Essays on Household Finance and Small Business Credit

by

Olivia S. Kim

B.A., Smith College, 2011
S.M., Massachusetts Institute of Technology, 2020

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Management

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

Author . . . .
Department of Management
April 26, 2022

Certified by . . . .
Jonathan Parker
Robert C. Merton (1970) Professor of Financial Economics
Thesis Supervisor

Accepted by . . . .
Catherine Tucker
Sloan Distinguished Professor of Management
Professor, Marketing
Faculty Chair, MIT Sloan PhD Program

Essays on Household Finance and Small Business Credit

by

Olivia S. Kim

Submitted to the Department of Management
on April 26, 2022, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Management

Abstract

Chapter 1 examines whether closing disparities in credit access between spouses can help reduce consumption inequality in the household. The 2013 reversal of the Truth-in-Lending Act increased the borrowing capacity of secondary earners in equitable-distribution states but not in community-property states, where division-of-property laws superseded the policy change. Using a matched difference-in-differences design and administrative financial-transaction records measuring the credit and consumption of each spouse, I show that this reversal closed the credit gap between spouses by increasing secondary earners' credit card limits. In turn, spouses shared consumption more equally, reducing their pre-reversal consumption gap. Delinquency rates were not measurably impacted, suggesting that household financial standing did not worsen. These results are consistent with a model of joint decision-making under limited commitment, in which credit causes a shift in marital bargaining power.

Chapter 2 explores the investment decisions of small business owners when their child goes to college using the linked financial accounts of small businesses and their owners. By comparing small business owner households with college-entering aged children to otherwise similar households with near college-entering aged children, I show that small business owners respond to the increase in education spending by downsizing business production and liquidating the business. These results suggest that business owners' family financial decisions affect the real economy as business owners struggle to separate business capital demands from personal finances.

Joint work with Natalie Cox and Constantine Yannelis in Chapter 3 uses notches in the loan guarantee rate schedule for Small Business Administration loans to estimate the elasticity of bank lending volume to loan guarantees. We show significant bunching in the loan distribution on the side of the size threshold that carries a more generous loan guarantee. The excess mass implies that increasing guarantee generosity by one percentage point of loan principal would increase per-loan lending volume by $19,000. Placebo results indicate that bunching disappears when the guarantee notch is eliminated. We conclude that federal guarantee programs have the potential to increase lending levels when borrowing is inefficiently low.

Thesis Supervisor: Jonathan Parker
Title: Robert C. Merton (1970) Professor of Financial Economics

Acknowledgments

I owe a tremendous debt to my thesis advisors, Jonathan Parker, Antoinette Schoar, Josh Angrist, and Frank Schilbach. As my dissertation chair, Jonathan has dedicated countless hours to ensuring that I produced the highest quality work possible with credible empirics and quantitative rigor. Jonathan has inspired me to pursue big and fundamental economic questions grounded in theory. Antoinette has challenged me to sharpen my arguments through her piercing insights and deep institutional knowledge. Her guidance has been a unique source of inspiration to me as a rookie female economist in finance. Josh has given me the courage to pursue a Ph.D. even before arriving at MIT. Building on my apprenticeship as a pre-doc at Blueprint Labs, Josh has taught me how to apply and improve on this training during my Ph.D. Finally, it is thanks to Frank that I dotted all the i's and crossed the t's during the job market season. Frank has overseen every step of my research process and continuously reminded me the most important lesson: first, my well-being (mental and physical); then, research.

Beyond my committee, the support of the MIT finance and economics faculty has been instrumental to my accomplishments. I am especially grateful to Taha Choukhmane for his enthusiasm and feedback on my job market paper, David Thesmar for creating a fun and collaborative research environment for Ph.D. students by organizing a bi-weekly reading group, and Adrien Verdelhan for his dedication to improving the Finance Ph.D. program. I also thank Daron Acemoglu, David Autor, Amy Finkelstein, Daniel Greenwald, Debbie Lucas, Christopher Palmer, Larry Schmidt, Emil Verner, and Haoxiang Zhu for their comments and suggestions. Outside of MIT, I want to thank Natalie Cox and Constantine Yannelis for their advice and encouragement.

The JPMorgan Chase Institute has played a unique role in my Ph.D. journey. I first joined the Institute as a graduate research assistant to Antoinette Schoar and Jonathan Parker at the end of my first year of Ph.D. Over the years, the Institute has become my second intellectual home. I learned tremendously working alongside a team of brilliant and dedicated economists who care deeply about public policy. The year I spent in Washington D.C. as an on-site academic fellow was one of the best highlights of my Ph.D. I would like to thank Diana Farrell, Fiona Grieg, and Chris Wheat for their mentorship and guidance at the Institute.

