in magnitude to those in the overall sample (which includes other lenders that screen and then may strategically sell). This finding highlights that screening is the driving force behind our results.

3.4.6 Additional Variation From a Natural Experiment

Unrelated Optimal Rule Of Thumb

So far we have worked under the assumption that the 620 threshold is related to securitization. One could plausibly argue, in the spirit of Baumol and Quandt (1964), that this rule of thumb could have been set by lenders as an optimal cutoff for screening that is unrelated to differential

[23]For more details, see Keys et al. (2009).

[24]More specifically, in a specification similar to Column (2) in Panel C of Table 3.3, we find that the coefficient on the indicator T(FICO$\geq$ 620) is 0.67 (t=3.21).

securitization. Ruling this alternative out requires an examination of the effects of the threshold when the ease of securitization varies, everything else equal. To achieve this, we exploit a natural experiment that involves the passage of anti-predatory lending laws in two states which reduced securitization in the subprime market drastically. Subsequent to protests by market participants, the laws were substantially amended and the securitization market reverted to pre-law levels. We use these laws to examine how the main effects vary with the time series variation in the ease of securitization likelihood around the threshold in the two states.

In October 2002, the Georgia Fair Lending Act (GFLA) went into effect, imposing anti-predatory lending restrictions which at the time were considered the toughest in the United States. The law allowed for unlimited punitive damages when lenders did not comply with the provisions and that liability extended to holders in due course. Once GFLA was enacted, the market response was swift. Fitch, Moodys, and S&P were reluctant to rate securitized pools that included Georgia loans. In effect, the demand for the securitization of mortgage loans from Georgia fell drastically during the same period. In response to these actions, the Georgia Legislature amended GLFA in early 2003. The amendments removed many of the GFLAs ambiguities and eliminated covered loans. Subsequent to April 2003, the market revived in Georgia. Similarly, New Jersey enacted its law, the New Jersey Homeownership Security Act of 2002, with many provisions similar to those of the Georgia law. As in Georgia, lenders and ratings agencies expressed concerns when the New Jersey law was passed and decided to substantially reduce the number of loans that were securitized in these markets. The Act was later amended in June 2004 in a way that relaxed requirements and eased lenders' concerns.

If lenders use 620 as an optimal cutoff for screening unrelated to securitization, we expect the passage of these laws to have no effect on the differential screening standards around the threshold. However, if these laws affect the differential ease of securitization around the threshold, our hypothesis would predict an impact on the screening standards. As 620^{+} loans become relatively more difficult to securitize, lenders would internalize the cost of collecting soft information for these loans to a greater degree. Consequently, the screening differentials we observed earlier should attenuate during the period of enforcement. Moreover, we expect the results described in Section 3.4.4 to appear only during the periods when the differential ease of securitization around the threshold is high, i.e., before the law was passed as well as in the period after the law was amended.

Our experimental design examines the ease of securitization and performance of loans above and below the credit threshold in both Georgia and New Jersey during the period when the securitization market was affected and compares it with the period before the law was passed and the period after the law was amended. To do so, we estimate equations (3.1) and (3.2) with an additional dummy variable that captures whether or not the law is in effect (*NoLaw*). We also include time fixed effects to control for any macroeconomic factors independent of the laws.

The results are striking. Panel A of Table 3.4 confirms that the discontinuity in the number of loans around the threshold diminishes during a period of strict enforcement of anti-predatory lending laws. In particular, the difference in number of loans securitized around the credit thresholds fell by around 95% during the period when the law was passed in Georgia and New Jersey. This effectively nullified any meaningful difference in the ease of securitization above the FICO threshold. Another intuitive way to see this is to compare these jumps in the number of loans with jumps in states which had similar housing profiles as Georgia and New Jersey before the law was passed (e.g., Texas in 2001). For instance, relative to the discontinuity in Texas, the jump during the period when the law was passed is about 5%, whereas the jumps are of comparable size both before the law is passed and after the law was amended. In addition, the results also indicate a rapid return of a discontinuity after the law is revoked. It is notable that this time horizon is too brief for any meaningful change in the housing stock (Glaeser and Gyourko 2005), or in the underlying demand for home ownership.

Importantly, our performance results follow the same pattern as well. Columns (1) and (2) of Panel B show that the default rates for 620^{+} loans are below that of 620^{-} loans in both Georgia and New Jersey *only* when the law was in effect. In addition, when the law was either not passed or was amended, we find that default rates for loans above the credit threshold is similar to loans below the credit threshold. This upward shift in the default curve above the 620 threshold is consistent with the results reported in Section 3.4.4. Taken together, these results suggest that our findings are indeed related to differential securitization at the credit threshold and that lenders were not blindly following the rule of thumb in all instances.

Manipulation Of Credit Scores

Having confirmed that lenders are screening more at 620^{-} than 620^{+}, we assess whether borrowers were aware of the differential screening around the threshold. Even though there is no difference in contract terms around the cutoff, screening is weaker above the 620 score than below it, and this may create an incentive for borrowers to manipulate their credit score. If FICO scores could be manipulated, lower quality borrowers might artificially appear at higher credit scores. This behavior would be consistent with our central claim of differential screening around the threshold. Note that as per the rating agency (Fair Isaac), it is difficult to strategically manipulate one's FICO score in a targeted manner. Nevertheless, to examine the response of borrowers more closely, we exploit the variation generated from the same natural experiment.

If FICO scores tend to be quite sticky and it takes relatively long periods of time (more than 3 to 6 months) to improve credit scores, as Fair Isaac claims, we should observe that the difference in performance around the threshold should take time to appear after the laws are reversed. Restricting our analysis to loans originated within six months after the laws were reversed, Columns (3) and (4) of Panel B (Table 3.4) show that the reversal of anti-predatory lending laws has immediate effects on the performance of loans that are securitized. This result suggests that borrowers might not have been aware of the differential screening around the threshold or were unable to quickly manipulate their FICO scores. Overall the evidence in this section is consistent with Mayer and Pence (2008), who find no evidence of manipulation of FICO scores in their survey of the subprime market.[25]

3.4.7 Additional Confirmatory Tests

GSE Selection

Although the subprime market is dominated by the non-agency sector, one might worry that the GSEs may differentially influence the selection of borrowers into the subprime market through their actions in the prime market. For instance, the very best borrowers above the 620 threshold might select out of the subprime market in search of better terms in the prime market. We establish several facts to confirm that this is not the case.

First, the natural experiment we discuss in Section 3.4.6 suggests that prime-influenced selection is not at play. The anti-predatory laws were targeted primarily towards the subprime part of the market (Bostic et al. 2007), while leaving the prime part of the market relatively unaffected. To confirm the behavior of the prime market during the enforcement of anti-predatory laws, we rely on another dataset of mortgages in the US which also covers the agency loan market. The data is collected from the top 10 US servicers and covers the period 2001 to 2006. As reported in Panel A of Table 3.5, during the natural experiment, it was no more difficult to obtain an agency loan than

[25]As a further check, we obtained another dataset of subprime loans that continues to track the FICO scores of borrowers after loan origination. Borrowers who manipulate their FICO scores before loan issuance should experience a decline in FICO score shortly after receiving a loan (because a permanent change in the credit score cannot be considered manipulation). Consistent with evidence for no manipulation around the threshold, we find that both 620^+ and 620^- borrowers are as likely to experience such a reduction within a quarter of obtaining a loan. Results of this analysis are available upon request.

before or after the law was in effect. Similarly, in unreported tests we find that contractual terms (such as LTV ratios and interest rates) around 620 see no change across time periods. Furthermore, in the prime market, there were no differences in defaults around the 620 threshold across the time periods (Table 3.5, Panel B). Since borrower quality in the prime market did not change around 620 threshold across the two time periods, if there was indeed selection, the very best 620^{+} subprime borrowers should have selected out into the prime market even while the laws were in place. As a result, we should have found that 620^{+} borrowers in subprime market continue to default more than 620^{-} borrowers even when the law is in place. As we showed earlier in Table 3.4, this is not the case.

Second, the dataset confirms that Freddie Mac and Fannie Mae primarily do not buy loans with credit scores around FICO of 620 (especially low documentation loans). This is consistent with anecdotal evidence that the role of active subprime securitization in recent years had shifted to non-agency sector (Gramlich, 2007). In unreported permutation tests (see internet appendix Table 4, Panel A), we also find that the number of loans in the agency market is smooth around the 620 threshold. In addition, the loan terms and default rates are also smooth. Together these results suggest that, in general, there seems to be no differential selection in terms of number of loans or quality of loans around the 620 cutoff.

Third, if our results in the low documentation market around 620 threshold are driven by differential GSE selection, we should observe no differences in defaults when we combine the loans from agency with low documentation subprime loans around the 620 threshold. If it were purely selection, lower performance above the threshold in the low documentation subprime loans would be offset by differentially higher quality loans selected into the agency market. Unreported results (internet appendix Table 5) show that there are still differences in default rates around the 620 threshold when we examine the agency loans and low documentation subprime loans together.

Finally, we examine the set of borrowers in the subprime market (around 620) who are offered similar contractual terms as those offered in the prime market. If there is indeed selection into the prime market, it is likely based on contractual terms offered to borrowers. By examining borrowers who are offered similar contractual terms in the subprime market, we are able to isolate our analysis to borrowers of similar quality as those who are possibly attracted by GSEs (i.e., the good quality borrowers). For this subset of subprime borrowers, we are able to show that 620^{+} loans still default more than 620^{-} loans (internet appendix Table 4, Panel B). This evidence further suggests that selection by GSEs is unlikely to explain our results.

Other Thresholds

In the data, we also observe smaller jumps in other parts of the securitized loan FICO distribution as other ad-hoc cutoffs have appeared in the market in the past three years (e.g., 600 for low documentation in 2005 and 2006). We remain agnostic as to why or how these other cutoffs have appeared; either due to greater willingness to lend to riskier borrowers, or changing use of automated underwriting which generally included a matrix of qualifications and loan terms including FICO buckets. Several comments about why we focus on the 620 threshold are therefore in order.

First, the 620 cutoff is the only threshold that is actively discussed by the GSEs in their lending guidelines, where the ease of securitization is higher on the right side of the threshold (see internet appendix exhibit 1). This feature is essential for us to disentangle the effect of lax screening on defaults from what a change in FICO score might predict. As increasing FICO scores predict decreasing default rates, performing our analysis with any cut-off where ease of securitization is lower on the right side of threshold would not allow us to use this identification. For instance, consider the cutoff of 660 that is also discussed in the GSE guidelines and where we observe a jump in securitization. The ease of securitization is lower on the right hand side of this cutoff, i.e., the unconditional probability of securitization is lower at 660^{+} relative to 660^{-}, suggesting that 660^{+} loans would be more intensively screened and would default less frequently than 660^{-}. However, it would be impossible to disentangle this effect from just a mechanical effect of 660^{+} FICO loans being more creditworthy and thus defaulting less often than 660^{-} loans (by construction). This subtle advantage of the 620 cutoff is crucial to our identification strategy and rules out the use of several other ad-hoc thresholds.

Moreover, to identify the effects of securitization on screening by lenders, the liquidity differential for the loan portfolios around the threshold has to be large enough. Since 620 is the largest jump we observe in the loan distribution, it is a natural choice. This is confirmed in the permutation tests, which show that $FICO = 620$ has the smallest p-value (and is thus largest outlier) among *all* the visible discontinuities for *each year* in our sample. While other cutoffs may also induce slight differences in screening effort in some years, these differences may be small to make any meaningful inferences. In results not shown, we analyzed some of these other thresholds and find results for delinquencies that are consistent with those reported for the predominant cutoff (620), but are indeed quite small in magnitude.

Other Tests

We also conduct several falsification tests, repeating our analysis at other credit scores where there is no jump in securitization. In sharp contrast to the results reported in Section 3.4.4, the higher credit score bucket defaults *less* than the lower credit score bucket. This is consistent with the results of the permutation tests reported above, which estimate *every* false discontinuity and compare it to the discontinuity at 620. Moreover, as we will show in Section 3.5, full documentation loans do not see any jumps at this threshold. We plot the delinquency rates of 620^+ and 620^- for full documentation loans (2001-2006) in Figure **??** and find loans made at lower credit scores are more likely to default.[26]

As further tests of our hypothesis, we also conducted our tests in the refinance market, and find a similar rule of thumb and similar default outcomes around the 620 threshold in this market. Finally, we re-estimated our specifications with state, lender and pool fixed effects to account for multiple levels of potential variation in the housing market and find qualitatively similar results.[27]

3.5 Did Hard Information Matter?

The results presented above are for low documentation loans, which necessarily have an unobserved component of borrowers' creditworthiness. In the full documentation loan market, on the other hand, there is no omission of hard information on the borrower's ability to repay. In this market, we identify a credit threshold at the FICO score of 600, the score that Fair Isaac (and the three credit repositories) advises lenders as a bottom cutoff for low risk borrowers. They note "...anything below 600 is considered someone who probably has credit problems that need to be addressed..."(see www.myfico.com). Similarly Fannie Mae in its guidelines notes "...a borrower with credit score of 600 or less has a high primary risk..." (see www.allregs.com/efnma/doc/). The Consumer Federation of America along with Fair Isaac (survey report in March 2005) suggests that "...FICO credit scores range from 300-850, and a score above 700 indicates relatively low credit risk, while scores below 600 indicate relatively high risk which could make it harder to get credit or lead to higher loan rates." Einav, Jenkins and Levin (2007) make a similar observation when they note that "...a FICO score above 600 [is] a typical cut-off for obtaining a standard bank loan."

Figure **??** reveals that there is a substantial increase in the number of full documentation loans above the credit threshold of 600. This pattern is consistent with the notion that lenders are more willing to securitize at a lower credit threshold (600 vs. 620) for full documentation loans since

[26]This test can also provide insight into the issue of GSE selection discussed earlier. Since 620^+ full documentation loans do not default more than 620^- loans, differential selection into the agency market must account for this fact as well. One possibility is selection on the basis of debt-to-income ratios. To examine this, we compare DTI ratios in the full and low documentation markets. Unreported tests (internet appendix Table 3) show that the DTI ratios are similar around the threshold and thus cannot entirely explain results across the two types of loans.

[27]For additional information on tests across types of lenders and states, see Keys et al. (2009).

there is less uncertainty about these borrowers relative to those who provide less documentation. The magnitudes are again large – around 100% higher at 600^+ than at 600^- in 2004 – for full documentation loans. In Panel A of Table 3.6, we estimate regressions similar to equation (3.1) and find the coefficient estimate is also significant at 1% and is on average around 100% (from 80 to 141%) higher for 600^+ as compared to 600^- for post-2000 loans. Again, if the underlying creditworthiness and the demand for mortgage loans (at a given price) is the same for potential buyers with a credit score of 600^- or 600^+, as the credit bureaus claim, this result confirms that it is easier to securitize full documentation loans above the 600 FICO threshold. We repeated a similar analysis for loan characteristics (LTV and interest rates) and borrower demographics and find no differences for full documentation loans above and below the credit score of 600. Table 3.A3 in appendix presents the estimates from the regressions (Table 3.A4 provides permutation test estimates corresponding to these loan terms).

Interestingly, we find that for full documentation loans, those with credit scores of 600^- (FICO between 595 and 599) are about as likely to default after a year as compared to loans of credit score 600^+ (FICO between 601 and 605) for the post-2000 period. Both Figures **??** and **??** and results in Panels B, C and D of Table 3.6 support this conjecture. Following the methodology used in Figures **??** and 3.18, we show the default rates annually across the FICO distribution (Figure **??**) and across the age of the loans (Figure **??**). The estimated effects of the ad-hoc rule on defaults are negligible in all specifications.

The absence of differences in default rates around the credit threshold, while maintaining the same magnitude of the jump in the number of loans, is consistent with the notion that the pattern of delinquencies around the low-documentation threshold are primarily due to the soft information of the borrower. With so much information collected by the lender for full documentation loans, there is less value to collecting soft information. Consequently, for full documentation loans there is no difference in how the loans perform subsequently after hard information has been controlled for. Put another way, differences in returns to screening are attenuated due to the presence of more hard information.

3.6 Discussion

In the wake of the subprime mortgage crisis, a central question confronting market participants and policymakers is whether securitization had an adverse effect on the ex-ante screening effort of loan originators. Comparing characteristics of the loan market above and below the ad-hoc credit threshold, we show that a doubling of securitization volume is on average associated with about a 10-25% increase in defaults. Notably, our empirical strategy delivers only inferences on differences in the performance of loans around this threshold. While we cannot infer what the optimal level of screening at each credit score ought to be, we conclude from our empirical analysis that there was a causal link between ease of securitization and screening. That we find any effect on default behavior in one portfolio compared to another with virtually identical risk profiles, demographic characteristics, and loan terms suggests that the ease of securitization may have a direct impact on incentives elsewhere in the subprime housing market, as well as in other securitized markets.

The results of this paper, in particular from the anti-predatory lending laws' natural experiment, confirm that lender behavior in the subprime market did change based on the ease of securitization. This suggests that existing securitization practices did not ensure that a decline in screening standards would be counteracted by requiring originators to hold more of the loans' default risk. If lenders were in fact holding on to optimal risk where it was easier to securitize, there should have been no differences in defaults around the threshold. This finding resonates well with concerns surrounding the subprime crisis that, in an environment with limited disclosure on who holds what in the originate-to-distribute chain, there may have been insufficient 'skin in the game' for some lenders (Blinder 2007; Stiglitz 2007). At the same time, the results further suggest that the breakdown in the process only occurred for loans where soft information was particulary important. With enough hard information, as in the full documentation market, there may be less value in requiring market participants to hold additional risk to counteract the potential moral hazard of reduced screening standards.

In a market as competitive as the market for mortgage-backed securities, our results on interest rates are puzzling. Lenders' compensation on either side of the threshold should reflect differences in default rates, and yet we find that the interest rates to borrowers are similar on either side of 620. The difference in defaults, despite similar compensation around the threshold, suggests that there may have been some efficiency losses. Of course, it is possible that from the lenders' perspective, a higher propensity to default above the threshold could have exactly offset the benefits of additional liquidity – resulting in identical interest rates around the threshold.

Our analysis remains agnostic about whether investors accurately priced the moral hazard aspects of securitization. It may have been the case that moral hazard existed in this market though investors appropriately priced persistent differences in performance around the threshold (see Rajan et al. 2008). On the other hand, developing an arbitrage strategy to exploit this opportunity may have been prohibitively difficult given that loans are pooled across the FICO spectrum before they are traded. In addition, these fine differences in performance around the FICO threshold could have been obscured by the performance of other complex loan products in the pool. Understanding these aspects of investor behavior warrants additional investigation.

It is important to note that we refrain from making any welfare claims. Our conclusions should be directed at securitization practices as they were during the subprime boom rather than at the optimally designed originate-to-distribute model. We believe securitization is an important innovation and has several merits. It is often asserted that securitization improves the efficiency of credit markets. The underlying assumption behind this assertion is that there is no information loss in transmission, even though securitization increases the distance between borrowers and investors. The benefits of securitization are limited by information loss, and in particular the costs we document in the paper. More generally, what types of credit products should be securitized? We conjecture that the answer depends crucially on the information structure: loans with more hard information are likely to benefit from securitization as compared to loans that involve soft information. A careful investigation of this question is a promising area for future research.

More broadly, our findings caution against policy that emphasizes excessive reliance on default models. Our research suggests that by relying entirely on hard information variables like FICO scores, these models ignore essential elements of strategic behavior on the part of lenders which are likely to be important. The formation of a rule of thumb, even if optimal (Baumol and Quandt 1964), has an undesirable effect on the incentives of lenders to collect and process soft information. As in Lucas (1976), this strategic behavior can alter the relationship between observable borrower characteristics and default likelihood, rather than moving along the previous predicted relationship. Incorporating these strategic elements into default models, although challenging, is another important direction for future research.

Figure 3.1: FICO distribution (US population)

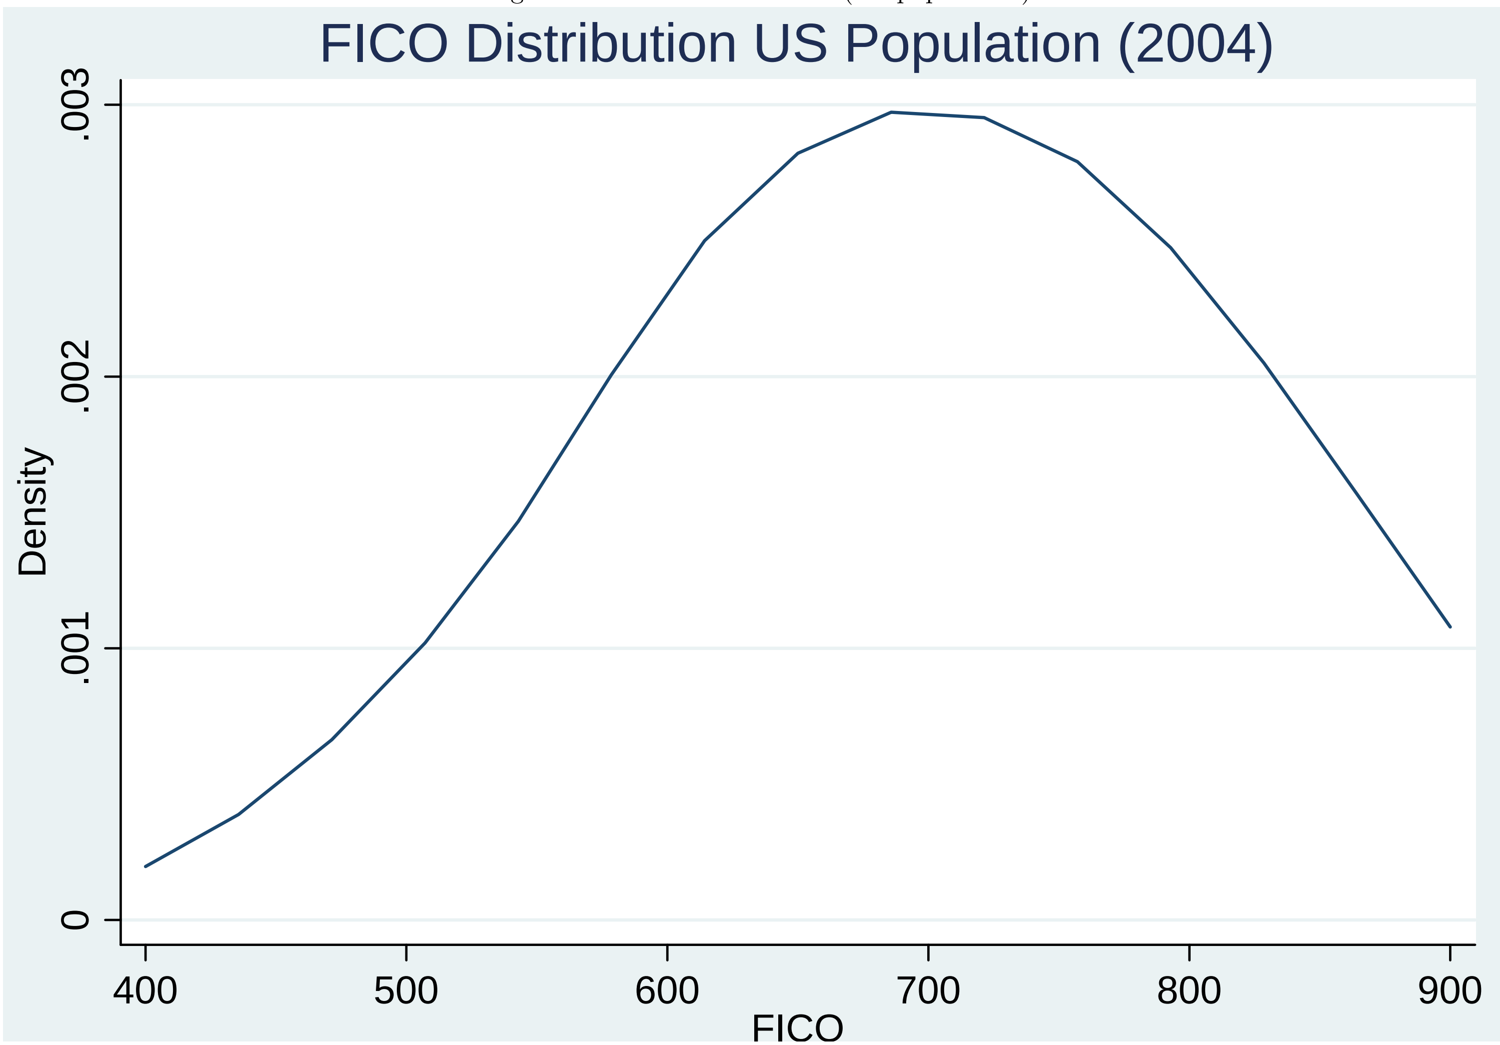

Figure 3.1 presents the FICO distribution in the U.S. population for 2004. This data is from an anonymous credit bureau which assures us that the data exhibits similar patterns during the other years of our sample. The FICO distribution across the population is smooth, so the number of prospective borrowers in the local vicinity of a given credit score is similar.

Figure 3.2: Number of loans (low documentation)

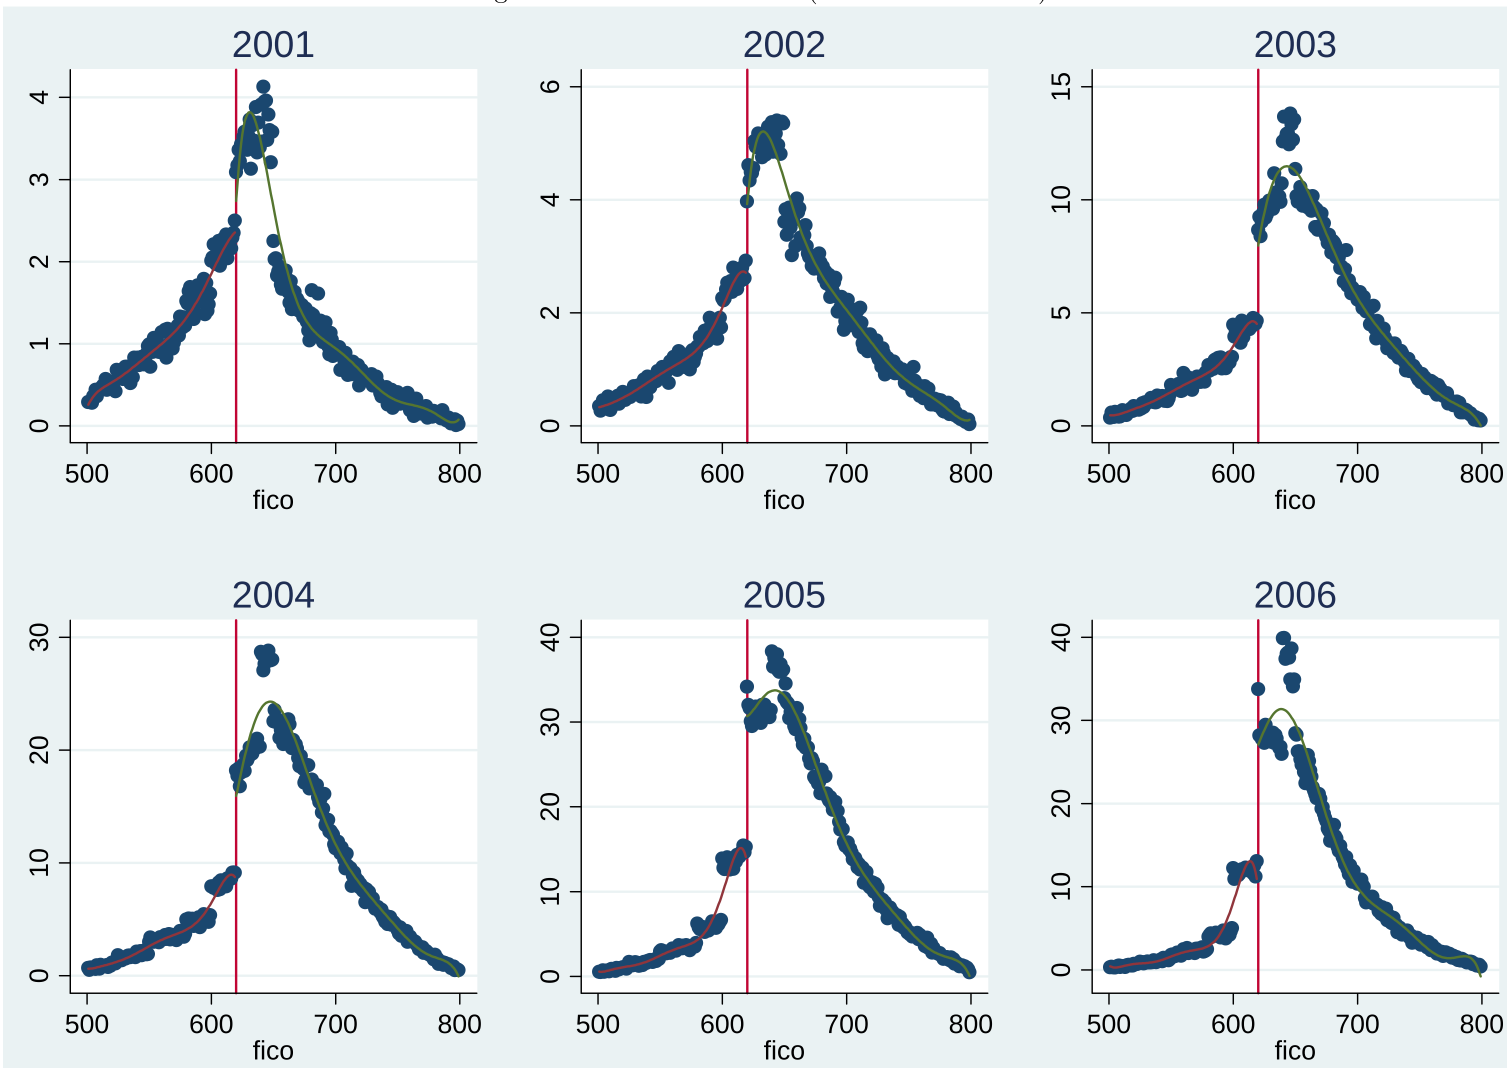

Figure 3.2 presents the data for number of low documentation loans (in '00s). We plot the average number of loans at each FICO score between 500 and 800. As can be seen from the graphs, there is a large increase in the number of loans around the 620 credit threshold (i.e., more loans at 620^+ as compared to 620^-) from 2001 onwards. Data is for the period 2001 to 2006.