I am very fortunate to have had such wonderful peers at MIT. Hector Blanco, Mary Gong, Jeremy Majerovitz, and I spent countless hours debating economics papers, cooking, playing board games, and working out together. Jian Sun has the sharpest mind and is one of the kindest people I know with whom I have had the pleasure to share an office. Sydnee Caldwell has been an amazing mentor and continues to be an inspiration. I am also grateful to Natalie Berfeld, Aileen Devlin, Mayara Felix, Peter Hansen, Mazi Kazemi, Madhav Kumar, Seung Kwak, Anton Petukhov, Pari Sastry, and Yupeng Wang. Outside of MIT, I thank my friends from Smith College (including Jenn's adorable puppy, Nugget), Sonia Cheung, Jung Choi, Miree Eun, and Helen Ho. They have been a source of tremendous encouragement and support through the years.

My family has been the only constant in my life through all the ups and downs. My parents, Yeounhee and Seonheung, have been my best cheerleaders, even at my lowest points. My sister Amy was only a phone call away whenever I needed her affirmation. Finally, my partner Antonio stood behind all my endeavors. His encouragement, humor, and unyielding confidence in me gave me the strength to persevere and the peace of mind to celebrate even the smallest wins. I could not have gotten this far without their unconditional love and support. I dedicate this dissertation to them.

Olivia S. Kim April 26, 2022 Cambridge, MA

Contents

List of Figures

List of Tables

Chapter 1

Credit and the Family: The Economic Consequences of Closing the Credit Gap of U.S. Couples

1.1 Introduction

Promoting fair and equal access to consumer credit has long been a policy goal in the United States.[1] But is credit shared equally in the household? And do disparities in access to credit between spouses lead to disparities in consumption? There are reasons to believe that disparities in banking services and credit persist within the household. Survey evidence shows that perceived financial inequity in the household is among the top predictors of divorce, and roughly half of marriages in the U.S. actually end in divorce.[2] Moreover, for married couples with a single household income–roughly half of U.S. couples (U.S. Bureau of Labor Statistics, 2020*a*) – breadwinners likely have higher borrowing capacity than their spouses, because income determines at least part of one's ability to borrow. Even for dual-income households, gender norms make Americans see

[1]Examples of financial policies aiming at equalizing access to credit include the Fair Housing Act of 1968, Equal Credit Opportunity Act of 1974, and Community Reinvestment Act of 1977.

[2]See Dew, Britt and Huston (2012) for survey evidence. The divorce to marriage rate was 44 percent in 2019 (Centers for Disease Control and Prevention, 2021).

men as the financial providers (Pew Research Center, 2017). However, the role of credit in the family has been understudied in finance and we know little about the extent and implications of credit disparities in the household. And while consumption inequality at least in part reflects differences in access to credit markets (Krueger and Perri, 2006; Blundell, Pistaferri and Preston, 2008), whether this association is causal remains an open question.

In this paper, I examine how policies that aim at reducing credit disparities affect the well-being of U.S. couples. Specifically, I analyze the 2013 reversal of the Truth-in-Lending-Act (TILA) as a source of exogenous variation in the amount of credit card limits extended to secondary earners. Before 2013, TILA required card issuers to evaluate card applicants' *independent* income in their lending decisions. The statute was reversed in 2013 to allow card issuers to consider *household* income, facilitating access to credit for secondary earners and stay-at-home spouses. Using detailed data on spouse-level financial accounts and a matched difference-in-differences design, I show that the reversal had the intended effect of increasing secondary earners' borrowing capacity. My central finding is that spouses shared consumption more equally, narrowing their pre-reversal consumption gap by 10 percent. Specifically, secondary earners' spending on "private" goods (for example, clothing) increased while that of primary earners decreased. Household spending on "public" goods (for example, home improvement) increased, suggesting that primary earners indirectly benefited from changes in household consumption patterns that reflected altruistic preferences of secondary earners. Household credit card debt increased moderately, with no material impact on delinquency rates and overdraft probabilities.