Figure 3.3: Interest rates (low documentation)

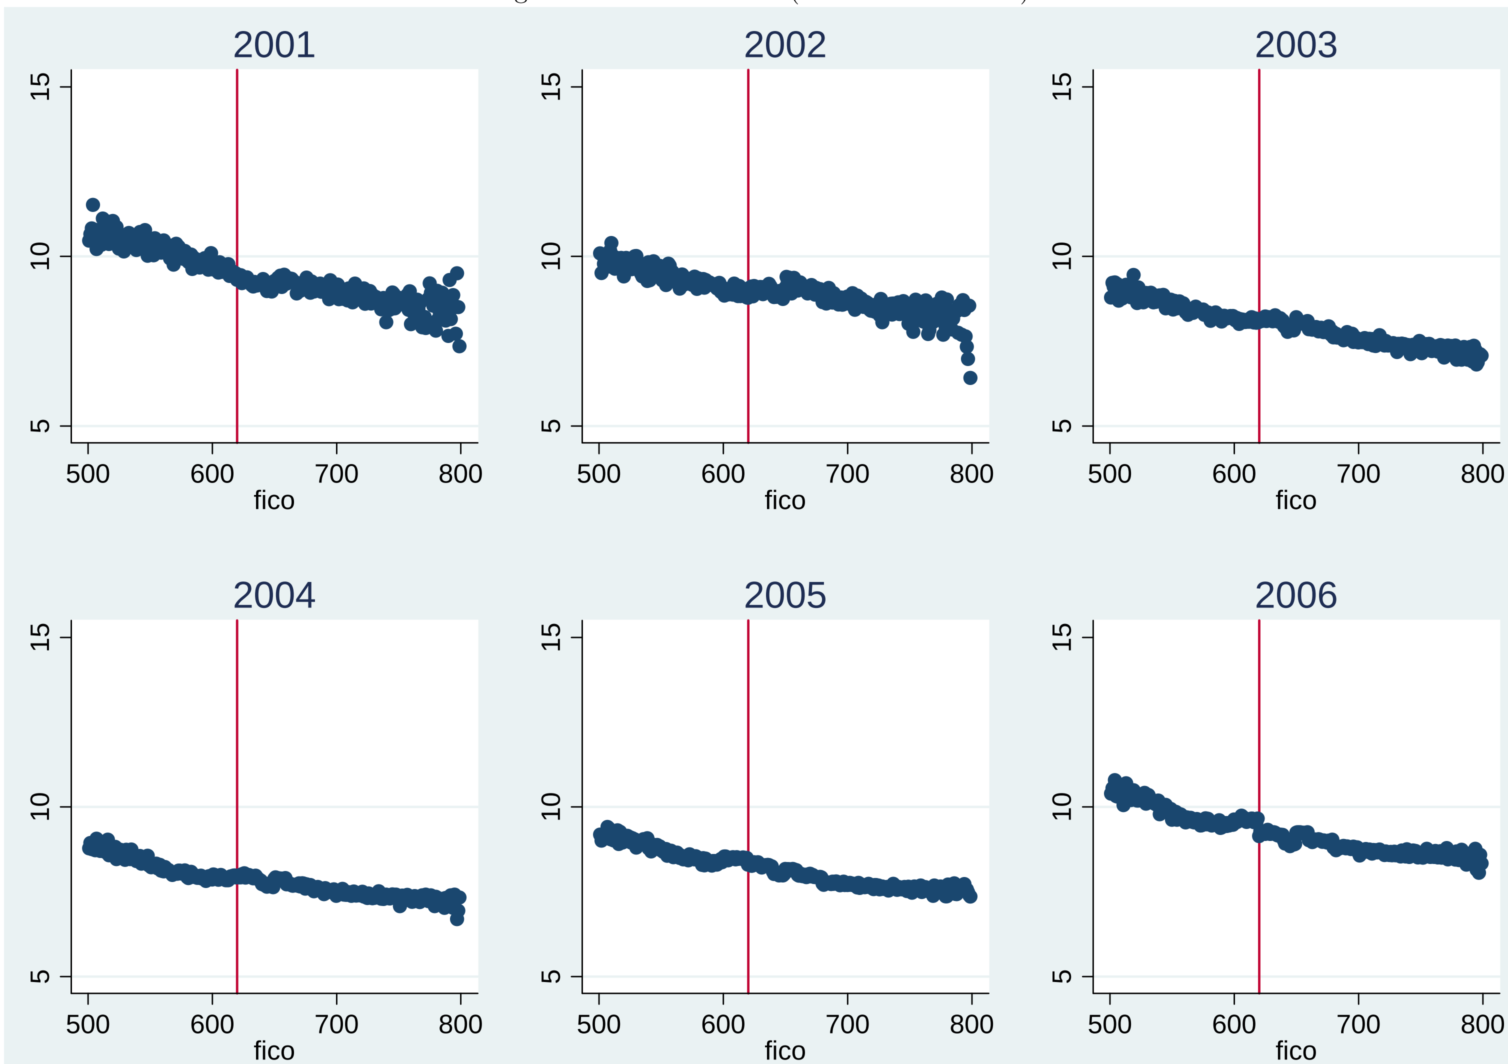

Figure 3.3 presents the data for interest rate (in %) on low documentation loans. We plot average interest rates on loans at each FICO score between 500 and 800. As can be seen from the graphs, there is no change in interest rates around the 620 credit threshold (i.e., more loans at 620^+ as compared to 620^-) from 2001 onwards. Data is for the period 2001 to 2006.

Figure 3.4: Loan-to-Value ratios (low documentation)

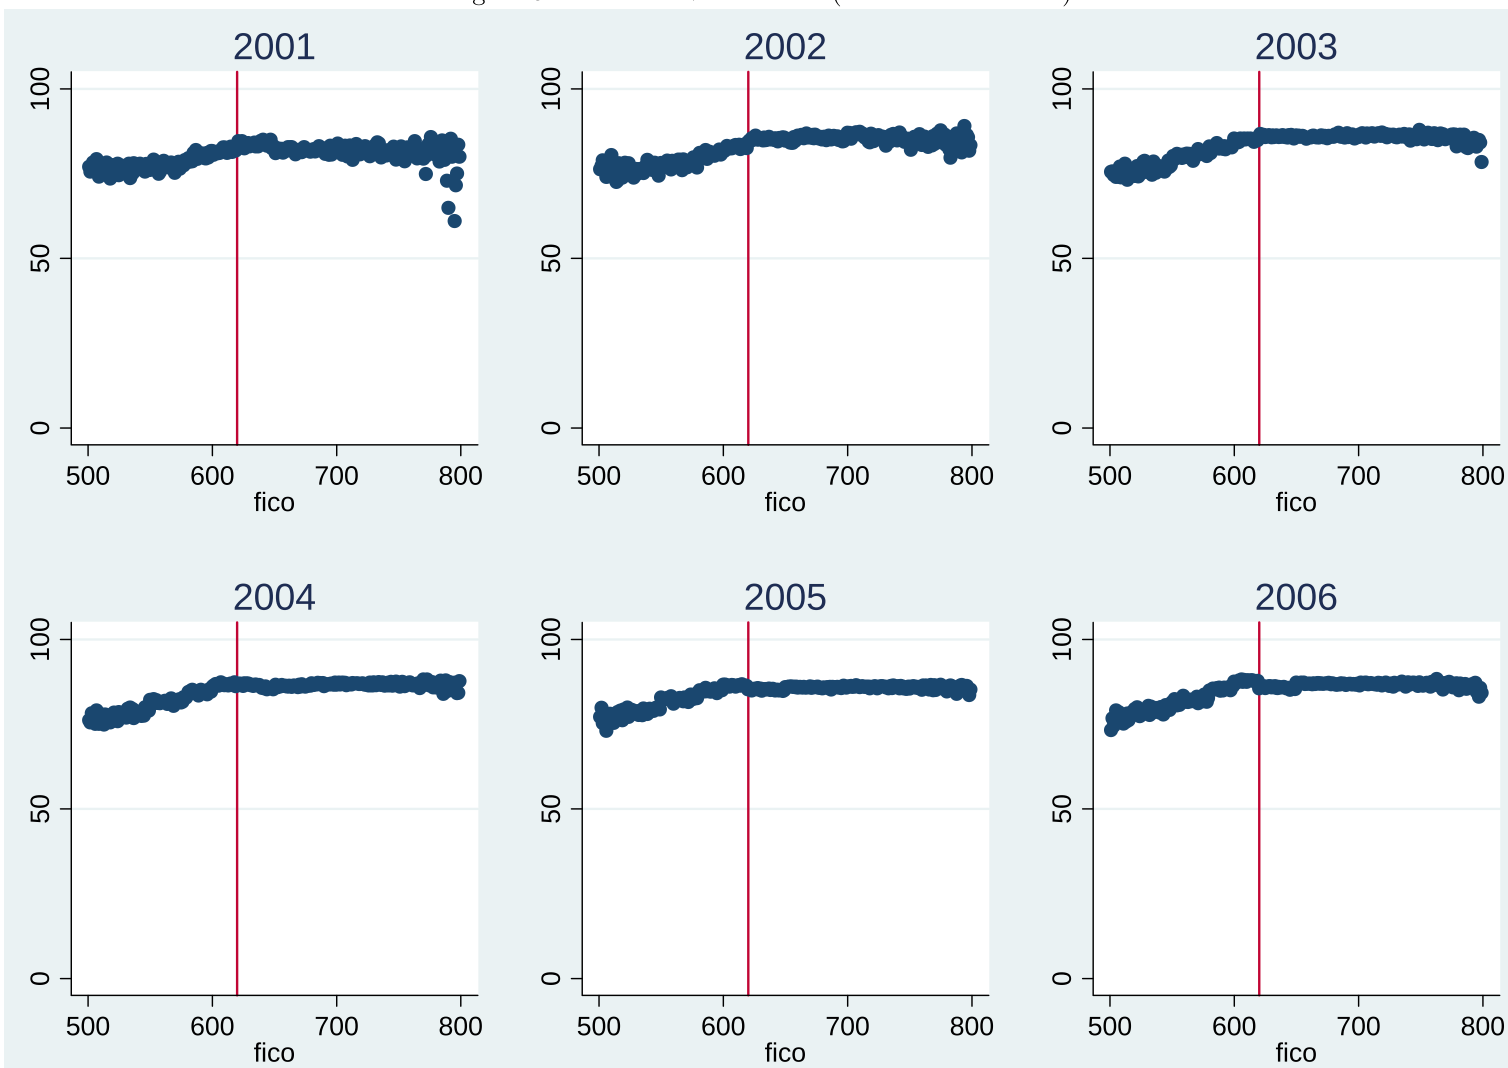

Figure 3.4 presents the data for loan-to-value ratio (in %) on low documentation loans. We plot average loan-to-value ratios on loans at each FICO score between 500 and 800. As can be seen from the graphs, there is no change in loan-to-value around the 620 credit threshold (i.e., more loans at 620^+ as compared to 620^-) from 2001 onwards. Data is for the period 2001 to 2006.

Figure 3.5: Median household income (low documentation)

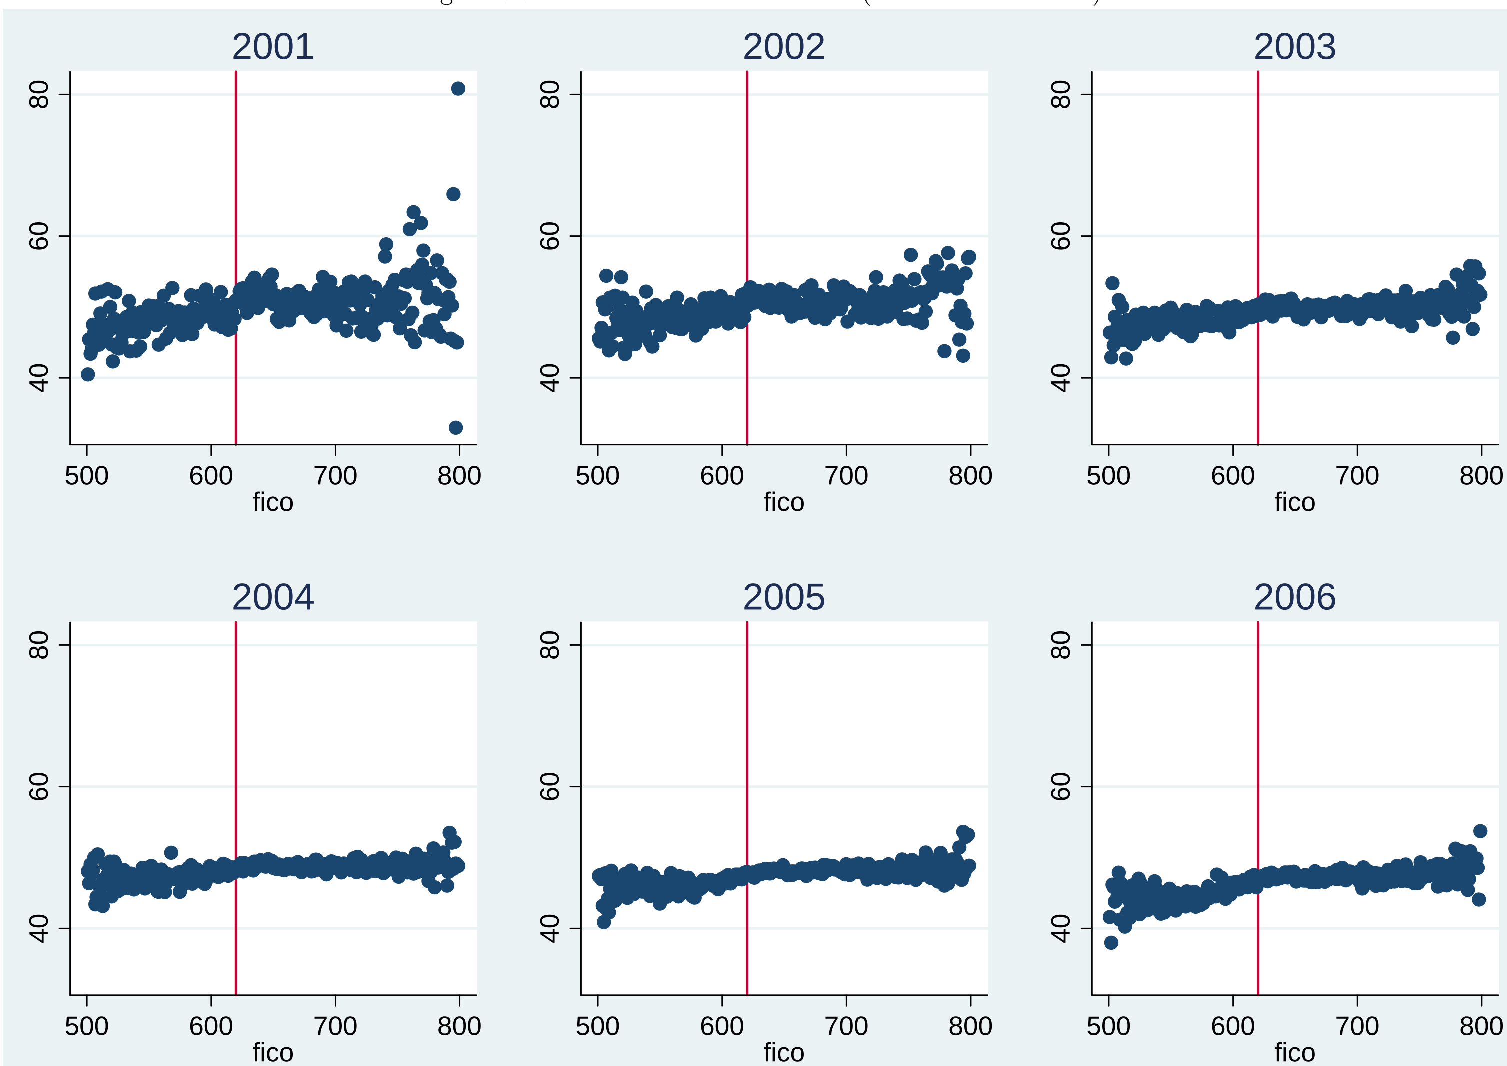

Figure 3.5 presents median household income (in '000s) of zip codes in which loans are made at each FICO score between 500 and 800. As can be seen from the graphs, there is no change in median household income around the 620 credit threshold (i.e., more loans at 620^+ as compared to 620^-) from 2001 onwards. We plotted similar distributions for average percent minorities taking loans, and average house size and find no differences around the credit thresholds. Data is for the period 2001 to 2006.

Figure 3.6: Annual delinquencies for low documentation loans originated in 2001

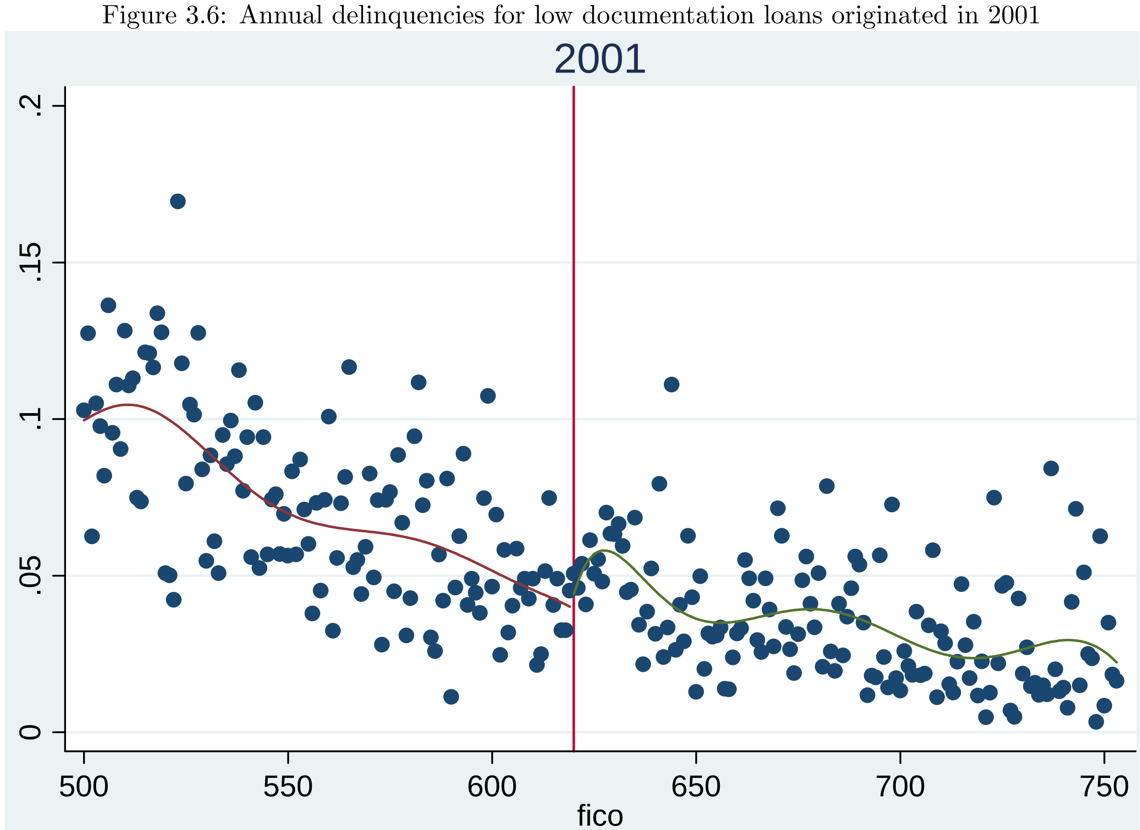

Figure 3.6 presents the percent of low documentation loans originated in 2001 that became delinquent. We plot the dollar weighted fraction of the pool that becomes delinquent for one-point FICO bins between score of 500 and 750. The vertical line denotes the 620 cutoff, and a seventh order polynomial is fit to the data on either side of the threshold. Delinquencies are reported between 10-15 months for loans originated in the year.

Figure 3.7: Annual delinquencies for low documentation loans originated in 2002

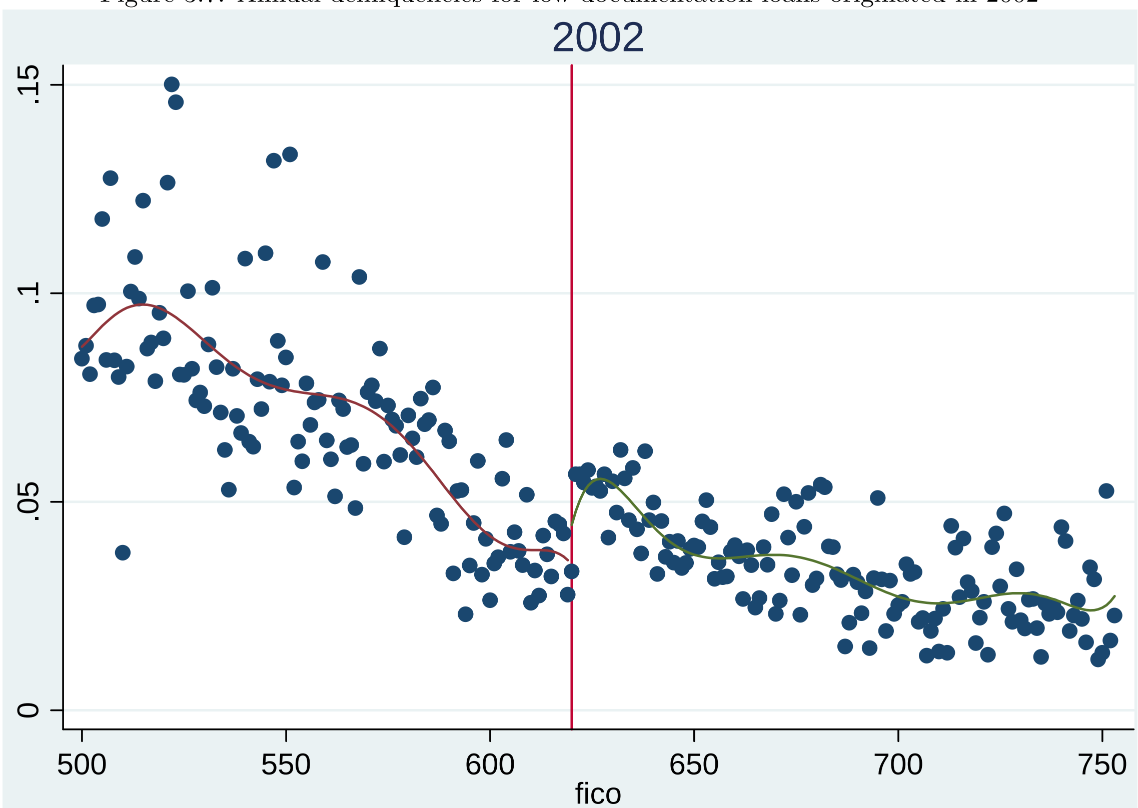

Figure 3.7 presents the percent of low documentation loans originated in 2002 that became delinquent. We plot the dollar weighted fraction of the pool that becomes delinquent for one-point FICO bins between score of 500 and 750. The vertical line denotes the 620 cutoff, and a seventh order polynomial is fit to the data on either side of the threshold. Delinquencies are reported between 10-15 months for loans originated in the year.

Figure 3.8: Annual delinquencies for low documentation loans originated in 2003

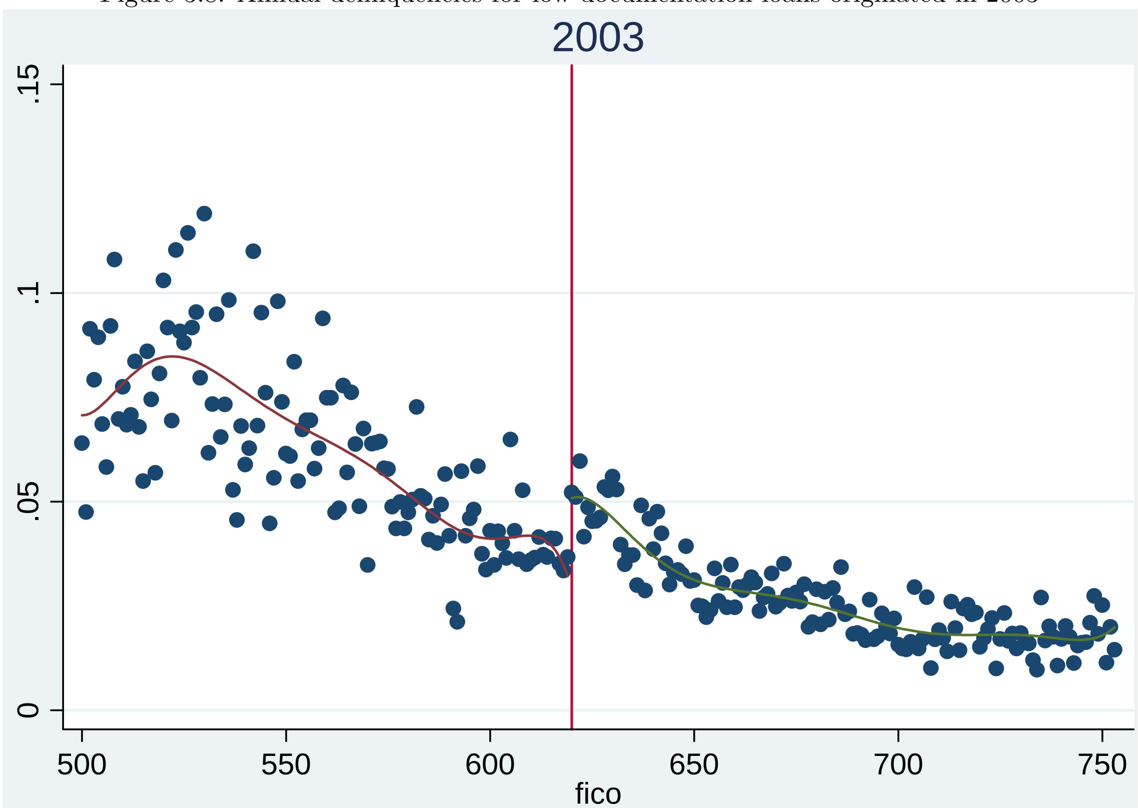

Figure 3.8 presents the percent of low documentation loans originated in 2003 that became delinquent. We plot the dollar weighted fraction of the pool that becomes delinquent for one-point FICO bins between score of 500 and 750. The vertical line denotes the 620 cutoff, and a seventh order polynomial is fit to the data on either side of the threshold. Delinquencies are reported between 10-15 months for loans originated in the year.

Figure 3.9: Annual delinquencies for low documentation loans originated in 2004

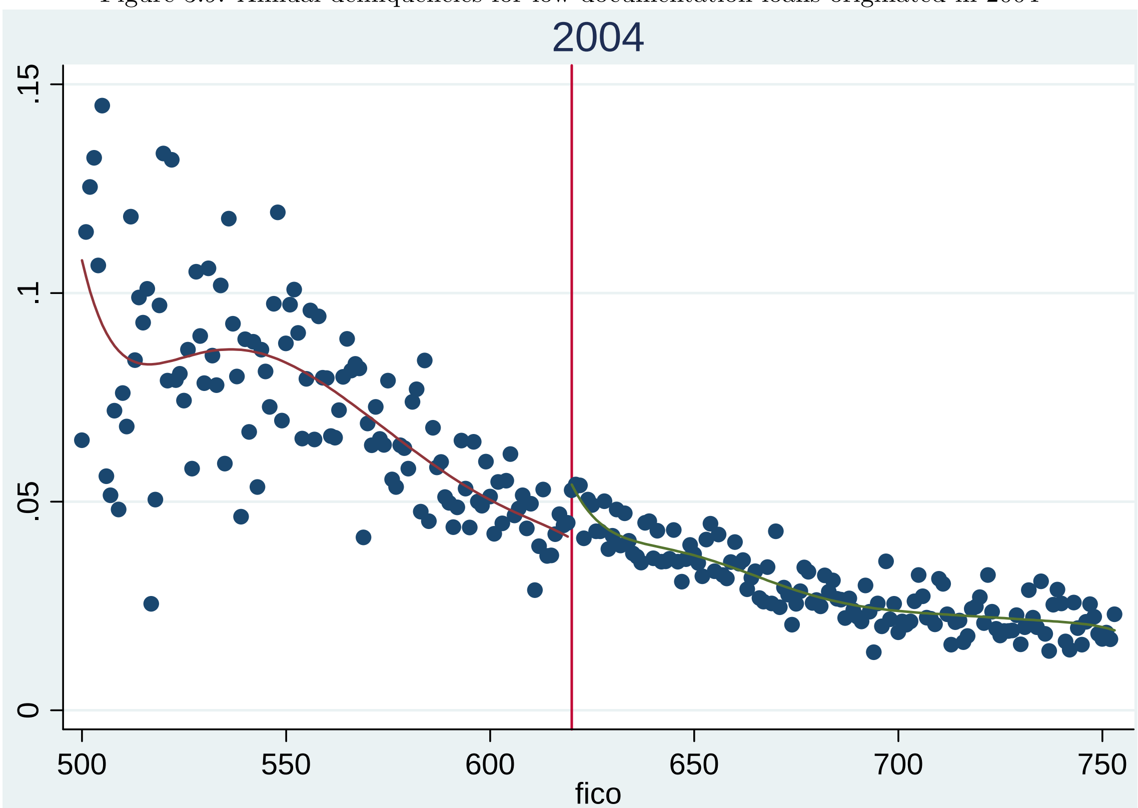

Figure 3.9 presents the percent of low documentation loans originated in 2004 that became delinquent. We plot the dollar weighted fraction of the pool that becomes delinquent for one-point FICO bins between score of 500 and 750. The vertical line denotes the 620 cutoff, and a seventh order polynomial is fit to the data on either side of the threshold. Delinquencies are reported between 10-15 months for loans originated in the year.