After establishing the causal link between credit and consumption disparities in the household, I use cross-sectional analysis and a calibrated model to clarify the economic mechanism. The limited-commitment (LC) channel posits that factors that improve the outside option of secondary earners (that is, the value of being divorced) should shift consumption allocation in their favor because a better outside option strengthens their bargaining power in the marriage (Chiappori and Mazzocco, 2017; Kocherlakota, 1996).[3] In fact, several institutional features make credit a plau-

[3]Hertzberg (2016) showed that strategic motives (rather than bargaining) under LC can lead individual family members to overconsume.

sible factor that can increase outside options.[4] Alternatively, imperfect information (Dubois and Ligon, 2011; Wang, 1995) or self-control stemming from differences in spouses' time preferences (Ashraf, Karlan and Yin, 2006; Bertaut, Haliassos and Reiter, 2009) can also lead to higher consumption shares for secondary earners. A heterogeneity analysis reveals patterns consistent with the predictions of the LC channel and at odds with other plausible channels. Motivated by this empirical result, I calibrate a dynamic model of household decision-making and show that the LC channel is quantitatively important, as it explains roughly one-third of the observed increase in secondary earners' consumption share.

My empirical strategy is a difference-in-differences design that compares secondary earners in equitable-distribution (ED) U.S. states, the treatment group, with those in community-property (CP) states, the control group. Secondary earners in CP states are a valid control group because card issuers were allowed to consider household income even before the reversal under marital division-of-property laws, which recognize household income as joint property regardless of who earns it. The identifying assumption is that, in the absence of the reversal, secondary-earner and household outcomes for the two groups would have evolved in parallel. To strengthen this parallel-trends assumption, I conduct nearest-neighbor propensity score matching to ensure that the treated and control groups have similar pretreatment characteristics that are thought to be associated with the dynamics of the outcome variables (Abadie, 2005). Because there is a never-treated group and a simultaneously absorbing treatment, my estimation does not suffer from the negative-weighting or underidentification problems that can arise in difference-in-differences setups with variation in treatment timing (Borusyak, Jaravel and Spiess, 2021; Callaway and Sant'Anna, 2020; Sun and Abraham, 2020).

I use a panel data set of monthly spending, income, and credit card borrowing covering roughly 66,200 opposite-sex couples, constructed using de-identified financial-account records from the JP-Morgan Chase Institute (JPMCI). This data set has the unique advantage of tracking the spending

[4]First, credit limits are portable in the sense that secondary earners or stay-at-home divorcées can keep the high credit limits that they obtained using household income while they were married because credit card issuers are prohibited from making lending decisions based on one's marital status. Second, since having a sole credit card account helps to build one's own credit history, secondary earners' access to credit can improve after divorce. Finally, while debt obligations are divided between spouses upon divorce according to marital division-of-property laws, credit *limits* are not considered marital property and do not get contested in divorce proceedings.

and credit use of individual spouses, which allows me to overcome the key measurement hurdle in the intrahousehold literature that spending is only observed at the household level. I proxy for each spouse's consumption by summing spending on their debit card, credit card, and checking-account transactions, such as cash withdrawals or electronic transfers;[5] I proxy their independent credit by summing the credit limits on their sole credit card accounts; and I proxy their total credit by summing the limits on all the credit card accounts the spouse has access to, either as a primary account holder or as an authorized user. I define spouse-specific consumption shares as each spouse's spending relative to total household spending, and I define the household's consumption gap as the difference in the two spouses' spending shares. Credit shares and gaps are constructed in the same manner.

An important measurement concern is that the broad spending measure I describe above may be a poor proxy for individual consumption. For instance, if spouses spend individually but consume a purchased good together, such as when there is a designated shopper, individual spending will not accurately reflect individual consumption. I address this concern by constructing a battery of alternative consumption measures. First, I construct a narrow consumption measure that only captures spending on gender-assignable goods, such as women's clothing or men's footwear. This measure provides a more precise proxy for consumption, under the assumption that gender-specific goods can be consumed by only one member of the household regardless of who purchased them. I also consider an alternative broad measure that includes credit card payments to other financial institutions and an alternative narrow measure that assumes that all cash withdrawals were used to purchase gender-specific goods. These measures address the potential concern that changes in payment behavior–such as increasing spending on credit cards that I observe while reducing cash spending or spending on cards that I don't observe–can lead to an upward bias of consumption gaps.

I find four main results. First, the TILA reversal had the intended effect of increasing access to credit for secondary earners. The estimated increase in secondary earners' sole credit card

[5]For joint checking accounts that are shared between spouses, I measure who spent what on these accounts by identifying which debit card is assigned to whom. For any transactions for which the identity of the spender cannot be clearly assigned, I assume that the spending was split equally by the spouses.

limits is 40 percent of their average pre-reversal monthly consumption, or roughly \$1,025. Put in context, the effect corresponds to 15 percent of the typical credit limit for card holders in my sample. This can be considered a first-stage effect since the reversal would not trigger a change in the consumption allocation between spouses without also affecting borrowing capacity.[6] I find that the treatment effect is driven by changes in secondary earners' income-reporting behavior rather than differential extensive-margin effects. Specifically, the reversal did not differentially affect secondary earners' propensity to open credit card accounts in the treated group relative to the control group; rather, conditional on opening an account, it differentially increased the credit limit for secondary earners in the treated group relative to the control group.[7] My estimates are not confounded by selection bias, as borrowing costs (that is, the annual percentage rate) and joint-account opening were invariant to the reversal.

Second, the central finding of the paper is that the TILA reversal reduced the consumption gap between spouses by shifting consumption toward secondary earners. I find that the reversal increased secondary earners' consumption by 14 percent relative to their pre-reversal monthly mean, or \$343. The reversal not only increased the secondary earners' *level* of consumption, but also increased their *share* of consumption in the household by 5 percent relative to their pre-reversal mean. The increase in household consumption is smaller (\$170) than the increase observed for secondary earners (\$343), suggesting that the consumption reallocation between spouses operated through primary earners cutting back their consumption. As a result, the consumption gap in the household narrowed by 10 percent relative to the pre-reversal mean. These results are robust to using a battery of alternative specifications, samples, and measures, such as the measure of gender-assignable consumption.

Third, despite increasing household credit card borrowing, the reversal did not worsen the financial standing of the household. The reversal increased household interest-accruing credit card revolving debt by 0.9 percent of pre-reversal average monthly household consumption, or \$51. Despite the increase, over the two-year post-reversal period, a variety of financial-solvency outcomes

[6]Placebo analysis confirms that there is no consumption effect for households in which secondary earners' borrowing capacity did not change.

[7]This result is consistent with the credit card industry's practice of using income to determine credit limits rather than to decide whether to issue credit.

were not materially impacted, including delinquency rates, overdraft probabilities, and the propensity to take out high-interest loans, such as payday or subprime loans. In addition, households with multiple credit card accounts become more likely to pay down more expensive debt first while borrowing more on lower-interest cards.[8] Overall, the size of the increase in household borrowing was smaller than the increase in secondary earners' consumption. That, combined with the fact that the indicators of financial solvency did not deteriorate, is consistent with spouses coordinating their consumption decisions to satisfy the family budget constraint. Accordingly, the results indicate that the narrowing of a consumption gap in the household did not come at the cost of worsened financial standing.

Finally, I find support for the LC channel as an explanation for these findings but no support for other potential channels. Under the canonical collective-household model with LC, since spouses cannot precommit to future allocations of resources, credit can empower secondary earners to act in their best interest and voice their opinions in the marriage, to the extent that higher borrowing capacity improves their outside options. This channel is relevant even for couples who are not on the verge of divorce insofar as there is some risk of divorce. That said, in the cross-section, the channel predicts that the effect of the reversal on secondary earners' consumption share should be larger (smaller) for couples with weaker (stronger) marital commitment because they will be more (less) sensitive to changes in the outside option. Consistent with this prediction, I find that the estimated effect is 40 percent larger for couples that are most likely to divorce and 50 percent smaller for couples that are least likely to divorce.[9] I find no clear-cut support for other channels, such as imperfect information, financial constraint, self-control, or limited attention.

Motivated by the reduced-form estimates showing support for the LC channel, I analyze the TILA reversal through the lens of the household-decision-making model under limited commitment and borrowing constraints. I find that the LC channel is quantitatively important. In this model, primary and secondary earners jointly decide how much to save and consume, whether to

[8]This measure of optimal debt prioritization is similar in spirit to the measure of cost-minimizing credit card repayment used in Ponce, Seira and Zamarripa (2017) and Gathergood, Mahoney, Stewart and Weber (2019).

[9]I define couples as most likely to divorce as those with positive pre-reversal spending on counseling, including couple counseling, or on dating services. Conversely, I define couples as least likely to divorce as those with positive pre-reversal spending on their children.

work, and whether to divorce by maximizing the weighted sum of their utilities, where the weights are their bargaining powers. The key feature of the model is that a spouse's bargaining power can change over time: whenever a spouse's outside option increases to the point at which the value of being divorced exceeds the value of staying married, the bargaining power adjusts just enough to make the spouse who prefers to divorce indifferent between divorcing and staying married. Higher bargaining power, in turn, leads to a higher consumption share in the household. Using this standard setup (Mazzocco, Ruiz and Yamaguchi, 2014; Voena, 2015), I test the quantitative importance of the LC channel by incorporating key feature of the reversal–namely, the secondary earners' expanded borrowing limits, which they can keep even after divorce–and track how secondary earners' share of consumption evolves in equilibrium. For realistic parameter values, I show that the model-generated consumption path can account for at least one-third of the observed increase in secondary earners' consumption share.