Figure 3.10: Annual delinquencies for low documentation loans originated in 2005

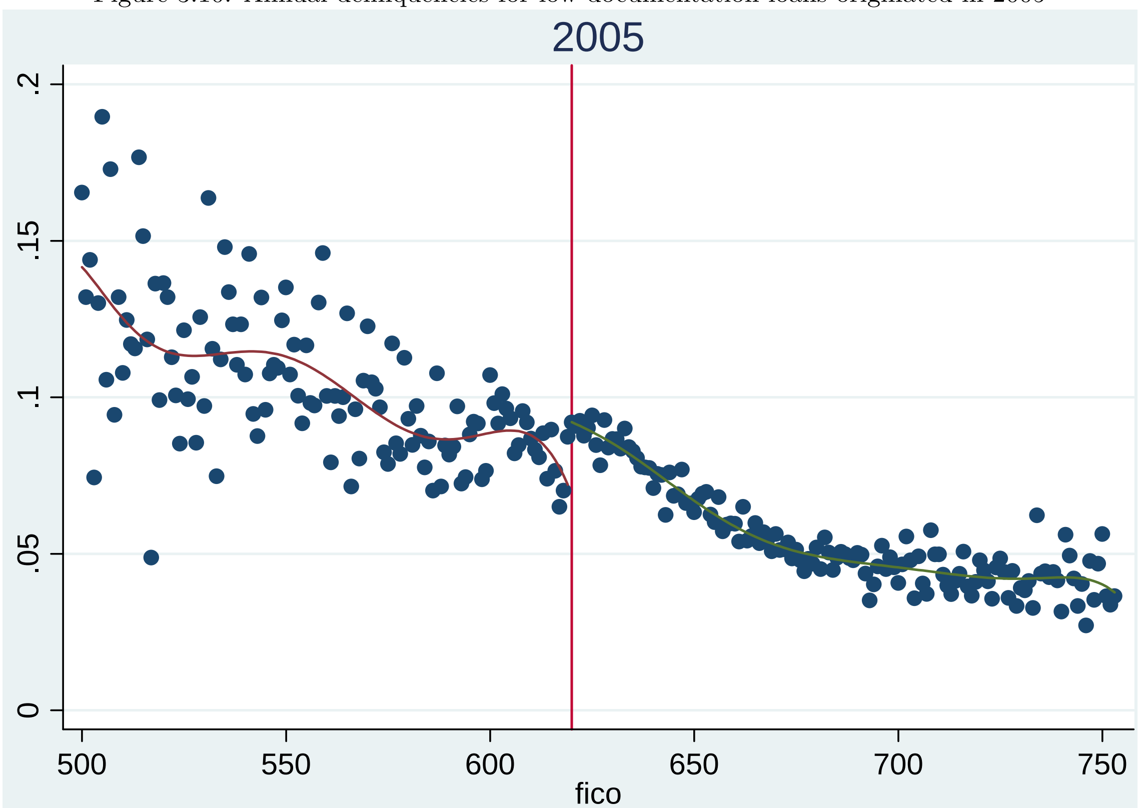

Figure 3.10 presents the percent of low documentation loans originated in 2005 that became delinquent. We plot the dollar weighted fraction of the pool that becomes delinquent for one-point FICO bins between score of 500 and 750. The vertical line denotes the 620 cutoff, and a seventh order polynomial is fit to the data on either side of the threshold. Delinquencies are reported between 10-15 months for loans originated in the year.

Figure 3.11: Annual delinquencies for low documentation loans originated in 2006

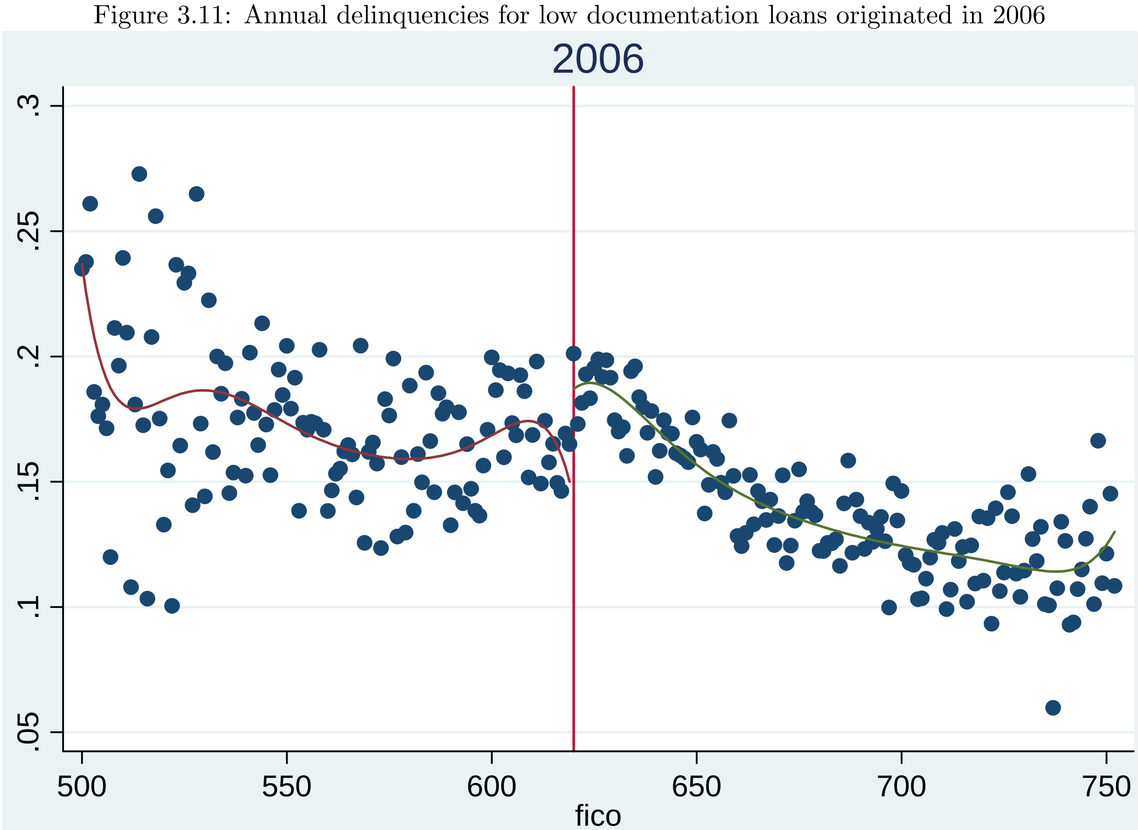

Figure 3.11 presents the percent of low documentation loans originated in 2006 that became delinquent. We plot the dollar weighted fraction of the pool that becomes delinquent for one-point FICO bins between score of 500 and 750. The vertical line denotes the 620 cutoff, and a seventh order polynomial is fit to the data on either side of the threshold. Delinquencies are reported between 10-15 months for loans originated in the year.

Figure 3.12: Delinquencies for low documentation loans (2001-2006)

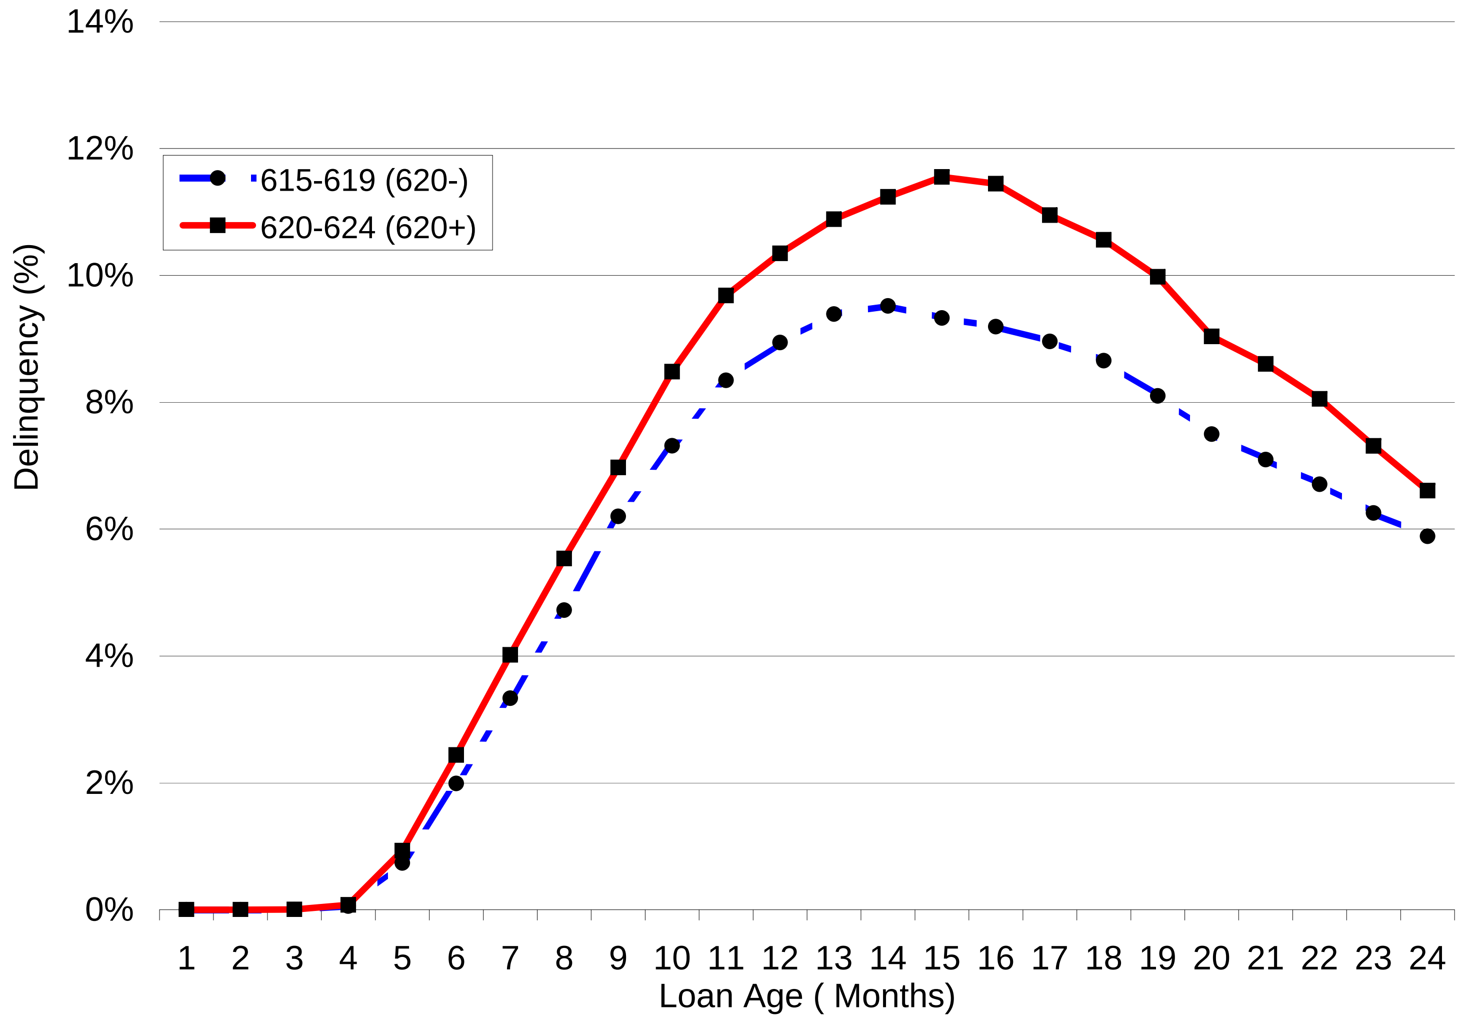

Figure 3.12 presents the percent of low documentation loans (dollar weighted) originated between 2001 and 2006 that subsequently became delinquent. We track loans in two FICO buckets – 615-619 (620^-) in dotted blue and 620-624 (620^+) in red – from their origination date and plot the average loans that become delinquent each month after the origination date. As can be seen, the higher credit score bucket defaults *more* than the lower credit score bucket for post 2000 period. For brevity, we do not report plots separately for each year of origination. The effects shown here in the pooled 2001-2006 plot are apparent in every year.

Figure 3.13: Actual prepayments for low documentation loans (2001-2006)

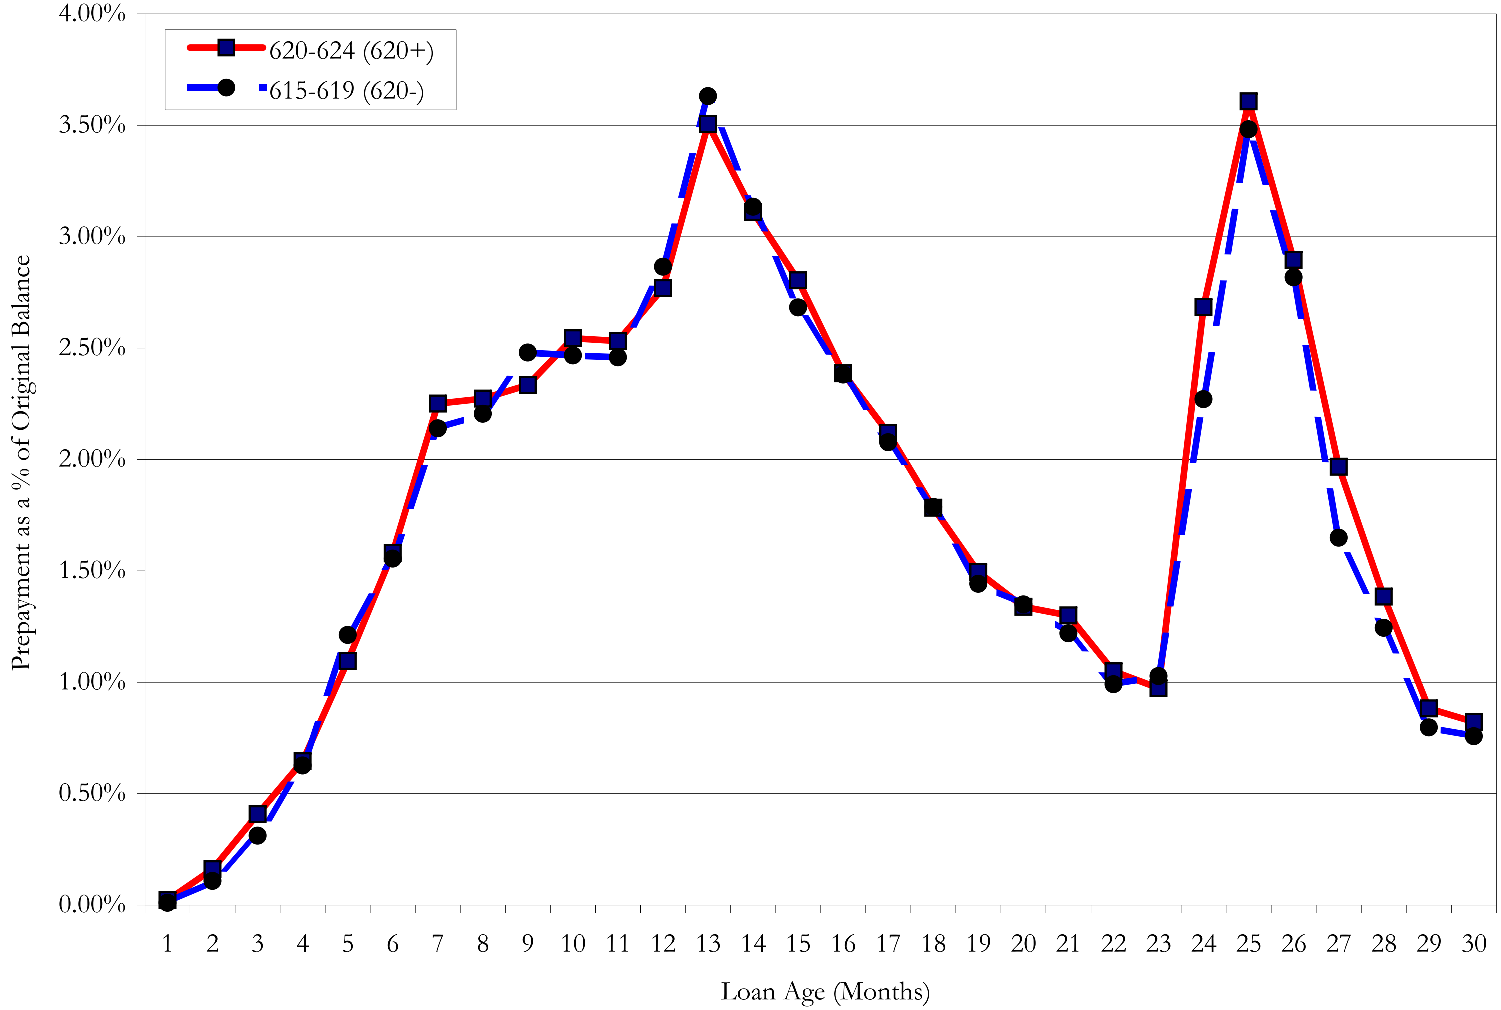

Figure 3.13 presents the percent of low documentation loans (dollar weighted) originated between 2001 and 2006 that subsequently were prepaid. We track loans in two FICO buckets – 615-619 (620^{-}) in dotted blue and 620-624 (620^{+}) in red – from their origination date and plot the average loans that prepaid each month after the origination date. As can be seen, there are no differences in prepayments between the higher and lower credit score bucket. For brevity, we do not report plots separately for each year of origination. The effects shown here in the pooled 2001-2006 plot are apparent in every year.

Figure 3.14: Dispersion of interest rates (low documentation)

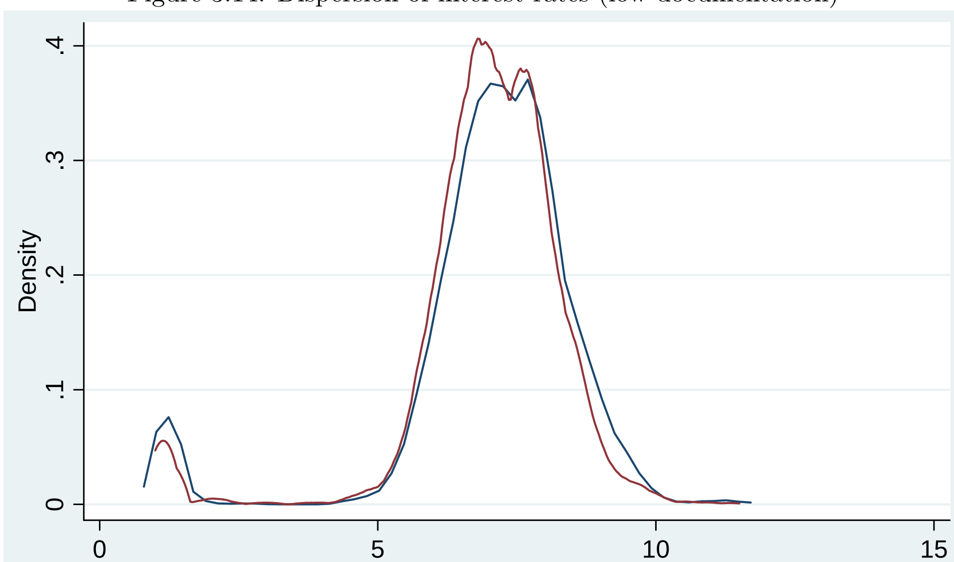

Figure 3.14 depicts the Epanechnikov kernel density of interest rate for two FICO groups for low documentation loans – 620^{-} (615-619) in blue and 620^{+} (620-624) in red. The bandwidth for the density estimation is selected using the plug-in formula of Sheather and Jones (1991). The figure shows that the density of interest rates on loans are similar for both the groups. A Kolmogorov-Smirnov test for equality of distribution functions cannot be rejected at the 1% level. Data for loans originated in 2004 is reported here. We find similar patterns for 2001-2006 originations. We do not report those graphs for brevity.

Figure 3.15: Dispersion of loan-to-value ratios (low documentation)

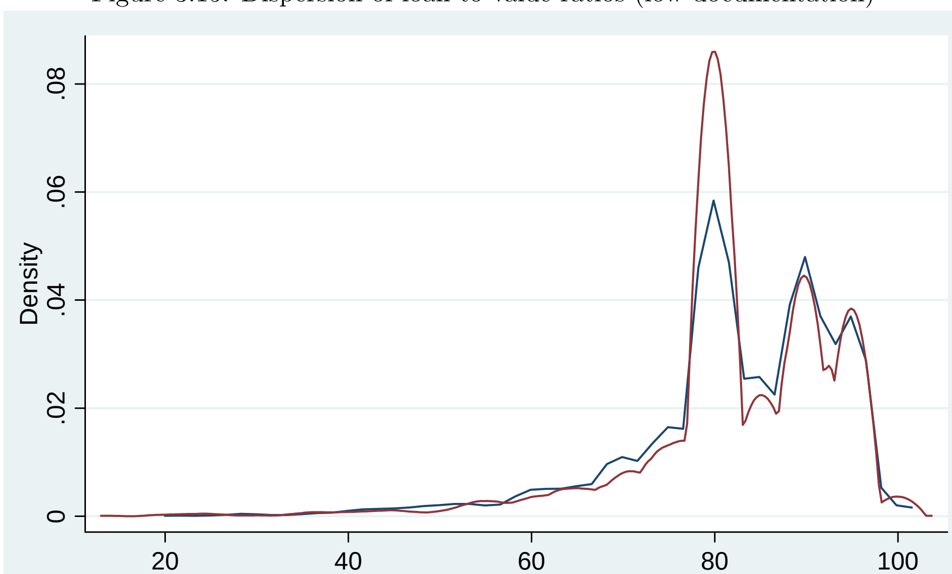

Figure 3.15 depicts the Epanechnikov kernel density of loan-to-value ratio for two FICO groups for low documentation loans – 620^{-} (615-619) in blue and 620^{+} (620-624) in red. The bandwidth for the density estimation is selected using the plug-in formula of Sheather and Jones (1991). The figure shows that the density of interest rates on loans are similar for both the groups. A Kolmogorov-Smirnov test for equality of distribution functions cannot be rejected at the 1% level. Data for loans originated in 2004 is reported here. We find similar patterns for 2001-2006 originations. We do not report those graphs for brevity.

Figure 3.16: Falsification test - Delinquencies for full documentation loans around FICO of 620

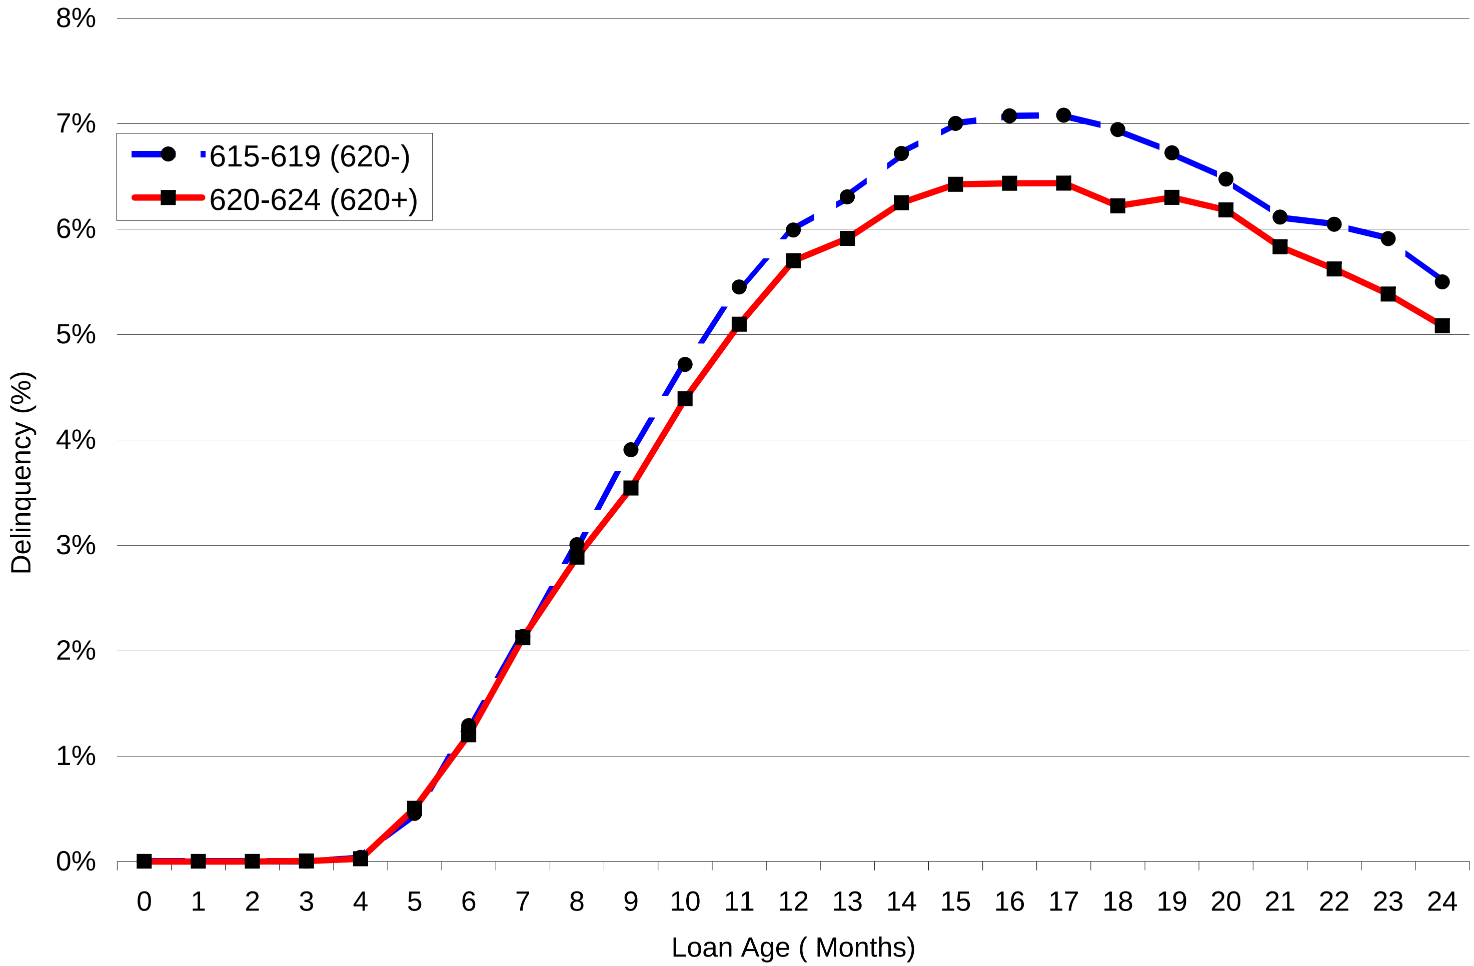

Figure 3.16 presents the falsification test by examining the percent of full documentation loans (dollar weighted) originated between 2001 and 2006 that became delinquent. We track loans in two FICO buckets – 615-619 (620^{-}) in dotted blue and 620-624 (620^{+}) in red – from their origination date and plot the average loans that become delinquent each month after the origination date. As can be seen, the higher credit score bucket defaults *less* than the lower credit score bucket for post 2000 period. For brevity, we do not report plots separately for each year of origination. The effects shown here in the pooled 2001-2006 plot show up for every year.

Figure 3.17: Number of loans (full documentation)

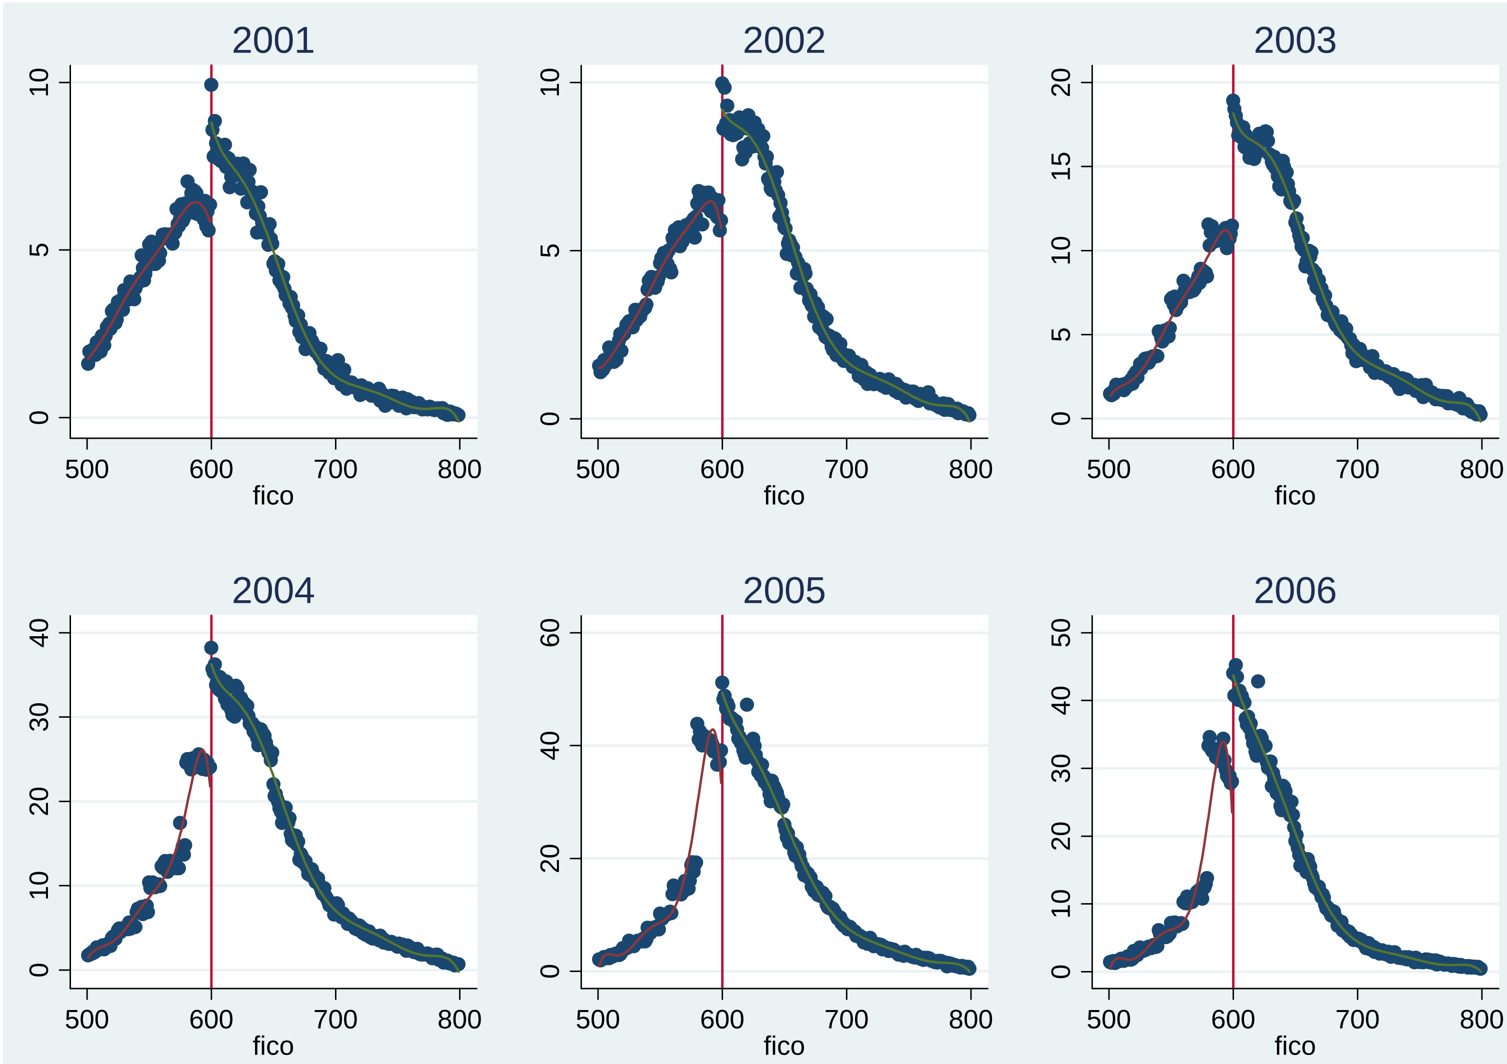

Figure 3.17 presents the data for the number of full documentation loans (in '00s). We plot the average number of loans at each FICO score between 500 and 800. As can be seen from the graphs, there is a large increase in number of loans around the 600 credit threshold (i.e., more loans at 600^{+} as compared to 600^{-}) from 2001 onwards. Data is for loans originated between 2001 to 2006.

Figure 3.18: Annual delinquencies for full documentation loans

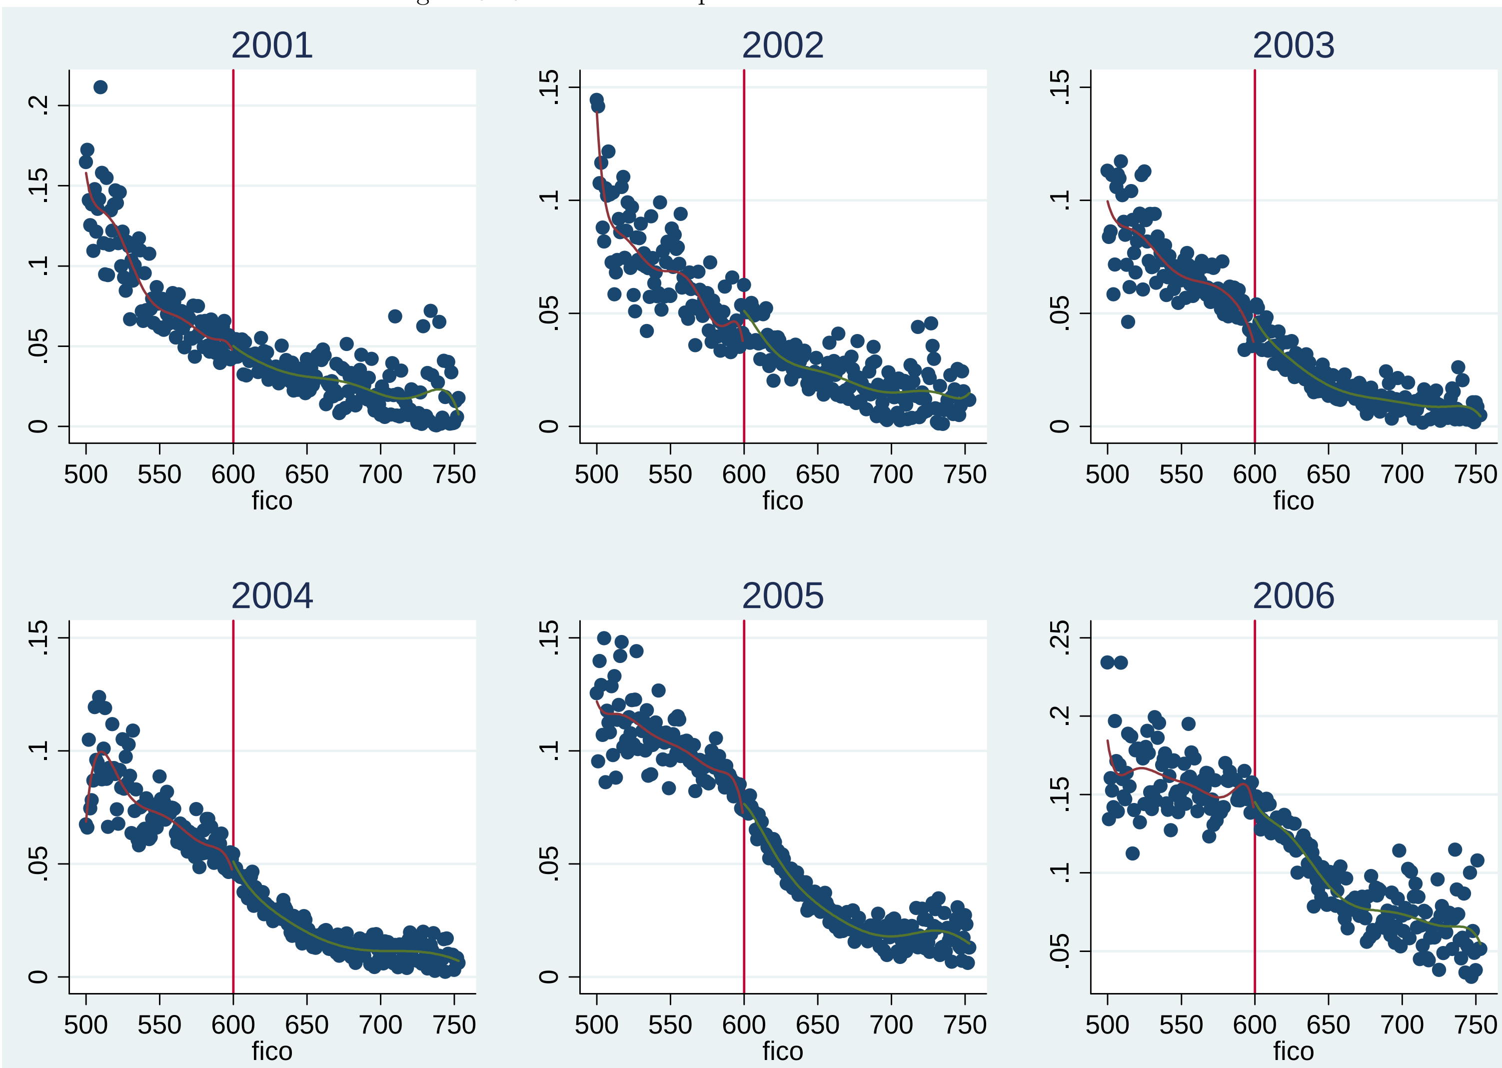

Figure 3.18 presents the percent of full documentation loans originated between 2001 and 2006 that became delinquent. We plot the dollar weighted fraction of the pool that becomes delinquent for one-point FICO bins between score of 500 and 750. The vertical line denotes the 600 cutoff, and a seventh order polynomial is fit to the data on either side of the threshold. Delinquencies are reported between 10-15 months for loans originated in all years.

Figure 3.19: Delinquencies for full documentation loans (2001-2006)

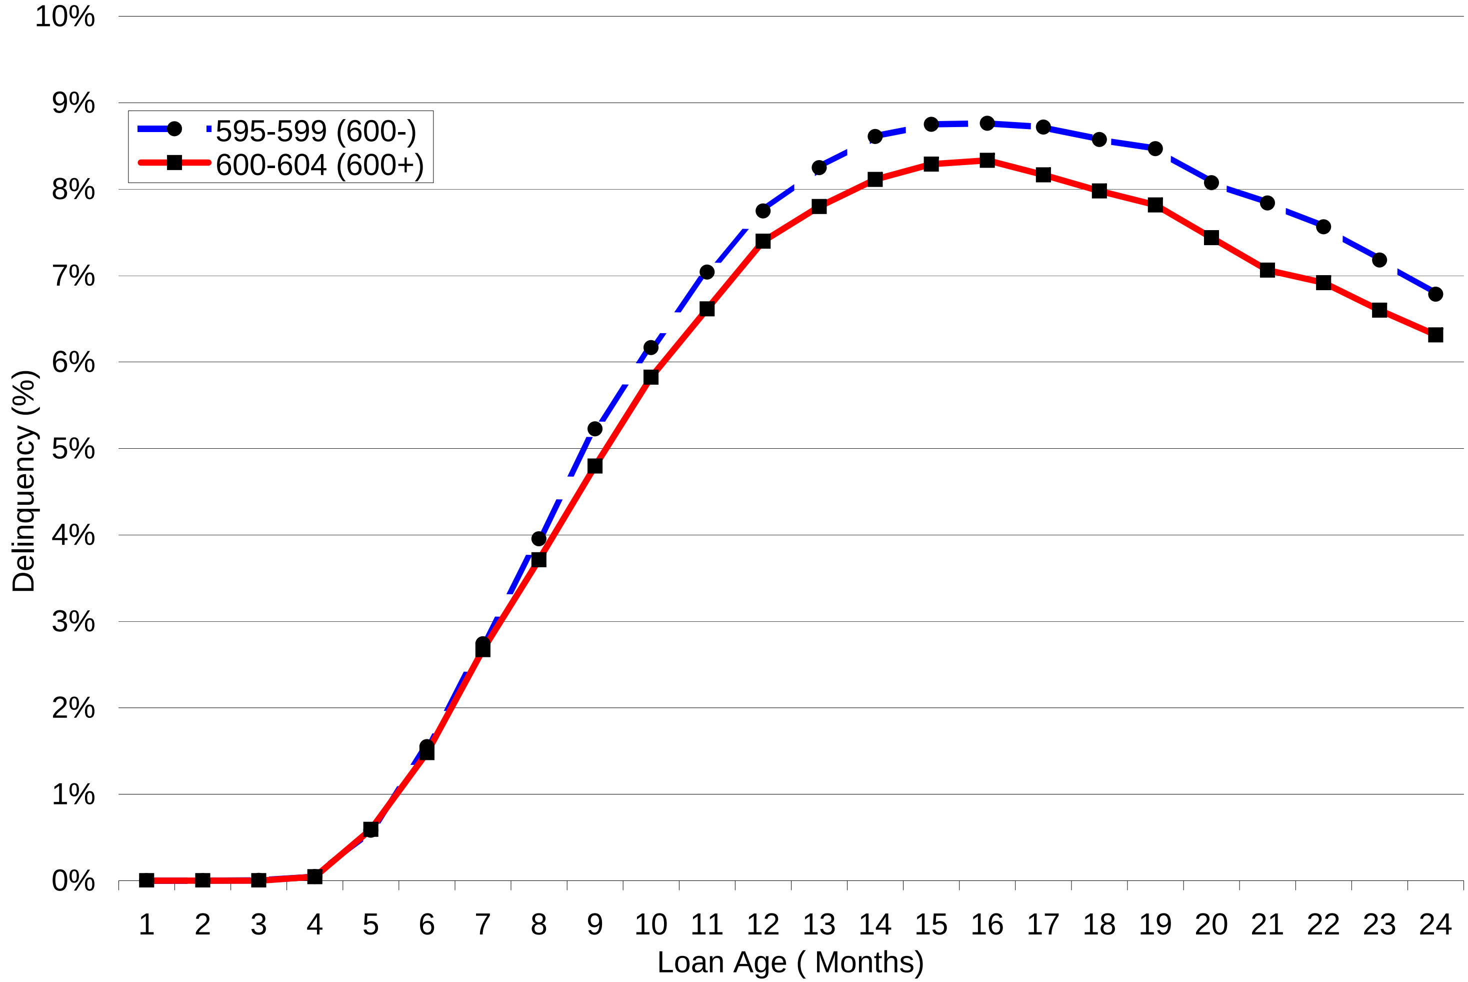

Figure 3.19 presents the percent of full documentation loans (dollar weighted) originated between 2001 and 2006 that became delinquent. We track loans in two FICO buckets – 595-599 (600^{-}) in dotted blue and 600-604 (600^{+}) in red – from their origination date and plot the average loans that become delinquent each month after the origination date. As can be seen, the higher credit score bucket defaults *more* than the lower credit score bucket for post 2000 period. For brevity, we do not report plots separately for each year of origination. The effects shown here in the pooled 2001-2006 plot show up for every year.

Table 3.1: Summary statistics

Information on subprime home purchase loans comes from LoanPerformance. Sample period is 2001-2006. See text for sample selection.

Panel A: Summary Statistics By Year

	Low Documentation			Full Documentation		
	Number of Loans	Mean Loan-To-Value	Mean FICO	Number of Loans	Mean Loan-To-Value	Mean FICO
2001	35,427	81.4	630	101,056	85.7	604
2002	53,275	83.9	646	109,226	86.4	613
2003	124,039	85.2	657	194,827	88.1	624
2004	249,298	86.0	658	361,455	87.0	626
2005	344,308	85.5	659	449,417	86.9	623
2006	270,751	86.3	655	344,069	87.5	621

Panel B: Summary Statistics Of Key Variables

	Low Documentation		Full Documentation	
	Mean	Std. Dev.	Mean	Std. Dev.
Average loan size ($000)	189.4	132.8	148.5	116.9
FICO score	656.0	50.0	621.5	51.9
Loan-to-Value ratio	85.6	9.8	87.1	9.9
Initial Interest Rate	8.3	1.8	8.2	1.9
ARM (%)	48.5	50.0	52.7	49.9
Prepayment penalty (%)	72.1	44.8	74.7	43.4

Table 3.2: Discontinuity in number of low documentation loans

This table reports estimates from a regression which uses the number of low documentation loans at each FICO score as the dependent variable. In order to estimate the discontinuity (FICO$\geq$ 620) for each year, we collapse the number of loans at each FICO score and estimate flexible seventh-order polynomials on either side of the 620 cutoff, allowing for a discontinuity at 620. We report t-statistics in parentheses. Permutation tests, which allow for a discontinuity at every point in the FICO distribution, confirm that jumps for each year are significantly larger than those found elsewhere in the distribution (see Section 3.4.2 for more details). For brevity, we report a permutation test estimate from pooled regressions with time fixed effects removed to account for vintage effects. $FICO = 620$ has the smallest permutation test p-value (and is thus largest outlier) among *all* the visible discontinuities in our sample.

Number of Low Documentation Loans

Year	FICO$\geq$ 620 (β)	t-stat	Observations	R^2	Mean
2001	36.83	(2.10)	299	0.96	117
2002	124.41	(6.31)	299	0.98	177
2003	354.75	(8.61)	299	0.98	413
2004	737.01	(7.30)	299	0.98	831
2005	1,721.64	(11.78)	299	0.99	1,148
2006	1,716.49	(6.69)	299	0.97	903
Pooled Estimate (t-stat) [Permutation Test p-value]			781.87 (4.14) [0.003]		

Table 3.3: Delinquencies in low documentation loans around the credit threshold

In Panel A, we estimate the differences in default rates using a flexible seventh-order polynomial on either side of the 620 cutoff, allowing for a discontinuity at 620. 60+ dollar-weighted fraction of loans defaulted within $10-15$ months is the dependent variable. In Panel B, we present estimates from permutation tests from pooled regressions with time fixed effects removed to account for vintage effects using specification similar to Panel A. Permutation tests confirm that the discontinuity at 620 has the smallest p-value (and is thus largest outlier) in our sample. We use alternative definitions of defaults as the dependent variable. In Panel C, we estimate differences in default rates on either side of the 620 FICO cut off using a logit regression. The dependent variable is the delinquency status of a loan in a given month that takes a value 1 if the loan is classified as under default, as defined in the text. Controls include borrower and loan terms discussed in Section 3.4. We report t-statistics are reported in parenthesis (marginal effects are reported in square brackets).

Panel A: Dollar Weighted Fraction Of Loans Defaulted (60+ Delinquent)

Year	FICO$\geq$ 620 (β)	t-stat	Observations	R^2	Mean
2001	0.005	(0.44)	254	0.58	0.053
2002	0.010	(2.24)	254	0.75	0.051
2003	0.022	(3.47)	254	0.83	0.043
2004	0.013	(1.86)	254	0.79	0.049
2005	0.023	(2.10)	254	0.81	0.078
2006	0.044	(2.68)	253	0.57	0.155

Panel B: Permutation Tests For Alternative Default Definitions
(Pooled 2001-2006 with Time Fixed Effects)

Dependent Variable (Default Definition)	FICO$\geq$ 620 (β)	t-stat	Permutation Test p-value	Obs.	R^2	Mean
60+ ($-weighted)	0.019	(3.32)	0.020	1523	0.66	0.072
90+ ($-weighted)	0.028	(4.67)	0.006	1525	0.70	0.065
Foreclosure+ ($-weighted)	0.025	(6.25)	0.004	1525	0.71	0.048
60+ (unweighted)	0.025	(5.00)	0.004	1525	0.65	0.073

Panel C: Delinquency Status Of Loans

	Pr(Delinquency)=1			
	(1)	(2)	(3)	(4)
FICO$\geq$620	0.12	0.48	0.12	0.48
	[0.004]	[0.011]	[0.004]	[0.011]
	(3.42)	(2.46)	(2.10)	(2.48)
Observations	1,393,655	1,393,655	1,393,655	1,393,655
Pseudo R^2	0.088	0.116	0.088	0.116
Other Controls	Yes	Yes	Yes	Yes
FICO$\geq$620*Other Controls	No	Yes	No	Yes
Time Fixed Effects	No	Yes	No	Yes
Clustering Unit	Loan id	Loan id	Vintage	Vintage
Mean Delinquency (%)	4.45			

Table 3.4: Number of loans and delinquencies in low documentation loans around the credit threshold: Evidence from a natural experiment

This table reports the estimates of the regressions on differences in number of loans and performance of loans around the credit thresholds. We use specifications similar to Table 3.2 in Panel A to estimate the number of loans regressions and Table 3.3 (Panel C) in Panel B to estimate delinquency regressions. We restrict our analysis to loans made in Georgia and New Jersey. *NoLaw* is a dummy that takes a value 1 if the anti-predatory law was not passed in a given year or was amended and a value 0 during the time period when then the law was passed. Permutation tests confirm that the discontinuity in number of loans at 620 when the law is not passed has the smallest p-value (and is thus largest outlier) in the Georgia and New Jersey sample. We report t-statistics in parentheses (marginal effects are reported in square brackets).

Panel A: Number of Low Documentation Loans

Year	FICO$\geq$ 620 (β)	t-stat	Observations	R^2	Mean
During Law	10.71	(2.30)	294	0.90	16
Pre & Post Law	211.50	(5.29)	299	0.96	150

Panel B: Delinquency Status Of Low Documentation Loans

	Pr(Delinquency)=1			
	Entire Period 2001-2006		During Law and Six months After	
	(1)	(2)	(3)	(4)
FICO$\geq$620	-0.91	-0.91	-1.02	-1.02
	[0.043]	[0.043]	[0.030]	[0.030]
	(1.78)	(2.00)	(1.69)	(2.12)
FICO$\geq$620*NoLaw	.88	.88	1.13	1.13
	[0.040]	[0.040]	[0.034]	[0.034]
	(1.90)	(1.94)	(1.79)	(1.93)
Observations	109,536	109,536	14,883	14,883
Other Controls	Yes	Yes	Yes	Yes
FICO$\geq$620* Other Controls	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes
Pseudo R^2	0.06	0.06	0.05	0.05
Clustering Unit	Vintage	Loan id	Vintage	Loan id
Mean Delinquency (%)	6.1		4.2	

Table 3.5: Number of loans and delinquencies in agency (GSE/prime) loans around the credit threshold: Evidence from a natural experiment

This table reports the estimates of the regressions on differences in number of loans and performance of loans around the credit thresholds. The analysis is restricted to prime loans made in Georgia and New Jersey. *NoLaw* is a dummy that takes a value 1 if the anti-predatory law was not passed in a given year or was amended and a value 0 during the time period when then the law was passed. Permutation tests confirm that the discontinuity in number of loans at 620 when the law is not passed or passed is no different from estimated jumps at other locations in the distribution in the Georgia and New Jersey sample. We report t-statistics in parentheses (marginal effects are reported in square brackets).

Panel A: Number of Prime Loans

Year	FICO$\geq$ 620 (β)	t-stat	Observations	R^2	Mean
During Law	4.80	(2.70)	249	0.88	20.30
Pre & Post Law	2.33	(1.02)	268	0.92	22.80

Panel B: Delinquency Status Of Prime Loans

	Pr(Delinquency)=1	
	60+ Delinquent 2001-2006	90+ Delinquent 2001-2006
	(1)	(2)
FICO$\geq$620	-0.026	-0.029
	[0.001]	[0.001]
	(0.19)	(0.10)
FICO$\geq$620*NoLaw	-0.004	-0.003
	[0.0004]	[0.0004]
	(0.03)	(0.05)
Observations	56,300	56,300
Other Controls	Yes	Yes
FICO$\geq$620* Other Controls	Yes	Yes
Time Fixed Effects	Yes	Yes
Clustering Unit	Vintage	Vintage
Pseudo R^2	0.01	0.02
Mean Delinquency (%)	5.2	3.1

Table 3.6: Number of loans and delinquencies around the credit threshold for full documentation loans

This table reports the estimates of the regressions on differences in number of loans and performance of loans around the credit threshold of 600 for full documentation loans. We use specifications similar to Table 3.2 in Panel A to estimate the number of loans regressions and Table 3.3 (Panels A, B and C) in Panels B, C and D to estimate delinquency regressions. Permutation tests confirm that $FICO = 600$ has the smallest permutation test p-value (and is thus largest outlier) among *all* the visible discontinuities in the full documentation loan sample. We report t-statistics in parentheses (marginal effects are reported in square brackets).

Panel A: Number of Full Documentation Loans

Year	FICO≥ 600 (β)	t-stat	Observations	R^2	Mean
2001	306.85	(5.70)	299	0.99	330
2002	378.49	(9.33)	299	0.99	360
2003	780.72	(11.73)	299	0.99	648
2004	1,629.82	(8.91)	299	0.99	1205
2005	1,956.69	(4.72)	299	0.98	1499
2006	2,399.48	(6.97)	299	0.98	1148
Pooled Estimate (t-stat) [Permutation Test p-value]			1241.75 (3.23) [0.000]		

Panel B: Dollar Weighted Fraction Of Loans Defaulted

Year	FICO≥ 600 (β)	t-stat	Observations	R^2	Mean
2001	0.005	(0.63)	250	0.87	0.052
2002	0.018	(1.74)	250	0.87	0.041
2003	0.013	(1.93)	250	0.94	0.039
2004	0.006	(1.01)	254	0.94	0.040
2005	0.008	(1.82)	254	0.96	0.059
2006	0.010	(0.89)	254	0.86	0.116

Panel C: Permutation Tests For Alternative Default Definitions
(Pooled 2001-2006 with Time Fixed Effects)

Dependent Variable (Default Definition)	FICO≥ 600 (β)	t-stat	Permutation Test p-value	Obs.	R^2	Mean
60+ ($-weighted)	0.010	(1.66)	0.240	1512	0.84	0.058
90+ ($-weighted)	0.006	(1.00)	0.314	1525	0.75	0.046
Foreclosure+ ($-weighted)	0.005	(1.25)	0.265	1525	0.77	0.032
60+ (unweighted)	0.011	(1.50)	0.150	1525	0.70	0.056

Panel D: Delinquency Status Of Loans

	Pr(Delinquency)=1			
	(1)	(2)	(3)	(4)
FICO≥600	-.06	-.02	-.06	-.02
	[0.002]	[0.0006]	[0.002]	[0.0006]
	(2.30)	(0.15)	(1.21)	(0.18)
Observations	3,125,818	3,125,818	3,125,818	3,125,818
Pseudo R^2	0.073	0.084	0.073	0.084
Other Controls	Yes	Yes	Yes	Yes
FICO≥600*Other Controls	No	Yes	No	Yes
Time Fixed Effects	No	Yes	No	Yes
Clustering Unit	Loan id	Loan id	Vintage	Vintage
Mean Delinquency (%)	4.54			

Table 3.A1: Loan characteristics around discontinuity in low documentation loans

This table reports estimates from a regression which uses the mean interest rate and LTV ratio of low documentation loans at each FICO score as the dependent variable. In order to estimate the discontinuity (FICO$\geq$ 620) for each year, we collapse the interest rate and LTV ratio at each FICO score and estimate flexible seventh-order polynomials on either side of the 620 cutoff, allowing for a discontinuity at 620. Because the measures of the interest rate and LTV are estimated means, we weight each observation by the inverse of the variance of the estimate. We report t-statistics in parentheses. Permutation tests, which allow for a discontinuity at every point in the FICO distribution, confirm that these jumps are <u>not</u> significantly larger than those found elsewhere in the distribution. For brevity, we report permutation test estimates from pooled regressions (with time fixed effects removed to account for vintage effects) and report them in Table 3.A4.

Low Documentation Loans

	Loan To Value					Interest Rate				
Year	FICO$\geq$ 620 (β)	t-stat	Obs.	R^2	Mean (%)	FICO$\geq$ 620 (β)	t-stat	Obs.	R^2	Mean (%)
2001	0.67	(0.93)	296	0.76	80.3	0.06	(0.59)	298	0.92	9.4
2002	1.53	(2.37)	299	0.91	82.6	0.15	(1.05)	299	0.89	8.9
2003	2.44	(4.27)	299	0.96	83.4	0.10	(1.50)	299	0.97	7.9
2004	0.30	(0.62)	299	0.96	84.5	0.03	(0.39)	299	0.97	7.8
2005	-0.33	(0.96)	299	0.95	84.1	-0.09	(1.74)	299	0.98	8.2
2006	-1.06	(2.53)	299	0.96	84.8	-0.21	(2.35)	299	0.98	9.2

Table 3.A2: Borrower demographics around discontinuity in low documentation loans

This table reports estimates from a regression which uses the mean demographic characteristics of borrowers of low documentation borrowers at each FICO score as the dependent variable. In order to estimate the discontinuity (FICO$\geq$ 620) for each year, we collapse the demographic variables at each FICO score and estimate flexible seventh-order polynomials on either side of the 620 cutoff, allowing for a discontinuity at 620. Because the demographic variables are estimated means, we weight each observation by the inverse of the variance of the estimate. We obtain the demographic variables from Census 2000, matched using the zip code of each loan. Permutation tests, which allow for a discontinuity at every point in the FICO distribution, confirm that these jumps are not significantly larger than those found elsewhere in the distribution. We report t-statistics in parentheses.

Panel A: Percent Black in Zip Code

Year	FICO$\geq$ 620 (β)	t-stat	Observations	R^2	Mean (%)
2001	1.54	(1.16)	297	0.79	11.2
2002	0.32	(0.28)	299	0.63	10.6
2003	1.70	(2.54)	299	0.70	11.1
2004	0.42	(0.53)	299	0.72	12.2
2005	-0.50	(0.75)	299	0.69	13.1
2006	0.25	(0.26)	299	0.59	14.7

Panel B: Median Income in Zip Code

Year	FICO$\geq$ 620 (β)	t-stat	Observations	R^2	Mean (%)
2001	1,963.23	(2.04)	297	0.33	49,873
2002	-197.21	(0.13)	299	0.35	50,109
2003	154.93	(0.23)	299	0.50	49,242
2004	699.90	(1.51)	299	0.46	48,221
2005	662.71	(1.08)	299	0.64	47,390
2006	-303.54	(0.34)	299	0.68	46,396

Panel C: Median House Value in Zip Code

Year	FICO$\geq$ 620 (β)	t-stat	Observations	R^2	Mean (%)
2001	3,943.30	(0.44)	297	0.66	163,151
2002	-599.72	(0.11)	299	0.79	165,049
2003	-1,594.51	(0.36)	299	0.89	160,592
2004	-2,420.01	(1.03)	299	0.91	150,679
2005	-342.04	(0.14)	299	0.93	143,499
2006	-3,446.06	(1.26)	299	0.92	138,556

Table 3.A3: Loan characteristics and borrower demographics around discontinuity in full documentation loans

This table reports the estimates of the regressions on loan characteristics and borrower demographics around the credit threshold of 600 for full documentation loans. We use specifications similar to Tables 3.A1 and 3.A2 for estimation. We report t-statistics in parentheses. Permutation tests, which allow for a discontinuity at every point in the FICO distribution, confirm that these jumps are <u>not</u> significantly larger than those found elsewhere in the distribution. For brevity, we report permutation test estimates from pooled regressions (with time fixed effects removed to account for vintage effects) and report them in Table 3.A4.