The main contribution of this paper is to apply the family-economics perspective to household finance. This perspective delivers two novel insights. First, there are substantial disparities in credit and consumption in the household. Second, a credit market policy that closed the credit gap in the household reduced consumption inequality. These results are important because promoting equal access to credit is an active and pressing policy agenda, but the consequences of reducing credit disparities are not well understood. Moreover, theories in the intrahousehold bargaining literature predict that empowering the spouse with lower initial bargaining power by increasing their financial means should shift consumption allocation in their favor, but directly measuring how credit and consumption are shared in the household is challenging because information on individual spouses' consumption is not typically available. I tackle these challenges by leveraging the unique institutional setting of the TILA reversal and rich administrative data that track individual spending behaviors of each spouse. By providing credible evidence that credit market policies can have an uneven impact on individual family members, this paper puts forward a new research agenda in the household finance literature that focuses on how finance shapes and is shaped by family dynamics. Section 1.2 discusses how this paper relates to the existing literature.

1.2 Related Literature

This paper relates to the growing literature on how financial regulations or practices in the financial sector affect inequality. Existing studies document how mortgage market policies can lead to racial disparities in lending outcomes, with risk-equivalent Latinx/African-American borrowers paying 4.9 basis points higher interest rates relative to White borrowers (Bartlett, Morse, Stanton and Wallace, 2021), and quantity-focused lending policy leading to higher incidence of fraud to minority customers (Begley and Purnanandam, 2021) and greater racial segregation (Malmendier and Kulkarni, 2021).[10] I contribute to this literature by documenting the uneven impact of a financial policy on different members in the same household. An implication of my study is that policies aimed at equalizing the credit gap between spouses can reduce the within-household inequality.[11]

This paper also contributes to a vast literature on intra-household bargaining. Existing studies show that giving spouses with low initial bargaining power (typically women) more control over income (Schultz, 1990; Thomas, 1990; Browning, Bourguignon, Chiappori and Lechene, 1994; Duflo and Udry, 2004; Blundell, Chiappori and Meghir, 2007; Bobonis, 2009), cash transfers (Lundberg, Pollak and Wales, 1997; Attanasio and Lechene, 2014), savings accounts (Ashraf, Karlan and Yin, 2010), or better outside options in marriage markets in terms of gender ratio (Angrist, 2002; Chiappori, Fortin and Lacroix, 2002) reduces their labor force participation and changes household consumption patterns in a way that better reflect preferences of the wives, with greater spending on education, housing, and nutrition for children. Other studies find that changes in divorce laws can change savings behavior depending on post-divorce asset allocation (Voena, 2015; Lafortune and Low, 2020). A related of set of studies that evaluate the effect of microcredit programs targeting women's financial independence find that improving women's financial control

[10] Related studies show that bias, technological innovation, or information disparity in consumer lending markets can lead to inequality. See, for example, Butler, Mayer and Weston (2021); Lanning (2021); Argyle, Indarte, Iverson and Palmer (2021) for the role of bias of loan officers or bankruptcy stakeholders; Morse and Pence (2020); Fuster, Goldsmith-Pinkham, Ramadorai and Walther (2021) for technological innovation in underwriting; Blattner and Nelson (2021) for information disparity in assessing consumer's default risk; and Cox (2019) and Catherine and Yannelis (2021) for distributional consequences in the student loan market. Goldsmith-Pinkham and Shue (2020) and Kermani and Wong (2021) document gender and racial disparities in housing returns.

[11] A distinct but related set of studies quantify the degree to which consumption inequality tracks income inequality. See, for example, Krueger and Perri (2006); Blundell, Pistaferri and Preston (2008); Aguiar and Bils (2015); Blundell, Pistaferri and Saporta-Eksten (2016).

reduces household consumption of temptation goods, but has no discernible effect on women's empowerment (Banerjee, Duflo, Glennerster and Kinnan, 2015).