Panel A: Loan Characteristics

	Loan To Value					Interest Rate				
Year	FICO$\geq$ 600 (β)	t-stat	Obs.	R^2	Mean (%)	FICO$\geq$ 600 (β)	t-stat	Obs.	R^2	Mean (%)
2001	0.820	(2.09)	299	0.73	85.1	-0.097	(0.87)	299	0.97	9.5
2002	-0.203	(0.65)	299	0.86	85.8	-0.279	(3.96)	299	0.97	8.6
2003	1.012	(3.45)	299	0.95	86.9	-0.189	(3.42)	299	0.99	7.7
2004	0.755	(2.00)	299	0.96	86	-0.244	(6.44)	299	0.99	7.3
2005	0.354	(1.82)	299	0.93	86.2	-0.308	(5.72)	299	0.99	7.7
2006	-0.454	(1.96)	299	0.94	86.7	-0.437	(9.93)	299	0.99	8.6

Panel B: Percent Black in Zip Code

Year	FICO$\geq$ 600 (β)	t-stat	Observations	R^2	Mean (%)
2001	2.32	(2.03)	299	0.86	13.6
2002	-0.79	(1.00)	299	0.82	12.5
2003	0.40	(0.48)	299	0.87	12.5
2004	0.54	(0.96)	299	0.92	12.9
2005	-0.38	(0.85)	299	0.86	13.4
2006	-0.86	(1.40)	299	0.81	14.3

Table 3.A4: Permutation test results for loan characteristics in low and full documentation loans

This table reports the estimates of the regressions on loan characteristics around the credit threshold of 620 for low documentation loans and credit threshold of 600 for full documentation loans. We pool the loans across all years and remove year effects to account for vintage effects. We use specifications similar to Table 3.A1 for estimation. Permutation tests, which allow for a discontinuity at every point in the FICO distribution, confirm that these jumps are <u>not</u> significantly larger than those found elsewhere in the distribution. We report p-values from these tests in the table.

Panel A: Low Documentation Loan Characteristics

	Interest Rate	Loan To Value Ratio	Debt To Income Ratio	Prepayment Penalty	Actual Prepayments	CLTV Ratio
Pooled FICO$\geq$ 620 (β)	0.02	0.54	0.42	-0.016	-0.0004	0.05
t-stat	0.32	1.40	1.25	1.23	0.44	0.73
Permutation Test p-value	0.90	0.46	0.32	0.55	0.84	0.96

Panel B: Full Documentation Loan Characteristics

	Interest Rate	Loan To Value Ratio	Debt To Income Ratio	Prepayment Penalty	Actual Prepayments	CLTV Ratio
Pooled FICO$\geq$ 600 (β)	0.39	-0.26	0.68	0.008	-0.0009	0.17
t-stat	1.63	1.91	1.83	0.73	1.51	0.39
Permutation Test p-value	0.20	0.07	0.11	0.45	0.35	0.62

CHAPTER IV

Can Self-Control Explain Avoiding Free Money? Evidence from Interest-Free Student Loans

> "Although it may be tempting to use student loan money for college football tickets, midnight pizza while cramming for finals, or a Florida spring break trip, try to resist this lure....If you receive a larger loan than you need, the temptation to spend the extra money on "fun" things can be hard or even impossible to resist."

-Dara Duguay, "Spend Student Loans Only on College Expenses" youngmoney.com (money management website for young adults), 2004

4.1 Introduction

This paper uses insights from behavioral economics to explain a particularly bizarre borrowing phenomenon: About one in six undergraduate students who are offered interest-free loans turn them down. The students we observe making these choices are not atypical: Our sample consists of full-time students enrolled at public or private non-profit four-year institutions. Upon filling out the application required for all forms of need-based aid, these students demonstrated sufficient financial need to qualify for interest-free loans sponsored by the federal government.

There are three principal reasons we should be surprised that one-sixth of eligible students turn down the subsidized loans that they are offered. First, these loans do not accrue interest until six months after students leave school. These interest payments represent a direct transfer to the student, and the amount is non-trivial. If a student eligible for the maximum in each year chose to accept the loan each year, with an interest rate of four percent, the government subsidy would be worth more than $1500. The "free money" aspect of below-market interest rates on student loans has long been a part of conventional economic wisdom. One classic undergraduate textbook explains the benefits of a $1000 interest-free loan as follows: "You could at least take the money and put it in a savings bank, where you will earn at least 4 percent per year. Each year you can draw out the $40 interest and throw a big party. Finally...you can draw out the $1,000, plus the last year's interest; repay the $1,000; and have $40 for a last party"(Alchian and Allen 1964). We are unaware, however, of any work that has tried to systematically understand why students do not take advantage of this potential $1500 "gift" from the government.[1]

Seeing students turn down interest-free loans is also surprising because government-sponsored loans help to make increasingly expensive educational costs more affordable. During a period when the return to higher education has dramatically increased, the rising costs of an undergraduate education have far outpaced the increase in the availability of grants and scholarships (Hoxby and Long 1999, Dynarski 2002, Avery and Hoxby 2003). In 1980, 41 percent of financial aid was provided through loans. Today, loans make up 59 percent of federal aid (Douthat 2005) and 45 percent of full-time students borrow to finance their education (NCES 2000, NCES 2002). In the absence of these programs, students would find it costly to borrow against their future earnings due to informational asymmetries between students and private lenders. The federal government has recognized this potential market failure and offers students grants and loans through large-scale programs which provided 90 billion dollars in total aid during the 2004-2005 school year (The College Board 2005).[2] The Stafford Loan Program was originally legislated through the Higher Education Act of 1965, and has been awarded based on a straightforward needs test since 1987 (Mumper and van Ark 1991). By rejecting their government-sponsored loans, students are effectively choosing to borrow at a significantly higher cost, if at all.

Finally, student aid offers are administered under the presumption that students will accept all of their need-based aid. Students must actively reduce or reject any amount they do not wish to borrow. In fact, if a student has borrowed before, she needs to do nothing at all to receive the full amount of any subsidized loan awarded by her financial aid office. As other researchers have shown, there is a significant mental barrier to making decisions which deviate from the default, known as "default bias" (c.f. Choi et al. 2004). In the absence of competing forces, therefore, students should

[1] Note that a student need not plan on "gaming the system" when she borrows for accepting the loan to be a good idea. If there is some uncertainty about the costs she will face over the school year, she may wish to borrow the money as a precautionary measure.

[2] Previous work on the changing nature of the financial aid system focuses on the characteristics of students who default on their loans (Knapp and Seaks 1992, Dynarski 1994). Other authors examine whether the size and type of student loans affect whether and where students enroll (McPherson and Schapiro 1991, van der Klaauw 2002, Kane 2003, Epple, Romano and Sieg 2003). Field (2004) investigates an NYU law school experiment and finds that the decision to enter public-interest law in exchange for a lower debt burden is sensitive to the timing of incurring debt. See also Orfield (1992) for a summary of the policy debate which led to the expansion of student loans in the 1990s.

rarely deviate from the default of accepting all of their need-based aid, including interest-free loans.

While the benefits of subsidized student loans are seemingly unambiguous, borrowing does increase a student's short-term liquidity. As the quotation at the beginning of this section suggests, interest-free loans are a double-edged sword in the hands of an easily tempted consumer. Despite the fact that these loans make it possible to smooth consumption over time, having such a large amount of liquidity can lead to overspending, i.e. consuming more out of current income than an agent with perfect willpower would desire.

We formalize this argument by modeling a college student choosing how much to borrow while in school. We show that a rational agent would not turn down interest-free student loans because doing so requires foregoing a significant government subsidy in addition to limiting future liquidity. We then discuss how rejecting the loan is consistent with models of self-control from the theoretical literature that allow rational consumers to prefer a subset of choices to the complete set. The debt-averse behavior we observe, therefore, may be the optimal choice a forward-thinking student can make knowing that in the following period she will be tempted to overspend.

There are, however, alternative reasons why a potential borrower could make the "wrong" decision. Certainly some students will reject the loan because they do not understand how the subsidy works or do not analyze the decision closely enough.[3] Students may also falsely believe that borrowing through student loan programs will hurt their credit score. In fact, each month while the student is in school the lender reports that the loan account is being paid as agreed, establishing a solid credit history. Apart from these information problems, some students may reject their loans because of the hassle borrowing creates, such as having to keep track of the documents associated with a loan or being required to make a payment each month after graduation.[4] Still others may reject the loans because they have acquired an anti-debt ethic such that indebtedness carries a psychological cost. Survey research in the United Kingdom finds that students who are uncomfortable with debt are less likely to pursue a college education, although they do not attempt to determine the source of this discomfort (Callender and Jackson 2005). Because any of these factors can potentially explain the significant fraction of students who turn down their interest-free loans, we cannot simply interpret high rejection rates as evidence of a self-control motive.

To determine whether self-control plays an important role, the ideal quasi-experimental setting would fix the benefits of borrowing while varying students' exposure to increased liquidity. A feature

[3]See Dynarski and Scott-Clayton (2006) for a discussion of the sometimes overwhelming complexity of the financial aid system.

[4]Another potential reason to turn down student loans is that they are not dischargeable under current bankruptcy law.

of financial aid disbursement does *exactly* this: Although the value of the subsidy is unchanged, needy on-campus students have their loans automatically applied to their educational expenses while similarly needy off-campus students receive a portion of their aid in cash. Comparing the take-up rates of these two groups provides us with a means to test whether self-control motives are responsible for some of the failure in take-up.

However, if students who reject their loans for other reasons tend to live in off-campus housing, this comparison may incorrectly attribute differences in take-up rates to differences in liquidity. To address these selection concerns, we form a difference-in-differences estimator, using students whose liquidity is unaffected by their housing location as a counterfactual. For these students, any loan funds will be applied directly to their tuition bill regardless of where they live. Importantly, each member of the counterfactual group is also eligible for the maximum subsidized loan. If students reject their loans to avoid excess liquidity, the difference between on- and off-campus rejection rates should be much larger for the group who potentially receive their loans in cash.

Our estimates from the 1999-2000 and 2003-2004 waves of the National Postsecondary Student Aid Study support a self-control explanation: Students who would have received cash from their loans turn down the subsidized loan seven percentage points more frequently than similarly needy students who live on-campus. Importantly, there is no significant difference in rejection rates across housing locations for students who would not receive cash regardless of where they live. These difference-in-differences results suggest that the increased liquidity created by living off-campus leads students to reject their loans in higher numbers.

In further support of this hypothesis, we then isolate the variation in living off-campus resulting from supply constraints at the school level. Specifically, we estimate the effect of liquidity on take-up using the university's dormitory capacity (number of beds per student) as an instrument for the housing location decision. To maintain the advantages of the difference-in-differences framework, we also instrument for the interaction of location and loans in excess of tuition, which determines whether the loan is distributed in cash. This exactly identified IV specification (two endogenous regressors and two instruments) thus continues to compare the on/off-campus differences in take-up between students whose loans pay only tuition and students whose loans also pay room and board. In contrast to the OLS results that potentially suffer from endogenous selection into on- or off-campus housing, the IV results isolate the variation in housing location and exposure to liquidity that derives from differences in the supply of on-campus housing units. The IV results complement the earlier findings - again demonstrating a differential willingness to borrow across housing locations, even when controlling for differences in school quality that are correlated with housing capacity. These sets of results are difficult to explain without self-control concerns affecting students' decisions.

These findings provide evidence that consumers choose to limit their available choices in a natural setting, i.e. one not generated by the researcher. While several laboratory and simulation studies have presented evidence consistent with consumers exercising self-control (Wertenbroch, Soman, and Nunes 2001, Ariely and Wertenbroch 2002, Laibson, Repetto and Tobacman 2003, Fernandez-Villaverde and Mukherji 2006), studies using data and situations not generated by the researcher have tended to find evidence of consumers succumbing to the temptation of earlier consumption (Stephens 2003, Shapiro 2005, DellaVigna and Malmendier 2006). In addition, while most field experiments are explicitly designed to hold constant any differences between two choices except for the level of commitment, our results reveal that some consumers are willing to pay a substantial amount of money in order restrict their future decisions.[5] These two features distinguish this study as particularly compelling evidence for the existence and importance of time-inconsistent preferences.

In the next section, we discuss the mechanics of financial aid and emphasize the case when impatient individuals might be most wary of taking out loans. We present a brief model of the financial aid process in Section 4.3 and show how rejecting an interest-free loan, while costly, can effectively serve as a mechanism to regulate impulsive consumption. In Section 4.4, we establish the phenomenon empirically, test the additional predictions suggested by the model, and address selection into off-campus housing. Section 4.5 concludes and discusses potential policy implications.

4.2 Overview of the Financial Aid Process

This section presents a sketch of the process that determines the financial aid a college student receives. Our discussion draws heavily from the Federal Student Aid Handbook published by the Department of Education for use by financial aid professionals (Department of Education 2003).

There are two primary components that determine a student's eligibility for all federal financial aid: a measure of the student's ability to pay, and the costs the student faces at the school where she is enrolled. A student interested in need-based financial aid must first file a Free Application for Federal Student Aid (FAFSA), which collects information on the student and her parents, including

[5]For example, Ashraf, Karlan and Yin (2006) are careful to note the equal interest rates paid in the experimental restricted bank account and the unrestricted account.

the value of their assets and incomes from the previous year.[6] These data are then entered into a federal formula that calculates the Expected Family Contribution (EFC), the dollar amount a family could reasonably be expected to pay for the student's educational costs in the upcoming school year.

In order to qualify for need-based aid, a student must have educational expenses in excess of her EFC.[7] Based on this level of need, the student may be eligible for grant money from the federal or state government or from the institution the student attends. The student may also receive merit-based institutional aid or private scholarships.

Students file the FAFSA to qualify for several types of aid including need-based grants and scholarships. If these forms of aid do not cover the student's entire need, she will be eligible for subsidized Stafford loans.[8] The federal government pays the interest on these loans as long as the student is enrolled at least half-time, and for a period of six months after the student is no longer enrolled. Students can borrow through this program up to a grade-level specific cap: $2,625 for first-year students, $3,500 for second-year students and $5,500 for upper-year students. A student is thus eligible for the maximum loan amount when the difference between her total costs and the sum of the EFC and other grants is greater than the loan limit for her grade level.

After filing a FAFSA, the student receives an award letter from the college or university she is attending (or from the schools to which she applied if she is a first-time student). The letter contains an itemized list of the amounts and types of aid the student has been awarded. Although students maintain the right to reject individual types of aid and even to change individual amounts if they desire, the default choice is to accept the amount of the interest-free loan awarded by the financial aid office.[9] Thus students must intentionally choose not to receive subsidized loans for which they are qualified.

Our identification relies on a feature of the loan disbursement process. Financial aid funds must first be applied to expenses billed directly by the school, including tuition and fees, dorm room

[6]Individuals who have previously been convicted of drug-related felonies, or males over the age of 18 who refuse to register for Selective Service, are generally not eligible for federal financial aid.

[7]The definition of educational expenses is quite broad and includes tuition and fees, room and board, books and supplies, transportation, and other miscellaneous expenses.

[8]Students with exceptional need are also given access to interest-free Perkins loans. The Perkins loan program affects far fewer students than the Stafford program, and the loans are administered by each institution separately, so we choose to focus our attention on the larger federal program. The student may also receive a work-study award which is a promise from the government to pay a portion of the students wages if she obtains employment. Because both of these awards are also need-based, we will be conservative in the empirical section by subtracting Perkins and work-study awards from need before categorizing a student as eligible or ineligible for the Stafford program.

[9]In order to receive their loans, first-time borrowers must sign a Master Promissory Note and receive loan counseling related to borrowing through student loans. A student can currently fulfill both of these requirements online. In subsequent years, the student does not need to take any additional action beyond the normal FAFSA application process to receive the entire amount of loan funds she has been offered.

rental, and cafeteria meal plans. Any aid funding in excess of the school's direct charges is then distributed to the student through a refund check. Because aid eligibility is determined based on the entire student budget and not just tuition, these refund checks are not uncommon.[10]

Figure 4.1 presents the combination of circumstances required to be eligible for a refund check. Students living off-campus receive refund checks earmarked for room and board whenever the sum of their financial aid funds exceeds the cost of tuition and fees. All else equal, students who attend schools with lower tuition or who receive larger grant awards are more likely to be eligible for refunds. On-campus students will never have direct control over the money because it is automatically applied to their educational expenses, including room and board. Thus, the disbursement process creates variation in the short-term liquidity to which students are exposed, even though the financial benefits of the loan are the same. Note that potential endogeneity arises from students' choice whether to live on- or off-campus, which we carefully address in Section 4.4.

4.3 A Self-Control Motive?

With these institutional details in mind, we explore a stylized version of the decision facing an enrolled student who receives an aid award that includes a subsidized Stafford loan. We begin by demonstrating that rejecting the loan cannot be the optimal decision for a student with stable, time-invariant preferences. We then discuss models from the literature under which rejecting a Stafford loan, and thus foregoing the government subsidy, can be utility improving if doing so serves as a constraint on the behavior of an impatient future self.

Students' borrowing and consumption decisions take place over three periods: prior to attendance, during school, and post-graduation. In the initial period, the student is offered a financial aid package, one component of which is a subsidized loan in the amount of $\bar{S}$. She decides whether to accept or reject the loan, and once she has chosen, no future actor can alter this decision.[11] We refer to the actor who makes these decisions as the "Borrower."

In the next period, during school, the student takes her previous loan-taking decisions as given, receives financial aid and other exogenous income (e.g. parental support), and pays tuition. We denote income available after paying tuition as I_2. The student then decides how much to con-

[10]Recent statistics based on administrative data from one of our institutions reveals that 36 percent of aid recipients were issued refund checks.

[11]We discuss the choice as binary, despite the fact that students can choose to borrow only a fraction of the amount they are offered. While we were originally interested in students who took partial loans because they appear to exhibit sophisticated behavior, data limitations do not allow us to distinguish between volitional partial borrowers and students who failed to receive the full amount because they dropped out or graduated. In addition, the structure of the award letter often frames the choice as an all-or-none decision, with the reduction option buried in the fine print.

sume while in school, c_2, and how much to save until the final period. Savings earn a (nominal) interest rate of r per period. We assume that, other than her access to student loans, the student cannot access alternative credit markets.[12] We refer to the actor who makes these decisions as the "Student."

In the final period, post-graduation, the student receives income I_3, repays the principal on any loan she has accepted (the government pays the interest), and consumes the remainder of her income. We refer to this actor as the "Graduate."

The decision whether to borrow is equivalent to a choice between two in-school budget sets, shown in Figure 4.2.[13] If the Borrower chooses to reject the loan, the Student will be faced with the budget set AB. Choosing to borrow provides the Student with budget set CD. Notice that borrowing has two effects on the choice set available to the Student. First, the loan relaxes the Student's credit constraint, allowing her to consume more than I_2 while in school by borrowing against future income, I_3.

Additionally, the government pays the interest that would normally accrue on the loan, $r\bar{S}$. This increase in lifetime income results in a vertical shift upward in the budget set. Proving that a rational student who is not subject to problems of temptation and self-control should strictly prefer to accept the loan requires nothing more than an assumption of locally non-satiated preferences as the choice set available without borrowing is a proper subset of the choice set induced by accepting the loan.

4.3.1 Behavioral motivations

For students who tend to indulge in immediate gratification, however, accepting the loan may not be the optimal choice. In order to take advantage of the government interest payment $r\bar{S}$, the Borrower must relax a constraint on in-school consumption, making it far easier for the Student to overspend.

Dating back to the pioneering work of Strotz (1955), consumer choice theory has developed models which allow for restrictions of future choices to improve lifetime utility. More recently, Laibson (1997), O'Donoghue and Rabin (1999) and Frederick, Loewenstein and O'Donoghue (2002) explore time-inconsistent preferences in which consumers consistently prefer immediate consumption, and discuss the utility improvement that commitment devices can create. The economic theory

[12] While allowing for students to borrow from higher-cost private lenders would add a degree of realism, the intuition underlying this section would be unchanged. We maintain the assumption for expositional simplicity.

[13] For a more general discussion of how self-control concerns cause consumers to prefer a subset of choices to the entire set, see Gul and Pesendorfer (2001).

of self-control in which consumers use personal rules to regulate the impulses of their current and future "selves" began with Thaler and Shefrin (1981), and was further developed by Loewenstein and co-authors (Loewenstein and Thaler 1989, Prelec and Loewenstein 1998) and Benabou and Tirole (2004). In each of these models, some consumers are willing to pay to restrict their future consumption because they anticipate that doing so will help them avoid overspending.

The appendix contains a more formal treatment of one of these potential mechanisms, using quasi-hyperbolic discounting to create a time inconsistency in preferences. Here we highlight the main insights of that analysis. When the Student is relatively patient, it is not worth discarding the government subsidy to change her consumption behavior. The Borrower allows the Student access to the loan funds even though she knows the Student will consume more than she would like her to. For a sufficiently impatient future actor, however, rejecting the loan can be utility improving.

Access to an additional commitment device that would lessen the temptation caused by the increased liquidity could mediate an agent's desire to reject a loan for self-control reasons. For example, as mentioned in the previous section, students who live on-campus will have their loan funds applied directly to their educational expenses. This aspect of the distribution process reduces the Student's short-term liquidity and guarantees that loan funds will be spent on "Borrower-approved" expenses. Either or both of these features may help the borrower control her consumption impulses. In contrast, Students who receive refund checks must manage these funds over the course of the semester, facing the constant lure to spend more out of a temporarily high bank account balance. Because of this increased temptation, we expect students eligible for refund checks to reject their loans more often than their on-campus counterparts.

Thus far we have argued that rejecting an interest-free loan can serve as an effective, albeit costly, method of reining in one's self-control problem and that students who would receive refund checks should be especially likely to make use of this mechanism. An important remaining concern is whether students will or should choose this method over a less-costly alternative. For concreteness, we consider the option of depositing the student loan funds into a certificate of deposit (CD).

There are several reasons to believe that such an alternative does not represent a viable option for our study population. First, depositing the money into a CD requires a significant amount of financial savvy that needy college students likely lack. We argue that these students are "sophisticated" only in that they perceive their own personalities and habits, not that they understand the nuances of the financial system as well as readers of this journal.[14] Additionally, the penalties

[14] According to the 2004 Survey of Consumer Finances, only 12.7% of all households held CD's, with rates much lower for moderate-income families who would likely qualify for financial aid.

on CD's (often only a few months' interest) are probably not large enough to deter students from accessing these accounts for perceived "emergencies." Finally, this strategy is vulnerable to the exact problem it seeks to correct. Because the money does not arrive until after school begins, the Borrower must rely upon the impatient *Student* to deposit the money into the CD. For each of these reasons, rejecting the loan represents the best choice for students aware of their own sufficiently severe self-control problems.

While we have described the model in the context of a within-person principal/agent problem rather than as a parent/child problem, the latter problem would have much the same flavor. In either case, seemingly irrational aversion to debt can be a rational response to a difference in the discount rates between the economic actor who makes the borrowing decision and the actor who makes the consumption decision. There are two reasons to prefer the within-person interpretation. First, subsidized Stafford loans are issued to students without requiring a parent's signature, and the responsibility to repay the loan lies entirely with the student. Second, we control directly for parental financial support in the empirical section, and the results continue to show that students are least likely to borrow when it would result in greater liquidity.

4.4 Empirical Results

4.4.1 Data

We use the 1999-2000 and the 2003-2004 cross-sectional waves of the National Postsecondary Student Aid Study (NPSAS) to investigate the predictions of the self-control model. This unique data source combines administrative financial aid data from the school and from the National Student Loan Data System (NSLDS), information submitted by students and parents on their aid applications, and survey responses from the students during the school year.[15] In addition, the NPSAS contains detailed institution characteristics and individual student information, such as GPA, SAT scores, school location and selectivity, and demographic characteristics.[16]

To focus our analysis on the individuals toward whom the financial aid system is most directly targeted, we restrict our sample to full-time, full-year undergraduate students enrolled at one four-year public or private non-profit institution for a full academic year. The sample includes only those students who applied for financial aid and whose unmet need exceeded the subsidized loan

[15]One additional advantage of the NPSAS is that students make their financial aid decisions prior to being selected into the survey. Thus, there is no additional pressure to make the "correct" decision as a result of being in the study.

[16]We use the restricted version of the data for our analysis. A confidential data license agreement with NCES is required in order to obtain these data.

maximum. Recall that a student must submit an aid application to receive any form of need-based aid. Therefore, some students who were not specifically seeking loan funds will nevertheless receive loan offers.

These selection criteria introduce some heterogeneity by admitting needy students as well as more financially able students at high cost schools to our sample. To mitigate this issue, we further restrict our sample to students who, if they accepted their student loan, would owe no more than an additional $10,000 in tuition.[17] Therefore, within a grade level, all students are eligible for the same interest-free loan amount.[18] These students are usually those considered "representative" needy college students, and those most likely to be burdened with loans upon completion of college.

The fact that about one-sixth of our sample of needy students does not accept interest-free loans is striking. We refer to students who applied for financial aid and who were determined to be eligible for subsidized loans according to the federal formula but who do not receive any loan funds as having rejected the loan. Because this measure is all that our data allow, we were concerned that a significant fraction of our observed rejections might be the result of measurement error where we had incorrectly classified a student as eligible. We asked a senior financial aid administrator at one of our own institutions to compare this number with administrative data. She informed us that 18 percent of Stafford borrowers actively turned down their subsidized loans by logging on to the financial aid system and canceling the loans. Another significant fraction "passively" rejected the loans by failing to return the necessary paperwork for disbursal. This communication suggests that measurement error in eligibility does not comprise a large component of this descriptive statistic. Because we cannot distinguish between active and passive rejection in the data, we refer to all students who applied for aid and who qualified for loans but did not receive the loan as having turned it down.

Table 4.1 provides a descriptive look at the data, emphasizing that a significant fraction of students in each demographic group do not take the loan. The most dramatic differences in take-up rates are by race, where Hispanic and Asian students are nearly twice as likely to turn down the loan as white and African-American students. This could be due to persistent differences in distaste for debt by culture, differences in self-control, information, or other factors. These results serve as

[17] In addition, due to concerns regarding the quality of some responses in the NPSAS, we restrict our dataset further to exclude individuals whose values of student budget and Stafford loan amount were imputed. For similar accuracy concerns with the same variables, we also excluded individuals who were independent or lived with their parents, and students who were not born in the United States. Without these restrictions we risk making significant classification errors.

[18] Note that this sampling frame requires greater unmet need for upperclassmen to be included in the sample than for freshmen and sophomores. We have rerun our analysis using only students who have $5,500 in unmet need regardless of grade level, and the results are qualitatively unchanged.

a reminder that while self-control may be an important determinant of the borrowing decision, it is certainly not the only one. Racial differences in loan rejection are not the emphasis of this paper, though we investigate these racial gaps in more detail below.

Students with high unmet need are much more likely to take the loan. This difference confirms that, on average, the loans are being used by those who need them most. Students from families that earn less than $50,000, roughly the median in our sample, are actually more likely to turn down the loan than are students from wealthier families. Recall, however, that these families are also likely to be eligible for larger grant awards and scholarships. Because family income and need are negatively related as a result of the federal aid formula, it is difficult to determine whether either factor independently drives this result. More generally, this table reveals that students of all types reject the interest-free loans at non-trivial rates.

4.4.2 Evidence for a Self-Control Explanation

Less than full participation in the interest-free loan program is consistent with a number of hypotheses, including taste-based debt aversion, non-pecuniary "hassle" costs of borrowing, or a lack of information. The self-control discussion presented in section 4.3 provides a behavioral reason for rejecting subsidized loans. Unlike the other candidate explanations, this potential motivation provides an additional testable hypothesis: Students should be particularly unwilling to accept their loans when doing so would lead to a larger increase in short-term liquidity.

Recall from Figure 4.1 that some students living off-campus will receive financial aid funds earmarked for room and board in cash which they must manage over the course of the semester, while similarly needy on-campus students will have these funds applied directly to the cost of room and board. If students turn down loans solely because they dislike debt or because they do not understand the benefits, the form in which the loan funds are disbursed should make no difference. In contrast, if self-control is an important factor in students' take-up decisions, they should be especially reluctant to accept the loans if doing so results in a refund check.

The simplest test of the self-control hypothesis is to compare the take-up rates between on-campus and off-campus students whose loan funds would pay for room and board. The results of this comparison are shown in the first column of Table 4.2. Students who live off-campus are 8.0 percentage points less likely to accept their loans than are students in the same financial situation living on campus. As the overall take-up rate is 83 percent, this is nearly a 10 percent difference based solely between students who would receive a refund check and those who would have the funds applied directly to on-campus housing expenses.

However, living off-campus may be associated with greater loan rejection for reasons other than the "refund check" effect. To address this issue, we construct a counterfactual sample who are also eligible for the maximum subsidized loan amount, but whose financial aid benefits do not exceed tuition (see Figure 4.1). By calculating the difference in take-up rates between on- and off-campus students in this group, we can control for differences in take-up related to housing location, but unrelated to the increased liquidity of receiving a refund check. This difference-in-difference specification relies on the assumption that any omitted factors that affect both housing choice and loan take-up are similarly distributed across these two financial aid situations (those of the "treatment" and "control" groups).

As described in more detail below, assignment into the four categories of Figure 4.1 is at least partially endogenous because students have some degree of choice whether to live on- or off-campus at most institutions. To address this issue, we first present results based on this difference-in-difference specification, and then turn to results which exploit institutional housing capacity as an instrument for the ability to live on-campus and uses only variation across schools.

The results of the difference-in-differences specification are presented in the second column of Table 4.2. We estimate linear probability models of the form

$$y_i = \alpha_1(OFFCAMPUS)_i + \alpha_2(ROOMBOARD)_i + \alpha_3(OFFCAMPUS * ROOMBOARD)_i + X_i\beta + \nu_i.$$

The dependent variable is a dummy variable for whether a student accepted his/her interest-free loan (1=accept). The independent variables are indicators for residence (1=off-campus) and for whether loan funds, if accepted, would pay for room and board (1=yes).[19] The interaction of these two variables creates an indicator for whether the loans are distributed in cash (1=refund check). The coefficient on this variable, α_3, is our primary parameter of interest and our measure of the effect of a potential increase in liquidity on loan take-up.[20]

After netting out any on/off-campus differences unrelated to increased liquidity, the resulting coefficient remains strongly negative at 7.3 percentage points (column 2).[21] Notably, the estimated

[19] We limit the sample to students who live either in on-campus housing or off-campus, but not with their parents. Students who live with their parents are typically given a much smaller housing allowance than students living off-campus independently, and thus it is more difficult to determine their eligibility.

[20] We have run linear probability models because our primary parameter of interest is the interaction term, which can be difficult to interpret in probit and other MLE models. Most of the variables we include are categorical, and, as a result, none of the predicted values are greater than one or below zero.

[21] All of the reported standard errors are clustered to allow for arbitrary correlation at the school level.

on/off-campus difference for students who will not receive a refund check in either location is essentially zero (less than half of one percentage point). This reinforces our interpretation of the large and statistically significant difference among the refund-eligible group as the effect of anticipated additional liquidity.

Figure 4.3 presents an important specification check for the difference-in-differences methodology. The graph plots loan acceptance rates against the amount of aid in excess of tuition (including the loan). The continuous lines represent the results of local linear regression smoothing, while the individual points give unweighted averages of bins with a $1000 half-width. The darker lines and points represent the off-campus sample; the lighter plots represent the on-campus sample. For students whose loans pay only tuition (to the left of zero), the relationship between aid and acceptance is quite similar across housing situations. The similarity of the two lines to the left of zero supports using the on-campus students' take-up rate as a counterfactual for the off-campus students had they not been exposed to the prospect of additional liquidity.

The divergence of the acceptance rates between on- and off-campus students occurs immediately to the right of the zero-dollar cutoff as the amount of excess aid increases. These differential trends arise at this point even though the local linear regression does not impose any structure on the shape of the estimated relationship. This result supports the hypothesis that off-campus students are differentially rejecting the loans to avoid receiving large easy-to-spend refund checks. This non-parametric analysis suggests that whether the loan results in a refund is more important than the amount of the loan in determining take-up. Consequently, we continue to report categorical difference-in-differences results rather than results that include the size of the potential refund as a continuous variable.

Adding controls for race, gender and year in school (and thus indirectly for the amount of loan eligibility) in column 3 of Table 4.2 reduces this "refund check" effect only slightly, by less than one half of one percentage point. Importantly, we find little empirical evidence to support a selection across housing options argument as there is no significant difference in take-up between locations for those students whose loans pay only tuition.

These results provide strong evidence that liquidity concerns play a role in determining loan take-up decisions. There are certainly other factors affecting the borrowing choice. In particular, students with a smaller immediate need for funding (accepted loan funds cover more than tuition) are 6.5 percentage points less likely to accept the loan. Nonetheless, students with similar funding needs are an *additional* 7.0 percentage points less likely to take the loans when they would receive cash.

While we have described the student's self-control dilemma in the context of a within-person principal/agent framework rather than as a parent/child problem, the latter problem would have much the same flavor. To address parental influences, we include indicators for whether parents help pay tuition or other financial support, which includes housing expenses, in column 4, our preferred specification. While students whose parents help pay tuition are less likely to take the loan, the "refund check" result remains even after including measures of parental assistance. All else equal, students who would be exposed to additional short-term liquidity are 7.1 percentage points less likely to take the loan.

Table 4.3 presents a series of additional robustness checks on our main result. One alternative explanation for these results is that housing decisions and neediness are serving as proxies for other characteristics of the school the student attends. The NPSAS provides a broad range of school-level characteristics, which we add to our preferred specification from Table 4.2. We include a selectivity index constructed by the NCES based on admission standards and average standardized test scores, the school's Carnegie classification (e.g. research university, liberal arts college etc.), and a measure of the urbanicity of the school's location. The results are largely unchanged; the point estimate on the "refund check" effect falls by only 0.5 percentage points and remains statistically significant.

The second column of this table further addresses the question of whether the "loan funds in cash" effect can best be interpreted as evidence for a self-control or a parental control explanation. We exclude school attributes but include parents' education levels and measures of how involved the student's parents are in financing their educational and living expenses. The additional parental assistance measures are insignificant, and whether at least one parent has some college experience is insignificant as well. That our result still holds suggests that self-control concerns are independent of the role of parents, though a student's parents may still influence her take-up decision.[22] As a further test, we estimated our preferred specification separately on the sample of students who received parental assistance, and then using only those who did not, and generated nearly identical results.[23]

The stability of the point estimates when controlling for a host of potentially confounding

[22]Alternatively, because housing charges are traditionally due at the beginning of a semester, on-campus students may need to borrow in order to pay their housing bill on time. In order to rule out this explanation for the take-up patterns we observe, we restricted our sample only to those students with access to an installment payment plan at their university. The coefficients are slightly smaller (-4.4 percentage points for the interaction in our preferred specification), but remain statistically significant, suggesting that our results are not due to on-campus students' immediate liquidity needs.

[23]We have also estimated similar specifications including a cubic in parental income, which does not substantively affect the point estimate. All results not reported in the tables are available from the authors upon request.

factors suggests that the distribution of these covariates is roughly equal across the four housing location/financial situation categories. Table 4.4 investigates this balance directly. Students from all demographic types can be found in each category, usually in roughly the same proportions. Most of the demographic variation across these categories can be attributed exclusively to housing location or to financial situation, rather than to the "refund check" interaction. Additionally, the table reveals that our comparison group (those whose loans pay only tuition) are only somewhat better off than the group potentially eligible for a refund check. The difference in Estimated Family Contribution is only about $4,500. Our comparison, therefore, is not between poor and non-poor students, but rather between needy and somewhat less needy students.

There are a few cases, however, where demographic variables differ systematically by refund check status. This imbalance presents a challenge to the difference-in-differences specification. As an example, suppose that minority students are especially wary of borrowing to pay for expenses other than tuition. If more of these students live off-campus than on-campus, then the measured interaction would be negative even in the absence of any direct effect of receiving the loan in cash. An analogous argument can be made for any of the other unbalanced covariates.

To address this alternative explanation, we report a series of additional regressions that include the interaction of unbalanced covariates with both the off-campus and "loan pays room and board" dummies. The results are presented in four appendix tables. Appendix Table 4.A1 includes interactions with students' race and gender. Point estimates for the refund check coefficient range from -5.5 percentage points to -7.1 percentage points, and each is statistically significant at the .05 level. The second Appendix Table (4.A2) focuses on grade level. Here the coefficients of interest range from -7.1 percentage points to -8.7 percentage points and all estimates are significant at the 0.01 level.

The final two Appendix Tables (4.A3 and 4.A4) focus on whether differential characteristics of the schools the students attend across the four categories can explain the liquidity effect we identify. Appendix Table 3 interacts the two determinants of refund check status with the type of school, while the final table interacts these variables with measures of cost. In the third table, the interaction ranges from -6.9 percentage points to -7.9 percentage points. Each is significant at the 0.01 level. In the final table, the point estimate falls slightly, and we lose some precision, but the refund check effect is still significant at the .10 level. By helping to rule out simple composition effects, these robustness checks provide further evidence that variation in take-up across these groups is driven by exposure to different levels of short-term liquidity.[24]

[24]As an additional test that the covariates should enter the regression linearly, we estimated a propensity-score

4.4.3 Selection on observables and unobservables

Each of the previous specifications used both between-school and within-school variation in where students fall in the four categories listed in Figure 4.1. The third column of Table 4.3 adds college-specific fixed effects, eliminating the influence of between-school variation on the estimated coefficients. One might prefer this specification because it has the potential to remove unobserved institutional characteristics that affect loan take-up decisions and that are potentially correlated with students' housing and financial aid situations. The resulting point estimate of the coefficient of interest remains negative, but is no longer statistically different from zero, nor statistically different from the results of previous columns.

This difference in magnitude suggests that the majority of the observed "refund check" effect derives from between-school differences in where students live. One interpretation of this fact is that there are unobserved factors (beyond those we directly include as controls) at the school level that tend both to put students into off-campus housing and to reduce the probability that students accept the loan. Under this interpretation, the regressions including school fixed effects provide a more accurate estimate of the refund check effect.

We reject this interpretation because after removing the influence of school-level variables, the remaining within-school variation in housing location largely represents endogenous choices made by students. This endogeneity could easily generate the smaller point estimates observed in the third column of Table 4.3. If, for example, students with self-control problems choose to live on campus as a commitment device to ensure that aid funds go toward appropriate expenses, this selection will tend to minimize differences in on- and off-campus take-up rates, especially within schools.

To address this concern, we estimate just-identified instrumental variables (IV) regressions with two endogenous regressors and two instruments in the same difference-in-differences methodology as used above. We use the fraction of undergraduates that the student's school could place in on-campus housing as an instrument for living off-campus.[25] In addition, we use the interaction of dormitory capacity and whether the student's loan pays for room and board as an instrument for whether a student receives a "refund check." This instrument incorporates variation at both the individual level (financial aid package) and the school level (housing capacity). Importantly, this instrument allows us to avoid making potentially misleading comparisons between on- and

reweighted regression which used the above covariates to predict the likelihood of living off-campus. The resulting estimates were quantitatively similar.

[25] These data do not appear in the NPSAS files. We obtained the data from the Integrated Postsecondary Education Data System (IPEDS), http://nces.ed.gov/ipeds/, and merged them into our dataset using the school identifiers.

off-campus students at the same school, or between students at different institutions. Its estimated effect serves as the coefficient of interest in the specifications that follow.

Beds per student, viewed as a supply restriction for on-campus housing, is an excellent candidate for use as an instrument. With a smaller stock of available housing, some students who otherwise would have lived on-campus will be forced to move out of the dorms. Additionally, while the school's housing stock and its interaction with whether the accepted loan funds pay for room and board are not orthogonal to all other university characteristics that may affect a student's decision to live on-campus, there is little reason to think that these would *directly* affect borrowing decisions.

The required exclusion restriction for these instruments is that they not be correlated with any unobserved determinants of borrowing. For this reason, we explore the first-stage relationships between the two instruments, two potentially endogenous regressors, and observable institution-level characteristics in Table 4.5. These OLS regressions include individual-level characteristics and indicators for each school's urbanicity and Carnegie classification.

The first two columns of Table 4.5 show the strength of the each instrument's first-stage relationship. The strong statistical significance in these regressions eliminates any concern over weak instruments. The supply constraint appears to be binding for many students, as schools with smaller dormitory capacities clearly have more students living off-campus. In addition, students whose loans pay room and board are more likely to face the prospect of a refund check at schools with less capacity to house them on-campus. This combination of results allows us to continue to employ a difference-in-differences framework while relying only on the variation in housing location that is induced by the school's housing capacity.

The third column provides a crucial specification check for this IV estimation strategy by demonstrating that the refund check instrument (the interaction term) and the student's level of financial need (EFC) are uncorrelated. A student's EFC provides a comprehensive measure of her finances and ability to pay for school. The lack of a correlation demonstrates that students who end up in the "refund check" category as a result of housing capacity are similar to the comparison students whose schools have room for them on-campus. The remaining columns present additional specification checks on the interaction instrument using school characteristics. Importantly, there is no evidence that the refund check instrument is correlated with IPEDS measures of institutional quality, affordability, and wealth (endowment). The results contained in this table strongly support the validity of this instrument, although there is of course no way to eliminate every concern over its exogeneity.

Column 1 of Table 4.6 presents the results of the just-identified IV regressions using housing capacity as an instrument for on-campus housing and housing capacity interacted with whether the accepted loan funds pay for room and board as an instrument for the "refund check" effect. The results from this specification continue to demonstrate a strong "refund check" effect and suggest that, if anything, endogenous selection into off-campus housing attenuates the OLS difference-in-differences estimates of the effect of receiving aid funds in cash. Note that the previously reported specification that included school fixed effects effectively eliminated this exogenous variation in exposure to increased liquidity.

To further address concerns of instrument exogeneity, we re-estimated our IV specification including measures of institutional quality and the size of the school's endowment as additional controls. The specification presented in column 1 includes indicators for urbanicity and Carnegie classification, and the results for the "refund check" effect are largely unchanged when we use other measures of school quality, such as the 25th or 75th percentile of incoming students' SAT scores (column 2), or the institution's graduation rate (column 3).[26]

While these checks on exogeneity are not exhaustive, the additional IV specifications support the interpretation that housing capacity affects loan take-up only through the decision whether to live on or off campus. Taken together, the fully interacted OLS and instrumental variables difference-in-difference results in this section indicate that endogenous selection into off-campus housing does not artificially generate differential take-up rates by housing status.

4.4.4 Evidence of Planning Ahead

Thus far we have presented evidence that students who would receive a part of their loan funding in cash are less likely to accept the loan. Viewed through the lens of a self-control model, these empirical results support the hypothesis that students are rejecting their loans as a commitment device against overspending. These results, however, provide no direct evidence that students are rejecting their loans as an optimal forward-looking strategy. Table 4.7 presents an additional test to determine whether students are indeed planning ahead to reject these loans.

We examine whether students' stated borrowing intentions also show a difference by housing location. When filing a FAFSA in the spring, aid applicants must report whether they would like loans included in their financial aid package for the following school year, as well as where they expect to live in the fall. The residential choice and the preference for loans do not directly

[26] In results not shown, for a subset of our sample with available data, we also included the school's endowment; the IV results for this subsample were unaffected.

determine the aid package offered to students in most cases. Further correspondence with the financial aid administrator at one of our own institutions confirmed this fact. The primary concern among administrators is that students will respond that they are not interested in loans in an attempt to secure more grant funding. In addition, subsidized loans are an entitlement program, as students demonstrating eligibility cannot be denied these loans. As a result, subsidized loans are typically included in a qualified student's aid package regardless of that student's expressed interest. Notably, in our sample more than half of the students who report not being interested in loans eventually borrow their entire grade-level maximum. A student's response to this question, therefore, is a useful indicator of her borrowing intentions independent of her eventual borrowing behavior.

Table 4.7 confirms that the larger rejection of loans by potential refund recipients is intentional. These regressions use the student's stated desire for loans as a part of next year's aid package as the dependent variable in the difference-in-differences specification. Students who would get a refund are more likely to report that they are not interested in loans than are other groups. Their desire to avoid borrowing reveals itself not only in their eventual behavior but also in their stated intentions months before their loans are disbursed. These results provide further support for the hypothesis that students' failure to receive interest-free loans reflects the type of forward-thinking decisions made by "sophisticated" consumers aware of their self-control problems.

4.5 Implications for policy and further research

Our analysis suggests that self-control motives play a significant role in students' decisions to reject interest-free loans.[27] Students actively reject loans with an implied government subsidy of up to $1,500 over the student's career. Furthermore, they are particularly less likely to borrow when doing so provides them with a large amount of easy-to-spend cash. This behavior is consistent with the optimal choices of sophisticated economic actors with self-control concerns.[28] Other theories can explain the descriptive results, but no competing theory predicts that students will be exceptionally averse to borrowing when the funding is distributed in cash. These empirical results provide some of the first non-laboratory evidence of consumers choosing to limit their own borrowing and consumption despite the financial costs. In doing so, these results also suggest that many types of

[27] It is likely that we have identified only a portion of the behavior induced by self-control problems. Our estimates omit any effect resulting from students choosing not to apply for aid at all to avoid being faced with the temptation of loan funds.

[28] It is straightforward to show that students who do not anticipate their own impatience will also accept the entire loan for the same reason the rational student does. That their impatience leads to overconsumption is a standard result, and is also easily shown.

behavior previously thought to be "irrationally" debt-averse may, in fact, result from consumers trying to constrain their own impulses.

These results also have important implications for policy decisions related to subsidizing student borrowing. First, we find that a significant fraction of students who reject their loans would have used the money for living expenses rather than for tuition and fees. This finding suggests that the loans end up going to students who actually use the money for school rather than to those who are "gaming the system." However, many needy students, particularly from minority populations, do not reap the benefits of loans that our data suggest they would find unambiguously financially beneficial.

A second policy consideration concerns the efficiency of the design of the current loan system. Recent work on the optimal choice of default rules reveals that setting the default far away from decision-makers' true optima may be welfare improving

Third, potential policy solutions can directly reduce the burden of increased liquidity and increase student participation in this need-based program. For example, aid administrators could offer students access to educational spending accounts similar to flexible spending accounts currently used for medical expenses. Schools could place any aid in excess of tuition into these accounts, and students would need to provide evidence of approved education-related expenses in order to spend these funds. Account balances would earn interest. Upon leaving school, any remaining funds could be applied directly to the student's outstanding loan balances. In this way, all students could receive the benefits of the subsidized loans without needing to manage large increases in liquidity.[29]

The results presented in this paper are some of the first evidence that individuals are willing to pay to restrict their choices in a non-experimental setting, and suggest several avenues of additional research. While our data support an explanation for the surprisingly low take-up of interest-free student loans based on temptation and self-control, further survey work or a randomized experiment could directly confirm this channel. The shift from grant-based aid to loan-based aid may also affect educational decisions more broadly by influencing enrollment and school choice, and the extent of these effects is certainly worthy of further consideration.

Finally, by interacting hyperbolic discounting models with the particular features of this credit market, we have shown that impatience and a need for self-control can induce debt-averse behavior. Although the "rational" choice is less clearly defined in other contexts, we expect that this insight

[29]Alternatively, financial aid offices could offer the "excess" aid in monthly installments. When we compared schools on the semester calendar (where students receive two checks each year) to schools on the quarter calendar, we found that our results were stronger (though not significantly so) when the checks were delivered less frequently. To our knowledge, no university currently offers either educational spending accounts or monthly aid checks.

could help explain unresolved questions in similar economic situations, such as repaying car loans or home mortgages ahead of schedule. Further research is needed to determine how important a role self-control plays in other credit markets.

Figure 4.1: Student circumstances and refund eligibility

	Accepted loan funds:	
Housing location:	Do not pay room and board	Pay room and board
On-campus	Not eligible	Not eligible
Off-campus	Not eligible	Eligible for refund check

Figure 4.2: Choosing to borrow results in a larger choice set.

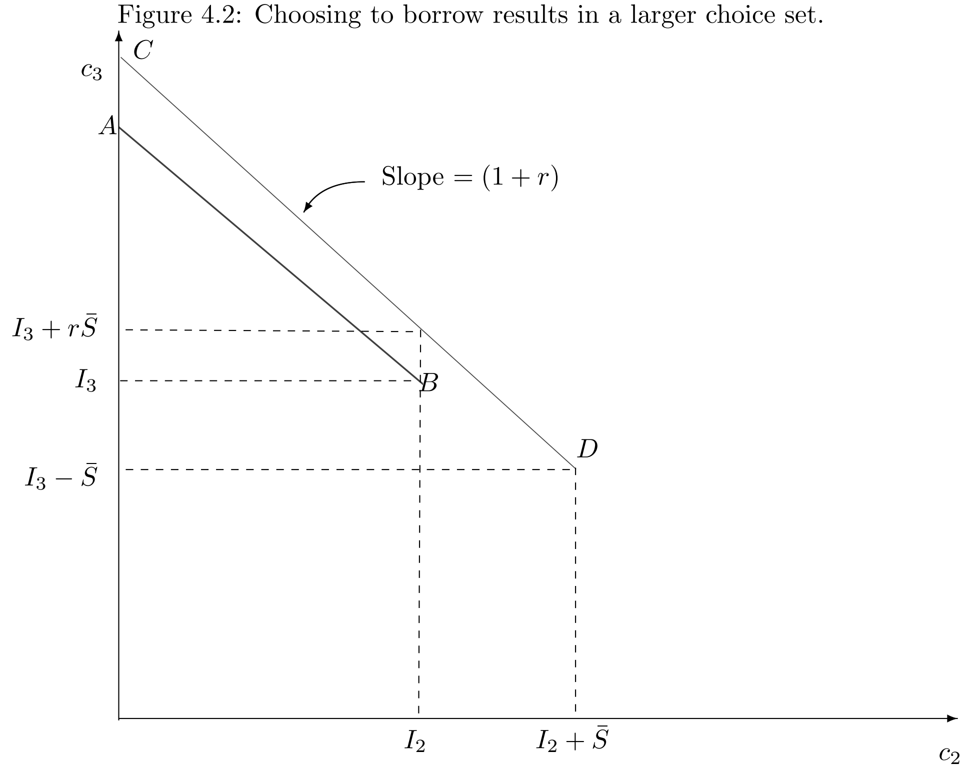

Figure 4.3: Loan acceptance rates by aid in excess of tuition.

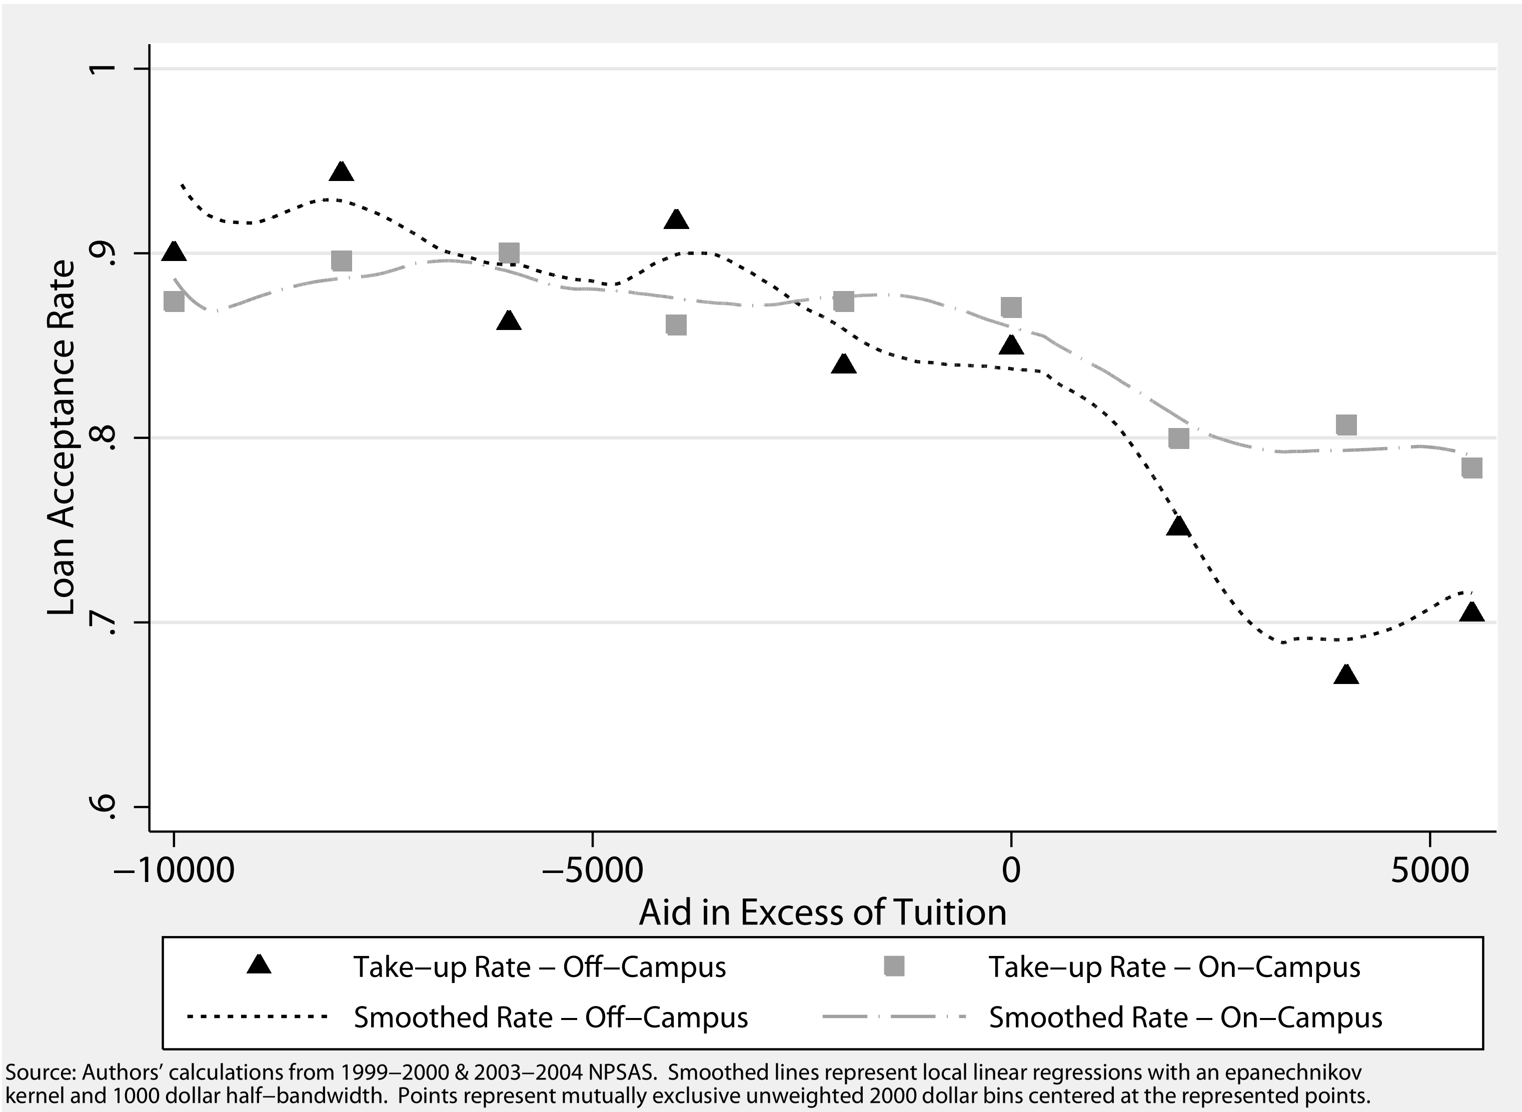

Table 4.1: Descriptive statistics of subsidized Stafford Loan take-up

	% Reject loan	N
Full sample	16.9%	5531
Grade level		
Freshmen	16.0%	2469
Sophmores	20.5%	1049
Juniors	14.5%	825
Seniors	17.0%	1188
Race		
White	15.8%	3938
African-American	15.2%	834
Hispanic	27.1%	402
Asian	27.4%	179
Other (incl. multiple)	14.0%	178
Gender		
Male	17.5%	2416
Female	16.4%	3115
Parental support		
Parents do not pay tuition	16.6%	2309
Parents pay tuition	17.4%	2662
Parental income		
Below $50,000/year	18.5%	3385
Above $50,000/year	14.3%	2146
Cost of attendance after grants/scholarships		
Below median	20.1%	3095
Above median	12.7%	2436
Parental education		
HS degree or less	17.0%	1516
Some college or higher	16.6%	3863
Standardized test scores		
Below median SAT / ACT	14.6%	982
Above median SAT / ACT	19.4%	1008
Survey Year		
1999-2000	16.2%	2170
2003-2004	17.3%	3361

Source: Authors' calculations using the NPSAS 99/00 and 03/04.

Note: We restrict the sample to full-year, full-time, US-born, dependent, undergraduate students at four-year public or private non-profit institutions who do not live with their parents, applied for financial aid, and demonstrated financial need exceeding their grade-level specific loan maximum. We additionally exclude students whose values of student budget and Stafford loan amount are imputed.

Table 4.2: Linear probability models for subsidized Stafford Loan take-up rates by direct access

Dependent variable:	Eligible[a]	Full Sample		
Accept/reject interest-free loans	(1)	(2)	(3)	(4)
Loan funds distributed in cash (offcampus*room and board)[b]	-0.080**	-0.073**	-0.070**	-0.071**
	(0.020)	(0.024)	(0.024)	(0.024)
Lives off-campus, not with parents		-0.004	-0.005	-0.011
		(0.016)	(0.017)	(0.017)
Accepted loan funds pay room and board		-0.066**	-0.065**	-0.072**
		(0.013)	(0.013)	(0.013)
Female			0.013	0.013
			(0.010)	(0.010)
African-American			0.017	0.013
			(0.016)	(0.016)
Asian-American			-0.107**	-0.105**
			(0.036)	(0.036)
Hispanic			-0.089**	-0.093**
			(0.028)	(0.028)
Other race			0.027	0.028
			(0.026)	(0.026)
Parents help pay tuition				-0.043**
				(0.012)
Financial support other than tuition				0.009
				(0.012)
Constant	0.840**	0.888**	0.888**	0.907**
	(0.015)	(0.011)	(0.012)	(0.014)
Controls for grade level (4 categories)	No	No	Yes	Yes
Observations	2771	5531	5531	5531
R^2	0.01	0.02	0.03	0.03

Source: Authors' calculations using the NPSAS 99/00 and 03/04.
Note: Standard errors clustered at the school level in parentheses.
+ significant at 10%; * significant at 5%; ** significant at 1%.
All models include a dummy for survey year.
a. The sample for the first column includes only students who would receive a refund check if they lived off-campus and accepted their loans. We maintain the sample restrictions from Table 1.
b. Loan funds are distributed in cash when the student BOTH lives off-campus and accepted loan funds pay room and board. See Figure 1.

Table 4.3: Subsidized Stafford Loan take-up rates by direct access - Robustness checks

Dependent variable: Accept/reject interest-free loans	(1)	(2)	(3)
Loan funds distributed in cash (offcampus*room and board)[a]	-0.066**	-0.073**	-0.038
	(0.024)	(0.025)	(0.027)
Lives off-campus, not with parents	-0.000	-0.010	0.020
	(0.017)	(0.018)	(0.020)
Accepted loan funds pay room and board	-0.065**	-0.073**	-0.044**
	(0.013)	(0.013)	(0.015)
Constant	0.934**	0.914**	0.863**
	(0.049)	(0.016)	(0.019)
Controls for grade level (4 categories)	Yes	Yes	Yes
Controls for race, gender, and parental assistance	Yes	Yes	Yes
Controls for Carnegie Classification, urbanicity[b]	Yes	No	No
Additional parental controls[c]	No	Yes	No
Institution-level fixed effects	No	No	Yes
Observations	5499	5379	5531
R^2	0.04	0.04	0.24

Source: Authors' calculations using the NPSAS 99/00 and 03/04.

Note: Standard errors clustered at the school level in parentheses. + significant at 10%; * significant at 5%; ** significant at 1%. All models include a dummy for survey year.

a. Loan funds are distributed in cash when the student BOTH lives off-campus and accepted loan funds pay room and board. See Figure 1. We maintain the sample restrictions from Table 1.

b. Includes 3 selectivity dummies, 5 categories of Carnegie classification, and 7 categories for degree of urbanicity. See text for details.

c. Parents' education, whether parents help pay educational expenses, whether parents pay non-housing living expenses.

Table 4.4: Balance of control variables

On-campus/off-campus?	On	On	Off	Off
Does loan cover some room&board expenses?	No	Yes	No	Yes
Borrower gets a refund check?[a]	No	No	No	Yes
Female	55.0%	56.8%	56.3%	58.4%
African-American	12.0%	22.7%	8.9%	12.1%
Hispanic	4.8%	7.4%	6.3%	13.0%
Asian-American	2.5%	3.7%	3.6%	3.7%
Other race	2.9%	3.6%	2.8%	3.3%
Masters U.	20.5%	22.4%	18.9%	17.8%
BA U.	20.0%	13.7%	7.6%	4.6%
Oth U.	40.3%	39.7%	55.3%	50.1%
Research U.	19.2%	24.2%	17.6%	27.2%
Highly selective	24.7%	21.1%	23.5%	19.2%
Moderately selective	71.4%	74.2%	68.5%	75.5%
Not selective	4.0%	4.7%	8.1%	5.3%
High parental education	77.5%	67.2%	71.6%	67.8%
Tuition above median	69.8%	36.4%	61.7%	11.3%
Any parental help with expenses	76.7%	67.2%	62.3%	52.8%
After grant cost of attendance above median	75.2%	10.1%	81.1%	14.4%
Parental income above median	58.6%	19.4%	55.6%	20.3%
Test scores above median	52.9%	49.1%	51.9%	48.4%
Demonstrated need above median	62.4%	28.9%	67.1%	31.6%
Has a credit card	44.2%	44.3%	56.6%	59.3%
Carries credit card balance	21.7%	26.6%	43.4%	42.6%
Average year in school	1.85	1.91	2.75	2.75
Expected Family Contribution (EFC) ($)	5,790	1,976	5,702	2,310
Number of Observations	2152	1755	608	1016

Source: Authors' calculations using the NPSAS 99/00 and 03/04.
a. Loan funds are distributed in cash when the student BOTH lives off-campus and accepted loan funds pay room and board. See Figure 1. We maintain the sample restrictions from Table 1.
Average year in school is coded as 1= Freshman, 2=Sophomore, etc.

Table 4.5: Relationship of instruments with endogenous regressors and other school characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable	Lives off-campus, not with parents	Loan funds distributed in cash	Estimated Family Contribution	25th Percentile of SAT/ACT	75th Percentile of SAT/ACT	Graduation Rate	Endowment	College Affordability Index
Dorm capacity * Funds pay room and board	-0.235**	-0.715**	422.056	-0.602	-0.527	1.377	161.139	0.315
	(0.055)	(0.045)	(425.661)	(0.581)	(0.506)	(3.096)	(100.153)	(1.121)
Dormitory capacity	-0.538**	-0.055**	-734.417+	5.423**	5.139**	32.446**	-33.938	2.431
	(0.048)	(0.018)	(445.003)	(0.761)	(0.686)	(3.574)	(100.124)	(1.888)
Accepted loan funds pay room and board	0.141**	0.598**	-3,200.325**	-0.033	-0.087	-3.523*	-73.663	-0.408
	(0.030)	(0.023)	(225.559)	(0.271)	(0.240)	(1.471)	(44.991)	(0.401)
Female	-0.011	-0.008	-158.680+	-0.189*	-0.178*	0.347	-13.999	0.022
	(0.011)	(0.009)	(95.576)	(0.081)	(0.080)	(0.468)	(10.691)	(0.111)
African American	-0.079**	-0.068**	-1,738.913**	-1.898**	-1.988**	-10.099**	-66.143*	-0.888*
	(0.019)	(0.015)	(129.568)	(0.412)	(0.415)	(2.153)	(28.971)	(0.343)
Asian American	0.002	-0.007	-1,078.019**	1.498**	1.338**	7.628**	63.358	0.841*
	(0.033)	(0.025)	(281.575)	(0.280)	(0.234)	(1.485)	(67.538)	(0.411)
Hispanic	0.021	0.013	-1,309.931**	-0.159	-0.186	-1.154	-42.430+	-0.117
	(0.024)	(0.019)	(161.120)	(0.182)	(0.169)	(1.078)	(25.589)	(0.381)
Other Race	0.001	0.000	-738.733**	-0.242	-0.208	-1.382	-49.531*	-0.273
	(0.036)	(0.030)	(244.987)	(0.286)	(0.268)	(1.321)	(19.777)	(0.434)
Observations	5485	5485	5485	4997	4997	5461	2592	2708
R^2	0.26	0.38	0.28	0.32	0.34	0.37	0.20	0.12

Source: Authors calculations using the NPSAS 99/00 and 03/04, and data from http://nces.ed.gov/ipeds/.

Note: Standard errors clustered at the school level in parentheses. + significant at 10%; * significant at 5%; ** significant at 1%.

Regressions also include an indicator for year of survey, and controls for urbanicity, Carnegie classification, and grade level listed in Table 3.

Table 4.6: Instrumental variable estimates of the liquidity effect

	(1)	(2)	(3)
Loan funds distributed in cash (offcampus*room and board)[a]	-0.171+	-0.181+	-0.186+
	(0.094)	(0.108)	(0.107)
Lives off-campus, not with parents[b]	-0.039	-0.096	-0.046
	(0.097)	(0.109)	(0.118)
Accepted loan funds pay room and board	-0.016	-0.013	-0.010
	(0.028)	(0.030)	(0.029)
Female	0.007	-0.000	-0.001
	(0.010)	(0.011)	(0.011)
African American	0.004	-0.015	-0.009
	(0.015)	(0.020)	(0.020)
Asian American	-0.107**	-0.101**	-0.104**
	(0.036)	(0.036)	(0.036)
Hispanic	-0.054*	-0.055*	-0.057*
	(0.026)	(0.027)	(0.027)
Other Race	0.026	0.028	0.029
	(0.025)	(0.027)	(0.027)
25th percentile for students' SAT/ACT scores		-0.001	-0.007
		(0.006)	(0.006)
75th percentile for students' SAT/ACT scores		-0.007	-0.009
		(0.005)	(0.005)
Graduation rate			0.002*
			(0.001)
First-stage F-test for Instrument Relevance - Off Campus	147.9	76.3	66.8
First-stage F-test for Instrument Relevance - Loan in cash	148.3	109.2	106.0
Observations	5485	4997	4996

Source: Authors calculations using the NPSAS 99/00 and 03/04, and data from http://nces.ed.gov/ipeds/.

Note: Standard errors clustered at the school level in parentheses. + significant at 10%; * significant at 5%; ** significant at 1%. Regressions also include an indicator for year of survey, and controls for urbanicity, Carnegie classification, and grade level listed in Table 3.

a. Loan funds are distributed in cash when the student BOTH lives off-campus and accepted loan funds pay room and board. Whether loan funds are distributed in cash is instrumented by the interaction of accepted loan funds paying for room and board and dorm capacity. See Figure 1. We maintain the sample restrictions from Table 1.

b. Whether student lives on campus is instrumented by dormitory capacity.

Table 4.7: Linear probability models of stated desire for loans

Dependent variable: *Want loans in financial aid package?*	(1)	(2)
Expected loan funds distributed in cash[a]	-0.035	-0.043+
	(0.022)	(0.022)
Expects to live off-campus, not with parents	0.012	-0.008
	(0.014)	(0.015)
Accepted loan funds pay room and board	-0.057**	-0.039**
	(0.011)	(0.011)
Constant	0.946**	0.980**
	(0.009)	(0.012)
Controls for grade level, gender, ethnicity, and parental help	No	Yes
Observations	4509	4509
R^2	0.02	0.06

Source: Authors' calculations using the NPSAS 99/00 and 03/04.
Standard errors clustered at the school level in parentheses. + significant at 10%; * significant at 5%; ** significant at 1%.
a. Loan funds are distributed in cash when the student BOTH lives off-campus and accepted loan funds pay room and board. See Figure 1. We maintain the sample restrictions from Table 1.

4.6 Appendix - Formal Treatment Using Hyperbolic Discounting

This appendix considers the borrowing decision whether to accept or reject a subsidized loan in the context of a student with quasi-hyperbolic preferences. In particular, we assume that the Student has $\beta\delta$ preferences, i.e. her discount rate between any two future periods, s and t is δ^{t-s}, while her discount rate between consumption today ($s = 0$) and consumption in a future period t is $\beta\delta^t$. Both β and δ are smaller than unity.

In deciding whether to accept or reject the loan, the Borrower knows that the Student solves the following problem:

$$\max_{c_2, c_3} u(c_2) + \beta\delta u(c_3)$$

subject to

$$c_2 \leq I_2 + S \tag{4.1}$$

$$c_2(1+r) + c_3 = I_2(1+r) + I_3 + rS \tag{4.2}$$

The loan amount, S, is set by the Borrower as either $\bar{S}$ (accept), or zero (reject), in period 1, prior to the Student's consumption decision. Consumption in each period, c_t, can take on any positive value, and I_t denotes income in period t.

We begin with a helpful proposition.

Proposition IV.1. *All else equal, a Student who is more patient consumes more later, i.e.* $\frac{dc_2^*}{d\beta} \leq 0$ *and* $\frac{dc_3^*}{d\beta} \geq 0$*, with* c_t^* *representing optimal consumption in period* t*. As a result, the Borrower always weakly prefers a more patient Student, i.e. a higher* β*. The Borrower strictly prefers a more patient Student whenever* $\frac{dc_2^*}{d\beta} < 0$*.*

Proof. When the constraint from Equation (4.1) binds, the consumption decision does not change with small changes in β. Thus, when the constraint is binding $\frac{dc_2^*}{d\beta} = \frac{dc_3^*}{d\beta} = 0$. At interior solutions, however, the following Euler equation holds:

$$u'(c_2^*) = (1+r)\delta\beta u'(c_3^*) \tag{4.3}$$

Taking total derivates and rearranging yields

$$\frac{dc_2^*}{d\beta} = \frac{u'(c_3^*)}{u''(c_2^*) + (1+r)u''(c_3^*)} < 0 \tag{4.4}$$

The budget constraint implies that

$$\frac{dc_3^*}{d\beta} = -\frac{dc_2^*}{d\beta}(1+r) > 0 \tag{4.5}$$

When Equation (4.1) does not bind, therefore, more patient Students consume less during school and more after graduation.

The second part of the proposition follows from this result. Letting $U(\cdot)$ denote lifetime utility, and taking the derivative of the Borrower's lifetime utility function with respect to β yields

$$\frac{dU}{d\beta} = u'(c_2^*)\frac{dc_2^*}{d\beta} + \delta u'(c_3^*)\frac{dc_3^*}{d\beta} \tag{4.6}$$

Rearranging, and plugging in $-\frac{dc_2^*}{d\beta}(1+r)$ for $\frac{dc_3^*}{d\beta}$ yields

$$\frac{dU}{d\beta} = (u'(c_2^*) - \delta(1+r)u'(c_3^*))\frac{dc_2^*}{d\beta} \tag{4.7}$$

Recall that c_2^* and c_3^* conform to Equation (4.3) whenever $\frac{dc_2^*}{d\beta} \neq 0$. Thus this can be simplified to

$$= -(1-\beta)\delta(1+r)u'(c_3^*))\frac{dc_2^*}{d\beta}$$

We have already established that $\frac{dc_2^*}{d\beta} \leq 0$; because $0 \leq \beta \leq 1$, it follows immediately that

$$\frac{dU}{d\beta} \geq 0 \tag{4.8}$$

In particular, $\frac{dU}{d\beta} > 0$ when Equation (4.1) does not hold with equality. When the constraint does bind, $\frac{dU}{d\beta} = 0$ holds with equality. □

□

Before characterizing the optimal borrowing behavior, we first determine situations in which the Borrower should always accept the loan.

Proposition IV.2. *Rejecting the loan can only be optimal if doing so causes $c_2 \leq I_2$ to hold with equality.*

Proof. If Equation (4.1) does not hold with equality, then c_2^* and c_3^* follow Equation (4.3). The only effect of taking the loan in this case is that doing so increases lifetime income. Taking the derivative of Equation (4.3) with respect to income yields:

$$u''(c_2^*)\frac{dc_2^*}{dI} = (1+r)\delta\beta u''(c_3^*)\frac{dc_3^*}{dI}$$

Because $u''(\cdot) < 0$, we know that $\frac{dc_2^*}{dI}$ and $\frac{dc_3^*}{dI}$ must be of the same sign. Because they also must sum to one (by Walras' Law), we know that each is positive. Therefore, consumption (and the utility level associated with consumption) increases in both periods when the Student is given access to the loan, and the Borrower will choose to accept the loan. □

□

Proposition IV.3. *If $u(I_2 + \bar{S}) + \delta u(I_3 - \bar{S}) > u(I_2) + \delta u(I_3)$, the Borrower will not reject the loan despite self-control problems.*

Proof. Proposition IV.2 implies that turning down a loan optimally for self-control reasons gives the Borrower utility of $u(I_2) + \delta u(I_3)$. Turning down the loan is therefore optimal only when the consumption choices made by the Student will result in a lower utility. It is trivial to show that when $\beta = 0$, the Student will choose to consume all available income. When the Borrower accepts the loan, therefore, a Student with $\beta = 0$ will consume $(I_2 + \bar{S}, I_3 - \bar{S})$. Proposition IV.1 demonstrates that the Borrower (weakly) prefers the consumption choices of a Student with higher levels of β. Therefore, if $u(I_2 + \bar{S}) + \delta u(I_3 - \bar{S}) > u(I_2) + \delta u(I_3)$ the Borrower should accept the loan for any level of β. □ □

The Borrower rejects the loan in order to limit the Student's consumption to I_2. If accepting the loan would not increase in-school consumption beyond I_2, she should accept the loan and take advantage of the government interest payment. Similarly, if the Borrower prefers shifting $\bar{S}$ in consumption from the post-graduation period to the in-school period over consuming out of current income in both periods, she should always accept the loan. This could occur, for example, if I_3 were significantly higher than I_2.[30]

[30] There are two factors that, in practice, attenuate the gap between I_2 and I_3. First, parental financial support after college is often significantly lower than while in school. Additionally, student loans must be repaid during the early years of a student's earnings trajectory.

We can now characterize the decision rule that determines whether the Borrower accepts or rejects the loan.

Proposition IV.4. *If* $u(I_2+\bar{S})+\delta u(I_3-\bar{S}) < u(I_2)+\delta u(I_3)$, *there exists a unique* $\beta^* \in [0,1]$ *such that* $u(c_2^*(\beta^*)) + \delta u(c_3^*(\beta^*)) = u(I_2) + \delta u(I_3)$. *The Borrower rejects the loan optimally iff* $\beta < \beta^*$.

Proof. We prove this proposition in three parts, beginning with the first claim.

1. At $\beta = 0$, the Student will consume $I_2+\bar{S}$ and leave $I_3-\bar{S}$ for the second period. By assumption, this yields a lower utility than consuming (I_2, I_3). When $\beta = 1$, there is no disagreement between the Student and the Borrower, and the additional income of $r\bar{S}$ ensures that the Borrower prefers (c_2^*, c_3^*) to (I_2, I_3). Combining these results with the (weak) monotonicity results of Proposition IV.1, we can draw Figure 4.A1. This graph establishes the existence and uniqueness of β^*. □

2. We next address the claim that if $\beta > \beta^*$, the Borrower should accept the loan. This result follows directly from Proposition IV.2. Any optimal rejection leads to the consumption pair (I_2, I_3), which is, by the definition of β^*, inferior to the bundle selected by the Student given access to the loan when $\beta > \beta^*$. □

3. Next we prove that $\beta < \beta^*$ is a sufficient condition for rejection to be optimal. To do so, we prove that if the Borrower rejects the loan the Student will consume (I_2, I_3). Because, by the definition of β^*, the Borrower prefers (I_2, I_3) to (c_2^*, c_3^*) the optimal choice is to reject the loan.

   We proceed by contradiction. Suppose this claim is false, i.e. that $c_2^* < I_2$, but that the Borrower prefers (I_2, I_3) to (c_2^*, c_3^*). This preference implies that the Borrower would like to consume more in the first period than would the Student. Recall that for the Borrower, $\beta = 1$ and we have assumed that the Student's $\beta < \beta^* \leq 1$. A preference for more consumption in the first period with a higher β contradicts Proposition IV.1 which established that $\frac{dc_2^*}{d\beta} \leq 0$. $\Rightarrow\Leftarrow$ □

□

When the Student is sufficiently impatient, the Borrower will reject the subsidized loan. Figure 4.A1 demonstrates this cutoff rule and how β^* is determined. The solid line shows the utility the Borrower receives if she accepts the loan and allows the Student to allocate consumption according to her preferences. The dotted line shows the utility the Borrower receives if she rejects the loan and constrains the Student to consume out of current income.

Figure 4.A1: The optimal borrowing decision follows a cutoff rule.

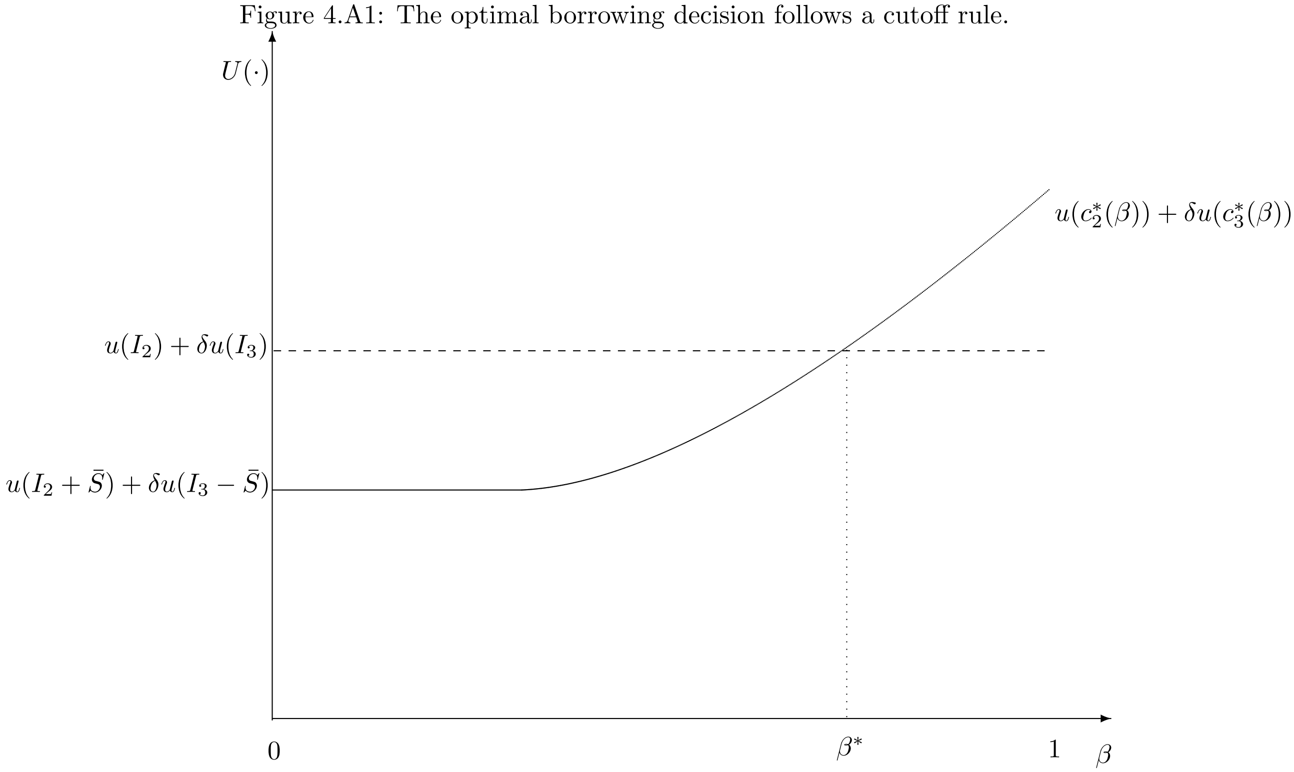

Table 4.A1: Difference-in-differences interacted with race and gender

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Loan funds distributed in cash[a]	-0.072**	-0.060**	-0.087**	-0.070**	-0.071**	-0.074**	-0.076**
	(0.013)	(0.017)	(0.014)	(0.013)	(0.014)	(0.013)	(0.018)
Lives off-campus, not with parents	-0.011	-0.022	-0.012	-0.011	-0.002	-0.010	-0.012
	(0.017)	(0.023)	(0.017)	(0.017)	(0.017)	(0.017)	(0.023)
Accepted loan funds pay room and board	-0.071**	-0.071**	-0.065**	-0.070**	-0.063**	-0.071**	-0.055*
	(0.024)	(0.024)	(0.025)	(0.024)	(0.024)	(0.024)	(0.025)
Female	0.013	0.017	0.013	0.013	0.012	0.013	0.017
	(0.010)	(0.013)	(0.010)	(0.010)	(0.010)	(0.010)	(0.012)
African American	0.013	0.013	-0.041	0.013	0.014	0.014	-0.035
	(0.016)	(0.016)	(0.025)	(0.016)	(0.016)	(0.016)	(0.025)
Asian American	-0.105**	-0.105**	-0.103**	-0.064	-0.106**	-0.105**	-0.063
	(0.036)	(0.036)	(0.036)	(0.053)	(0.036)	(0.036)	(0.054)
Hispanic	-0.093**	-0.093**	-0.091**	-0.093**	-0.013	-0.093**	-0.017
	(0.028)	(0.028)	(0.028)	(0.028)	(0.032)	(0.028)	(0.032)
Other Race	0.028	0.027	0.029	0.027	0.027	0.008	0.009
	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)	(0.037)	(0.037)
Parents help pay tuition	-0.043**	-0.043**	-0.043**	-0.043**	-0.043**	-0.043**	-0.043**
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Financial support other than tuition	0.009	0.009	0.009	0.009	0.008	0.009	0.009
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Female * Off-campus		0.020					-0.025
		(0.024)					(0.038)
Female * Room & board		-0.021					0.093**
		(0.020)					(0.029)
African-American * Off-campus			-0.005				0.024
			(0.037)				(0.024)
African-American*Funds pay room & board			0.092**				-0.023
			(0.028)				(0.020)
Asian-American * Off-campus				-0.014			-0.032
				(0.070)			(0.078)
Asian-American * Room & board				-0.064			-0.051
				(0.065)			(0.065)
Hispanic * Off-campus					-0.140**		-0.151*
					(0.048)		(0.059)
Hispanic * Room & board					-0.036		-0.019
					(0.045)		(0.045)
Other Race * Off-campus						-0.019	-0.041
						(0.065)	(0.066)
Other Race * Room & board						0.045	0.057
						(0.053)	(0.056)
Constant	0.907**	0.905**	0.914**	0.906**	0.904**	0.908**	0.907**
	(0.014)	(0.015)	(0.014)	(0.014)	(0.014)	(0.014)	(0.015)
Observations	5531	5531	5531	5531	5531	5531	5531
R-squared	0.03	0.04	0.04	0.04	0.04	0.04	0.04

Source: Authors' calculations using the NPSAS 99/00 and 03/04.

Note: Standard errors clustered at the school level in parentheses. + significant at 10%; * significant at 5%; ** significant at 1%. All models include a dummy for survey year.

a. Loan funds are distributed in cash when the student BOTH lives off-campus and accepted loan funds pay room and board. See Figure 1. We maintain the sample restrictions from Table 1.

Table 4.A2: Difference-in-differences interacted with grade level

	(1)	(2)	(3)	(4)	(5)
Loan funds distributed in cash[a]	-0.071**	-0.071**	-0.074**	-0.081**	-0.087**
	(0.024)	(0.024)	(0.025)	(0.025)	(0.025)
Lives off-campus, not with parents	-0.011	-0.009	-0.023	-0.006	-0.039
	(0.017)	(0.019)	(0.018)	(0.019)	(0.027)
Accepted loan funds pay room and board	-0.072**	-0.067**	-0.074**	-0.079**	-0.083**
	(0.013)	(0.013)	(0.013)	(0.014)	(0.015)
Female	0.013	0.013	0.013	0.012	0.012
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
African American	0.013	0.014	0.013	0.014	0.013
	(0.016)	(0.015)	(0.016)	(0.015)	(0.015)
Asian American	-0.105**	-0.105**	-0.105**	-0.105**	-0.105**
	(0.036)	(0.036)	(0.036)	(0.036)	(0.036)
Hispanic	-0.093**	-0.093**	-0.092**	-0.092**	-0.089**
	(0.028)	(0.028)	(0.028)	(0.028)	(0.028)
Other Race	0.028	0.028	0.029	0.028	0.031
	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)
Parents help pay tuition	-0.043**	-0.043**	-0.043**	-0.043**	-0.043**
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Financial support other than tuition	0.009	0.009	0.009	0.008	0.007
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Sophomore	-0.039**	-0.024	-0.037*	-0.039*	-0.039*
	(0.015)	(0.020)	(0.015)	(0.015)	(0.020)
Junior	0.034*	0.033*	-0.003	0.035*	-0.013
	(0.016)	(0.016)	(0.023)	(0.016)	(0.023)
Senior	0.007	0.006	0.012	-0.020	-0.028
	(0.016)	(0.016)	(0.016)	(0.020)	(0.021)
Sophomore * Off-campus		-0.009			0.030
		(0.033)			(0.039)
Sophomore * Funds pay room & board		-0.025			-0.005
		(0.027)			(0.029)
Junior * Off-campus			0.068*		0.092*
			(0.029)		(0.036)
Junior * Funds pay room & board			0.019		0.033
			(0.027)		(0.029)
Senior * Off-campus				0.004	0.040
				(0.027)	(0.033)
Senior * Funds pay room & board				0.047*	0.053*
				(0.024)	(0.026)
Constant	0.908**	0.906**	0.912**	0.912**	0.920**
	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
Observations	5531	5531	5531	5531	5531
R-squared	0.03	0.04	0.04	0.04	0.04

Source: Authors' calculations using the NPSAS 99/00 and 03/04.

Note: Standard errors clustered at the school level in parentheses. + significant at 10%; * significant at 5%; ** significant at 1%. All models include a dummy for survey year.

a. Loan funds are distributed in cash when the student BOTH lives off-campus and accepted loan funds pay room and board. See Figure 1. We maintain the sample restrictions from Table 1.

Table 4.A3: Difference-in-differences interacted with institution type

	(1)	(2)	(3)	(4)	(5)
Loan funds distributed in cash[a]	-0.072**	-0.074**	-0.069**	-0.077**	-0.079**
	(0.024)	(0.024)	(0.024)	(0.024)	(0.025)
Lives off-campus, not with parents	-0.002	0.013	-0.007	0.015	0.072+
	(0.017)	(0.019)	(0.018)	(0.023)	(0.038)
Accepted loan funds pay room and board	-0.066**	-0.051**	-0.068**	-0.077**	-0.056*
	(0.013)	(0.015)	(0.014)	(0.016)	(0.028)
Female	0.009	0.009	0.009	0.009	0.009
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
African American	0.009	0.010	0.008	0.009	0.010
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
Asian American	-0.100**	-0.100**	-0.100**	-0.100**	-0.098**
	(0.035)	(0.035)	(0.035)	(0.035)	(0.035)
Hispanic	-0.093**	-0.091**	-0.092**	-0.092**	-0.091**
	(0.027)	(0.028)	(0.027)	(0.027)	(0.027)
Other Race	0.026	0.028	0.026	0.027	0.029
	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)
Parents help pay tuition	-0.044**	-0.044**	-0.044**	-0.044**	-0.044**
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Financial support other than tuition	0.008	0.006	0.008	0.008	0.007
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Type_Masters	0.065**	0.121**	0.065**	0.065**	0.135**
	(0.020)	(0.024)	(0.020)	(0.020)	(0.030)
Type_Baccalaureate	0.101**	0.107**	0.090**	0.102**	0.119**
	(0.020)	(0.021)	(0.023)	(0.021)	(0.030)
Type_Other	-0.047	-0.046	-0.047	-0.052	-0.022
	(0.044)	(0.045)	(0.044)	(0.044)	(0.047)
Masters * Off-campus		-0.075*			-0.131**
		(0.035)			(0.046)
Masters * Funds pay room & board		-0.068**			-0.062+
		(0.023)			(0.032)
B.A. * Off-campus			0.038		-0.035
			(0.031)		(0.044)
B.A. * Funds pay room & board			0.016		0.005
			(0.027)		(0.036)
Other * Off-campus				-0.030	-0.085*
				(0.026)	(0.037)
Other * Funds pay room & board				0.028	0.008
				(0.021)	(0.031)
Constant	0.952**	0.939**	0.954**	0.954**	0.923**
	(0.047)	(0.049)	(0.048)	(0.047)	(0.050)
Observations	5531	5531	5531	5531	5531
R-squared	0.04	0.05	0.04	0.04	0.05

Source: Authors' calculations using the NPSAS 99/00 and 03/04.

Note: Standard errors clustered at the school level in parentheses. + significant at 10%; * significant at 5%; ** significant at 1%. All models include a dummy for survey year.

a. Loan funds are distributed in cash when the student BOTH lives off-campus and accepted loan funds pay room and board. See Figure 1. We maintain the sample restrictions from Table 1.

Table 4.A4: Difference-in-differences interacted with cost

	(1)	(2)
Loan funds distributed in cash[a]	-0.064*	-0.054
	(0.025)	(0.033)
Lives off-campus, not with parents	0.000	-0.015
	(0.017)	(0.033)
Accepted loan funds pay room and board	-0.043**	-0.082**
	(0.015)	(0.021)
Female	0.011	0.010
	(0.011)	(0.011)
African American	0.020	0.020
	(0.017)	(0.017)
Asian American	-0.098**	-0.097**
	(0.036)	(0.035)
Hispanic	-0.098**	-0.097**
	(0.029)	(0.029)
Other Race	0.050*	0.050*
	(0.025)	(0.025)
Parents help pay tuition	-0.050**	-0.050**
	(0.013)	(0.013)
Financial support other than tuition	0.012	0.012
	(0.012)	(0.012)
Tuition above median	0.067**	0.060**
	(0.014)	(0.020)
After grant cost of attendance above median	0.000	-0.034+
	(0.015)	(0.018)
High cost of attendance * Off-campus		0.000
		(0.036)
High cost of attendance * Funds pay room & board		0.094**
		(0.029)
High tuition * Off-campus		0.026
		(0.031)
High tuition * Funds pay room & board		0.009
		(0.024)
Constant	0.862**	0.894**
	(0.018)	(0.021)
Observations	5157	5157
R-squared	0.04	0.04

Source: Authors' calculations using the NPSAS 99/00 and 03/04.
Note: Standard errors clustered at the school level in parentheses. + significant at 10%; * significant at 5%; ** significant at 1%. All models include a dummy for survey year.
a. Loan funds are distributed in cash when the student BOTH lives off-campus and accepted loan funds pay room and board. See Figure 1. We maintain the sample restrictions from Table 1.

CHAPTER V

Conclusion

The first paper demonstrated the important role of job displacement in the household bankruptcy decision. Consistent with predicted filing behavior under persistent income shocks, I find that households in the NLSY are four times more likely to file in the year following job loss, with a smaller but significant response persisting two to three years. Aggregate patterns are also consistent with the model: At the county level, 1000 job losses are associated with 8-11 bankruptcies, the effects also last two to three years, and manufacturing job loss is more likely to induce bankruptcy than non-manufacturing job loss. The results suggest that providing credit counseling to vulnerable households at the time of displacement may be more effective than providing it at the time of bankruptcy.

Theories of financial intermediation suggest that securitization reduces financial intermediaries' incentives to screen borrowers. The second paper, co-authored with Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, examined this question using a unique dataset of securitized subprime mortgage loans. We exploit a rule of thumb in the lending market to generate exogenous variation in the ease of securitization and compare the composition and performance of lenders' portfolios around this threshold. Conditional on securitization, a portfolio that is more likely to be securitized defaults by 20% more than a similar risk profile group. Crucially, these two portfolios have similar observable risk characteristics and loan terms. Our results suggest that securitization can adversely affect lenders' screening incentives.

The third paper, co-authored with Brian C. Cadena, used insights from behavioral economics to offer an explanation for a surprising phenomenon: Nearly 20 percent of undergraduate students who are offered interest-free loans turn them down. We model the financial aid process and show that students facing self-control problems may optimally decline subsidized loans to avoid excessive consumption during school. Using the NPSAS, we investigate students' financial aid situations and subsidized loan take-up decisions and find that students who would receive their loans in cash are significantly more likely to reject the loan. These results suggest that consumers limit their liquidity in economically meaningful situations, consistent with the predictions of the behavioral model.

BIBLIOGRAPHY

Bibliography

Aghion, Philippe, Patrick Bolton, and Jean Tirole, "Exit Options in Corporate Finance: Liquidity versus Incentives," *Review of Finance*, 2004, *8*, 327–353.

Alchian, Armen and William Allen, *University Economics*, Belmont, Calif.,: Wadsworth Pub. Co., 1964.

Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber, "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools," *The Journal of Political Economy*, 2005, *113* (1), 151–184.

Ameriks, John, Andrew Caplin, John Leahy, and Tom Tyler, "Measuring Self-Control," 2004. NBER Working Papers: 10514.

Anderson, Patricia M. and Bruce D. Meyer, "The Extent and Consequences of Job Turnover," *Brookings Papers on Economic Activity: Microeconomics*, 1994, *1994*, 177–248.

Angeletos, George Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, "They Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," in George Loewenstein, Daniel Read, and Roy Baumeister, eds., *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*, New York: Russell Sage Foundation, 2003, pp. 517–43.

Ariely, Dan and Klaus Wertenbroch, "Procrastination, Deadlines, and Performance: Self-Control by Precommitment," *Psychological Science*, 2002, *13* (3), 219–224.

Ashcraft, Adam, Astrid K. Dick, and Donald P. Morgan, "The 2005 Bankruptcy Reform: Means-Testing or Just Mean?," 2007. Federal Reserve Bank of New York working paper.

Ashraf, Nava, Dean Karlan, and Wesley Yin, "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines," *Quarterly Journal of Economics*, 2006, *121* (2), 635–672.

Athreya, Kartik, "Equilibrium Models of Personal Bankruptcy: A Survey," *Federal Reserve Bank of Richmond Economic Quarterly*, 2005, *91*, 73–98.

_ **and Nicole B. Simpson**, "Personal Bankruptcy or Public Insurance?," Federal Reserve Bank of Richmond Working Paper 03-14 2003.

Avery, Christopher and Caroline M. Hoxby, "Do and Should Financial Aid Packages Affect Students' College Choices?," 2003. NBER Working Paper 9482.

Avery, Robert, Raphael Bostic, Paul Calem, and Glenn Canner, "Credit Risk, Credit Scoring and the Performance of Home Mortgages," *Federal Reserve Bulletin*, 1996, *82*, 621–648.

Barron, John M., Michael E. Staten, and Stephanie M. Wilhusen, "The Impact of Casino Gambling on Personal Bankruptcy Filing Rates," *Contemporary Economic Policy*, October 2002, *20* (4), 440–455.

Baumol, William J. and Richard E. Quandt, "Rules of Thumb and Optimally Imperfect Decisions," *American Economic Review*, 1964, *54*, 23–46.

Benabou, Roland and Jean Tirole, "Willpower and Personal Rules," *The Journal of Political Economy*, 2004, *112* (4), 848.

Bhide, Amar, "The Hidden Costs of Stock Market Liquidity," *Journal of Financial Economics*, 1993, *34*, 31–51.

Bostic, Raphael, Kathleen Engel, Patricia McCoy, Anthony Pennington-Cross, and Susan Wachter, "State and Local Anti-Predatory Lending Laws: The Effect of Legal Enforcement Mechanisms," 2007. working paper.

Brown, James N., "Why Do Wages Increase with Tenure? On-the-Job Training and Life-Cycle Wage Growth Observed within Firms," *American Economic Review*, 1989, *79*, 971–991.

Browning, Martin and Thomas F. Crossley, "Shocks, stocks, and socks: Smoothing Consumption over a Temporary Income Loss," McMaster University, Department of Economics, Working Paper, 200307 2003.

_ **and** _ , "The Long-Run Cost of Job Loss as Measured by Consumption Changes," *Journal of Econometrics*, 2008, *145*, 109–120.

Burkhauser, Richard V. and Mary C. Daly, "Employment and Economic Well-being Following the Onset of a Disability: The Role for Public Policy," in J.L. et al. Mashaw, ed., *Disability, Work, and Cash Benefits*, UpJohn Institute for Employment Research, 1996.

Callender, Claire and Jonathan Jackson, "Does the Fear of Debt Deter Students from Higher Education?," *Journal of Social Policy*, 2005, *34*, 509.

Cameron, Stephen and Christopher Taber, "Borrowing Constraints and the Returns to Schooling," 2000. NBER Working Papers: 7761.

Card, David, Alexandre Mas, and Jesse Rothstein, "Tipping and the Dynamics of Segregation," *Quarterly Journal of Economics*, 2008, *128*, 177–218.

Carneiro, Pedro and James J Heckman, "The Evidence on Credit Constraints in Post-secondary Schooling," *Economic Journal*, 2002, *112* (482), 705–34.

Carrington, William J., "Wage Losses for Displaced Workers: Is it Really the Firm that Matters?," *Journal of Human Resources*, 1993, *28*, 435–462.

Chakravarty, Sugato and Eun-Young Rhee, "Factors Affecting an Individual's Bankruptcy Filing Decision," Purdue University working paper 1999.

Charles, Kerwin Kofi and Melvin Stephens, "Job Displacement, Disability, and Divorce," *Journal of Labor Economics*, 2004, *22* (2), 489–522.

Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima, and Jose-Victor Rios-Rull, "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica*, 2007, *75*, 1525–1589.

Choi, James J, David Laibson, Brigitte C Madrian, and Andrew Metrick, "Optimal Defaults," *American Economic Review*, 2003, *93* (2), 180–85.

_ , _ , _ , and _ , "For Better or Worse: Default Effects and 401(k) Savings Behavior," in David A. Wise, ed., *Perspectives on the economics of aging. NBER Conference Report series*, Chicago and London: University of Chicago Press, 2004, pp. 81–121.

Chomsisengphet, Souphala and Anthony Pennington-Cross, "The Evolution of the Subprime Mortgage Market," *Federal Reserve Bank of St. Louis Review*, 2006, *88*, 31–56.

Coase, Ronald, "The Nature of the Firm," *Economica*, 1937, *4*, 386–405.

Coffee, John, "Liquidity Versus Control: The Institutional Investor as Corporate Monitor," *Columbia Law Review*, 1991, *91*, 1277–1368.

Congressional Budget Office, "Subsidy Estimates for Guaranteed and Direct Student Loans," 2005.

_ , "Long-Term Unemployment," Pub. No. 2765 October 2007.

Davis, Stephen J. and James A. Kahn, "Interpreting the Great Moderation: Changes in the Volatility of Economic Activity at the Macro and Micro Levels," NBER working paper No. 14048 2008.

Dawsey, Amanda E. and Lawrence M. Ausubel, "Informal Bankruptcy," 2004. UNC-Greensboro working paper.

Dell'Ariccia, Giovanni, Deniz Igan, and Luc A. Laeven, "Credit Booms and Lending Standards: Evidence from the Subprime Mortgage Market," 2008. working paper.

DellaVigna, Stefano and Ulrike Malmendier, "Paying Not to Go to the Gym," *American Economic Review*, 2006, *93* (3), 694–719.

DeMarzo, Peter and Branko Urosevic, "Optimal Trading and Asset Pricing With A Large Shareholder," *Journal of Political Economy*, 2006, *114*, 774–815.

Demyanyk, Yuliya and Otto Van Hemert, "Understanding the Subprime Mortgage Crisis," 2007. working paper.

Department of Education, "Federal Student Aid Handbook," 2003.

Diamond, Douglas, "Financial Intermediation and Delegated Monitoring," *Review of Economic Studies*, 1984, *51*, 393–414.

_ **and Raghuram Rajan**, "Liquidity Risk, Liquidity Creation and Financial Fragility: A Theory Of Banking," *Journal of Political Economy*, 2003, *109*, 287–327.

Dick, Astrid and Andreas Lehnert, "Personal Bankruptcy and Credit Market Competition," Federal Reserve Bank of New York Staff Report no. 272 January 2007.

DiNardo, John and David S. Lee, "Economic Impacts of New Unionization On Private Sector Employers: 1984-2001," *Quarterly Journal of Economics*, 2004, *119*, 1383–1441.

Domowitz, Ian and Robert L. Sartain, "Determinants of the Consumer Bankruptcy Decision," *Journal of Finance*, 1999, *54* (1), 403–420.

Doms, Mark, Fred Furlong, and John Krainer, "Subprime Mortgage Delinquency Rates," 2007. FRB San Francisco Working Paper.

Douthat, Ross, "Does Meritocracy Work?," *The Atlantic Monthly*, 2005, *296* (4), 120.

Drucker, Steven and Christopher Mayer, "Inside Information and Market Making in Secondary Mortgage Markets," 2007. working paper.

_ **and Manju Puri**, "On Loan Sales, Loan Contracting, and Lending Relationships," 2007. working paper.

Duguay, Dara, "Spend Student Loans Only on College Expenses," *www.youngmoney.com*, 2004.

Dynarski, Mark, "Who Defaults on Student Loans? Findings from the National Postsecondary Student Aid Study," *Economics of Education Review*, 1994, *13* (1), 55–68.

Dynarski, Susan, "The Behavioral and Distributional Implications of Aid for College," *American Economic Review*, 2002, *92* (2), 279–85.

Dynarski, Susan M. and Judith E. Scott-Clayton, "The Cost of Complexity in Federal Student Aid: Lessons from Optimal Tax Theory and Behavioral Economics," 2006. NBER Working Paper 12227.

Efron, Bradley, "Logistic Regression, Survival Analysis, and the Kaplan-Meier Curve," *Journal of the American Statistical Association*, 1988, *83*, 414–425.

Einav, Liran, Mark Jenkins, and Jonathan Levin, "Contract Pricing in Consumer Credit Markets," 2007. working paper.

Elul, Ronel and Narayanan Subramanian, "Forum-Shopping and Personal Bankruptcy," *Journal of Financial Services Research*, 2002, *21*, 233–255.

Epple, Dennis, Richard Romano, and Holger Sieg, "Peer Effects, Financial Aid and Selection of Students into Colleges and Universities: An Empirical Analysis," *Journal of Applied Econometrics*, 2003, *18* (5), 501–25.

Evans, William N. and Julie H. Topoleski, "The Social and Economic Impact of Native American Casinos," NBER working paper No. 9198 2002.

Fallick, Bruce C., "A Review of the Recent Empirical Literature on Displaced Workers," *Industrial and Labor Relations Review*, 1996, *50* (1), 5–16.

Fay, Scott, Erik Hurst, and Michelle J. White, "The Household Bankruptcy Decision," *American Economic Review*, June 2002, *92* (3), 706–718.

Fernandez-Villaverde, Jesus and Arijit Mukherji, "Can We Really Observe Hyperbolic Discounting?," 2006. Univ. of Pennsylvania and the Univ. of Minnesota Working Paper.

Field, Erica, "Educational Debt Burden and Career Choice: Evidence from a Financial Aid Experiment," 2004.

Fishelson-Holstein, Hollis, "Credit Scoring's Role in Increasing Homeownership for Underserved Populations," in "Building Assets, Building Credit: Creating Wealth in Low-Income Communities," Brookings Institution Press, 2005.

Fisher, Jonathan D., "Duration Dependence in Personal Bankruptcy," U.S. Department of Labor Working Paper 359 May 2002.

_ , "The Effect of Unemployment Benefits, Welfare Benefits, and Other Income on Personal Bankruptcy," *Contemporary Economic Policy*, October 2005, *23* (4), 483–492.

Fisher, Jonathan, Larry Filer, and Angela Lyons, "Is the Bankruptcy Flag Binding? Access to Credit Markets for Post-Bankruptcy Households," American Law and Economics Association Annual Meetings Paper 28 2004.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 2002, *40* (2), 351–401.

Gan, Li and Tarun Sabarwal, "A Simple Test of Adverse Events and Strategic Timing Theories of Consumer Bankruptcy," NBER working paper No. 11763 2005.

Glaeser, Edward L. and Joseph Gyourko, "Urban Decline and Durable Housing," *Journal of Political Economy*, 2005, *113*, 345–375.

Gorbachev, Olga, "Did Household Consumption Become More Volatile?," Columbia University working paper 2007.

Gorton, Gary B. and Nicholas S. Souleles, "Special Purpose Vehicles and Securitization," 2005. FRB Philadelphia Working Paper.

Government Accounting Office, "Personal Bankruptcy: Analysis of Four Reports on Chapter 7 Debtors' Ability to Pay," GAO/GGD-99-103 1999.

— , “Value of Credit Counseling Requirement Is Not Clear,” GAO-07-203 2007.

Gramlich, Edward, *Subprime Mortgages: America's Latest Boom and Bust*, The Urban Institute Press, 2007.

Gropp, Reint, John Karl Scholz, and Michelle J. White, “Personal Bankruptcy and Credit Supply and Demand,” *Quarterly Journal of Economics*, 1997, *112* (1), 217–251.

Gross, David B. and Nicholas S. Souleles, “An Empirical Analysis of Personal Bankruptcy and Delinquency,” *Review of Financial Studies*, 2002, *15* (1), 319–347.

Gul, Faruk and Wolfgang Pesendorfer, “Temptation and Self-Control,” *Econometrica*, 2001, *69* (6), 1403–1435.

Han, Song and Wenli Li, “Fresh Start or Head Start? The Effect of Filing for Personal Bankruptcy on Work Effort,” *Journal of Financial Services Research*, 2007, *31*, 123–152.

Handbook', 'Comptroller's, “Asset Securitization,” November 1997.

Harris, Christopher and David Laibson, “Dynamic Choices of Hyperbolic Consumers,” *Econometrica*, 2001, *69* (4), 935–57.

Heckman, J. and B. Singer, “A Method for Minimizing the Impact of Distributional Assumptions in the Econometric Models for Duration Data,” *Econometrica*, 1984, *52*, 271–320.

Hildreth, Andrew K. G., Till M. von Wachter, and Elizabeth Weber Handwerker, “Estimating the 'True' Cost of Job Loss: Evidence using Matched Data from California 1991-2000,” Columbia University working paper 2005.

Himmelstein, David U., Elizabeth Warren, Deborah Thorne, and Steffie Woolhander, “MarketWatch: Illness and Injury as Contributors to Bankruptcy,” *Health Affairs*, 2005, *web exclusive*, 63–73.

Holloway, Thomas, Gregor MacDonald, and John Straka, “Credit Scores, Early-Payment Mortgage Defaults, and Mortgage Loan Performance,” 1993. Freddie Mac working paper.

Holmstrom, Bengt and Jean Tirole, “Financial Intermediation, Loanable funds, and The Real Sector,” *Quarterly Journal of Economics*, 1997, *52*, 663–692.

Hoxby, Caroline M. and Bridget Terry Long, “Explaining Rising Income and Wage Inequality Among the College Educated,” 1999. NBER Working Paper 9482.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan, "Earnings Losses of Displaced Workers," *American Economic Review*, 1993, *83* (4), 685–709.

Jr., Robert E. Lucas, "Econometric Policy Evaluation: A Critique," in "The Phillips Curve and Labor Markets, Carnegie-Rochester Conferences on Public Policy" 1976.

Kane, Thomas J, "A Quasi-Experimental Estimate of the Impact of Financial Aid on College-Going," 2003. NBER Working Papers: 9703.

Kashyap, Anil K. and Jeremy C. Stein, "What Do a Million Observations on Banks Have To Say About the Monetary Transmission Mechanism?," *American Economic Review*, 2000, *90*, 407–428.

Keys, Benjamin J., "Trends in Income and Consumption Volatility, 1970-2000," in D. Joliffe and J. Ziliak, eds., *Income Volatility and Food Assistance in the United States*, Upjohn Institute for Employment Research, 2008.

_ , Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, "Financial Regulation and Securitization: Evidence from Subprime Mortgage Loans," *Journal of Monetary Economics*, 2009, *56.*

Kletzer, Lori G., "Job Displacement," *Journal of Economic Perspectives*, 1998, *12* (1), 115–136.

_ and Robert W. Fairlie, "The Long-Term Costs of Job Displacement for Young Adult Workers," *Industrial and Labor Relations Review*, 2003, *56* (4), 682–698.

Knapp, Laura Greene and Terry G. Seaks, "An Analysis of the Probability of Default on Federally Guaranteed Student Loans," *The Review of Economics and Statistics*, 1992, *74* (3), 404.

Kornfeld, Warren, "Testimony before the Subcommittee on Financial Institutions and Consumer Credit, U.S. House of Representatives, May 8," 2007.

Korniotis, George, "Rewards from Savings, Debt Aversion, and Mental Accounts," 2004. Manuscript, University of Notre Dame.

Kowalewski, K.J., "Personal Bankruptcy: Theory and Evidence," *Federal Reserve Bank of Cleveland Economic Review*, 1982, *Spring*, 1–29.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, *112* (2), 443–77.

_ , Andrea Repetto, and Jeremy Tobacman, "A Debt Puzzle," in Philippe Aghion, Roman Frydman, Joseph Stiglitz, and Michael Woodford, eds., *Knowledge, Information, and Expectations in*

Modern Economics: In Honor of Edmund S. Phelps, Princeton, NJ: Princeton University Press, 2003, pp. 228–66.

Lawrence, Emily C., "Consumer Default and the Life Cycle Model," *Journal of Money, Credit, and Banking*, 1995, *27* (4), 939–954.

Livshits, Igor, James MacGee, and Michele Tertilt, "Accounting for the Rise in Consumer Bankruptcies," NBER working paper no. 13363 September 2007.

_ , _ , and _ , "Consumer Bankruptcy: A Fresh Start," *American Economic Review*, 2007, *97* (1), 402–418.

Loewenstein, George and Richard H Thaler, "Intertemporal Choice," *Journal of Economic Perspectives*, 1989, *3* (4), 181–93.

Loutskina, Elena, "Does Securitization Affect Bank Lending: Evidence from Bank Responses to Funding Shocks," 2006. working paper.

_ and Philip Strahan, "Securitization And The Declining Impact Of Bank Finance On Loan Supply: Evidence From Mortgage Acceptance Rates," 2007. working paper.

Lucas, Douglas, Laurie Goodman, and Frank Fabozzi, *Collateralized Debt Obligations: Structures and Analysis*, Wiley Finance, 2006.

Luckett, Charles A., "Personal Bankruptcies," *Federal Reserve Bulletin*, 1988, *74* (9), 591–603.

_ , "Personal Bankruptcies," in T. Durkin and M. Staten, eds., *The Impact of Public Policy on Consumer Credit*, Boston: Kluwer Academic Publishers, 2002.

Maug, Ernst, "Large Shareholders as Monitors: Is There a Tradeoff between Liquidity and Control?," *Journal of Finance*, 1998, *53*, 65–98.

McCall, Brian and John McCall, *The Economics of Search*, Routledge, 2007.

McPherson, Michael S. and Morton Owen Schapiro, "Does Student Aid Affect College Enrollment? New Evidence on a Persistent Controversy," *The American Economic Review*, 1991, *81* (1), 309.

Meyer, Bruce D., "Unemployment Insurance and Unemployment Spells," *Econometrica*, 1990, *58*, 757–782.

Mian, Atif and Amir Sufi, "The Consequences of Mortgage Credit Expansion: Evidence from the 2007 Mortgage Default Crisis," 2008. working paper.

Moffitt, Robert A. and Peter Gottschalk, "Trends in the Transitory Variance of Earnings in the United States," *Economic Journal*, 2002, *112*, C68–C73.

Morrison, Alan D., "Credit Derivatives, Disintermediation and Investment Decisions," *Journal of Business*, 2005, *78*, 621–647.

Mullainathan, Sendhil and David Scharfstein, "Do Firm Boundaries Matter?," *American Economic Review Papers and Proceedings*, 2001, *91*, 195–199.

Mumper, Michael and Pamela Vander Ark, "Evaluating the Stafford Student Loan Program: Current Problems and Prospects for Reform," *The Journal of Higher Education*, 1991, *62* (1), 62–78.

Murphy, Kevin M. and Robert Topel, "Unemployment and Nonemployment," *American Economic Review Papers and Proceedings*, May 1997, *87* (2), 295–300.

Musto, David K., "The Reacquisition of Credit Following Chapter 7 Personal Bankruptcy," Wharton Financial Institutions Center Working Paper 99-22 1999.

Narajabad, Borghan N., "Information Technology and the Rise of Household Bankruptcy," Rice University working paper 2007.

National Center for Educational Statistics, "Trends in Undergraduate Borrowing: Federal Student Loans in 1989-90, 1992-93, and 1995-96," *Statistical Analysis Reports*, 2000.

_ , "Student Financing of Undergraduate Education: 1999-2000," 2002.

_ , "Characteristics of Undergraduate Borrowers: 1999-2000," 2003.

O'Donoghue, Ted and Matthew Rabin, "Doing it now or doing it later," *American Economic Review*, 1999, *89* (1), 103–124.

Orfield, Gary, "Money, Equity, and College Access," *Harvard Educational Review*, 1992, *62* (3), 337.

Parlour, Christine and Guillaume Plantin, "Loan Sales and Relationship Banking," *Journal of Finance*, forthcoming.

Pavan, Marina, "Consumer Durables and Risky Borrowing: the Effects of Bankruptcy Protection," Boston College Working Papers in Economics No. 573 2005.

Pennacchi, George, "Loan Sales and the Cost of Bank Capital," *Journal of Finance*, 1988, *43*, 375–396.

Petersen, Mitchell A. and Raghuram G. Rajan, "Does Distance Still Matter? The Information Revolution in Small Business Lending," *Journal of Finance*, 2002, *57*, 2533–2570.

Prelec, Drazen and George Loewenstein, "The Red and the Black: Mental Accounting of Savings and Debt," *Marketing Science*, 1998, *17* (1), 4–28.

Shapiro, Jesse M., "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle," *Journal of Public Economics*, 2005, *89* (2-3), 303–25.

Sheather, Simon and Chris Jones, "A Reliable Data-Based Bandwidth Selection Method for Kernel Density Estimation," *Journal of the Royal Statistical Society*, 1991, *53*, 683–690.

Shin, Donggyun and Gary Solon, "Trends in Men's Earnings Volatility: What Does the Panel Study of Income Dynamics Show?," Michigan State University working paper 2008.

Stein, Jeremy C., "Information Production and Capital Allocation: Decentralized versus Hierarchical Firms," *Journal of Finance*, 2002, *57*, 1891–1921.

Stephens, Melvin, ""3rd of tha Month": Do Social Security Recipients Smooth Consumption between Checks?," *American Economic Review*, 2003, *93* (1), 406–22.

Stephens, Melvin Jr., "The Long-Run Consumption Effects of Earnings Shocks," *Review of Economics and Statistics*, 2001, *83* (1), 28–36.

Stevens, Ann Huff, "Persistent Effects of Job Displacement: The Importance of Multiple Job Losses," *Journal of Labor Economics*, 1997, *15* (1), 165–188.

Stiglitz, Joseph, "Houses of Cards," The Guardian, October 9, 2007.

Strotz, Robert H., "Myopia and Inconsistency in Dynamic Utility Maximization," *The Review of Economic Studies*, 1955, *23* (3), 165–180.

Sufi, Amir, "Information Asymmetry and Financing Arrangements: Evidence from Syndicated Loans," *Journal of Finance*, 2006, *62*, 629–668.

Sullivan, Daniel and Till von Wachter, "Mortality, Mass-Layoffs, and Career Outcomes: An Analysis using Administrative Data," NBER working paper No. 13626 November 2007.

Sullivan, James X., "Borrowing During Unemployment: Unsecured Debt as a Safety Net," *Journal of Human Resources*, 2008, *43* (2), 383412.

Sullivan, Teresa A., Elizabeth Warren, and Jay Lawrence Westbrook, *As We Forgive Our Debtors: Bankruptcy and Consumer Credit in America*, BeardBooks, 1999.

_ , _ , **and** _ , "Less Stigma or More Financial Distress: An Empirical Analysis of the Extraordinary Increase in Bankruptcy Filings," *Stanford Law Review*, 2006, *59*, 213–256.

Temkin, Kenneth, Jennifer Johnson, and Diane Levy, "Subprime Markets, The Role of GSEs and Risk-Based Pricing," 2002. Prepared for U.S. Department of Housing and Urban Development Office of Policy Development and Research.

Thaler, Richard H and H M Shefrin, "An Economic Theory of Self-Control," *Journal of Political Economy*, 1981, *89* (2), 392–406.

The College Board, "Trends in Student Aid 2005," 2005.

Topel, Robert, "Specific Capital and Unemployment: Measuring the Costs and Consequences of Job Loss," *Carnegie-Rochester Conference Series on Public Policy*, 1990, *33*, 181–214.

Valletta, Robert G., "Rising Unemployment Duration in the United States: Causes and Consequences," Federal Reserve Bank of San Francisco working paper May 2005.

van der Klaauw, Wilbert, "Estimating the Effect of Financial Aid Offers on College Enrollment: A Regression-Discontinuity Approach," *International Economic Review*, 2002, *43* (4), 1249–87.

Wang, Hung-Jen and Michelle J. White, "An Optimal Personal Bankruptcy Procedure and Proposed Reforms," *Journal of Legal Studies*, 2000, *29* (1), 255–286.

Warren, Elizabeth and Amelia Warren Tyagi, *The Two-Income Trap: Why Middle-Class Mothers and Fathers are Going Broke*, Basic Books, 2003.

Wertenbroch, Klaus, Dilip Soman, and Joseph Nunes, "Debt aversion as self-control: Consumer self-management of liquidity constraints," 2001. INSEAD Working Paper.

White, Michelle J., "Why Don't More Households File for Bankruptcy?," *Journal of Law, Economics, and Organization*, 1998, *14* (2), 205–231.

_ , "Bankruptcy Reform and Credit Cards," *Journal of Economic Perspectives*, 2007, *21*, 175–199.

Zagorsky, Jay L., "Young Baby Boomers' Wealth," *Review of Income and Wealth*, June 1999, *45* (2), 135–156.

— , "Do you have to be smart to be rich? The impact of IQ on wealth, income and financial distress," *Intelligence*, 2007, *35*, 489–501.

— **and Lois R. Lupica**, "A Study of Consumers' Post-Discharge Finances: Struggle, Statis, or Fresh Start?," *ABI Law Review*, 2008, *16*, 1–37.

Zeldes, Stephen P, "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 1989, *97* (2), 305–46.