I contribute to this literature in two ways: measurement and novel policy setting. First, due to the paucity of spouse-level consumption data, most prior work either uses survey-based household expenditure data to proxy for household consumption or estimates intra-household consumption sharing rule using a structural approach by imposing restrictions on preferences to identify individual demand from household-level expenditure data. A key limitation of using household expenditure data is that it does not capture how spouses allocate consumption in the household. For this reason, the latter approach has become central to advancing the literature, but the lack of spouse-level data has limited researchers' ability to assess empirical credibility of this approach.[12] I overcome this challenge using spouse-level financial accounts data and complement existing research by providing reduced-form evidence that is consistent with the findings derived using a structural approach.[13]

Second, this paper examines how access to credit affects consumption allocation in the household in a novel policy setting. Existing studies that consider the relevance of credit mainly focus on the impact of microcredit extended to women in developing countries where gender and cultural norms make it difficult for women to obtain credit (Fletschner, 2009). This study examines the role of credit in a setting where credit use is widespread and cultural norms do not dictate access. While microcredit is a targeted policy tool with relatively low take up rates of 33 percent (Banerjee, Duflo, Glennerster and Kinnan, 2015), the U.S. credit card market is highly sophisticated and affects more than 80 percent of all American adults (CFPB, 2019). Since cultural and institutional differences play a central role in shaping decision-making within the family (Bau and Fernández, 2021), understanding the effect of credit access on decision-making of U.S. couples is an open, relatively under-explored question. I fill this gap by analyzing a novel policy event that changed the legal environment in credit card markets. To my knowledge, this paper is the first to examine

[12] Chiappori and Meghir (2015) highlight this concern: "The allocation of resources within the household cannot (in general) be directly observed; It has to be recovered from the household's (aggregate) behavior... It is evident from this discussion that better data would be important; and nothing is more important than detailed consumption and time use data. A renewed emphasis on such data is called for, given the importance of the issues at hand."

[13] See Chiappori, Fortin and Lacroix (2002) and Lise and Seitz (2011) for examples of the structural approach.

the effect of the TILA reversal.

This paper also relates to the literature that attributes various puzzles in household finance to coordination frictions between spouses. These studies find that the drop in consumption at retirement (Lundberg, Startz and Stillman, 2003); households simultaneously carrying high-interest debt and low-interest liquid assets (i.e., the co-holding puzzle) (Bertaut, Haliassos and Reiter, 2009; Vihriälä, 2020); over-consumption (Hertzberg, 2016; Olafsson and Pagel, 2017); inefficient retirement savings (Choukhmane, Goodman and O'Dea, 2021); couples' failure to aggregate information even when it's in their interest to do so (Ashraf, 2009; Conlon, Mani, Rao, Ridley and Schilbach, 2021); or low stock market participation (Ke, 2021) are driven by differences in time-preferences, bargaining power, private incentives, limited financial pooling, gender norms, or coordination frictions between spouses. I complement these studies by highlighting the importance of evaluating household behavior through the lens of individual family members. I show that household averages mask heterogeneity in spouse-level consumption response due to reallocation. If within-household dynamics is not accounted for, the reversal would be interpreted as having little effect on household behavior, despite the fact that it led to substantial consumption reallocation in the household. Thus, recognizing households as families can enrich our understanding of household behavior.

1.3 Institutional Background and Research Design

The 2013 reversal of the Truth-in-Lending Act Section 150, or the ability-to-pay provision, exogenously increased secondary earners' access to credit in the credit card market, providing an ideal setting to study how mitigating the credit gap affects the consumption gap in the household. Section 1.3.1 presents the institutional background on TILA and the 2013 reversal. Section 1.3.2 describes the empirical strategy, and Section 1.3.3 discusses the validity of this strategy.

1.3.1 The Truth-in-Lending Act

The 1968 Truth-in-Lending Act (TILA) is a federal statute that requires lenders to disclose terms and cost – such as the annual percentage rate (APR) – to consumers and bans lenders from using deceptive advertising practices (CFPB, 2021). The TILA governs a wide range of consumer credit products including credit cards, mortgages, auto, and installment loans.

This study examines the reversal of an amendment to TILA Section 150, which applies to the credit card market. In October 2011, roughly two years before the reversal, the Board introduced an amendment to Section 150, mandating credit card issuers to specifically consider the consumer's "independent" ability to pay when they issue credit.[14] Prior to the amendment, Section 150 did not offer any specific guidance and stated that:

> a card issuer may not open any credit card account for any consumer under an open end consumer credit plan, or increase any credit limit applicable to such account, unless the card issuer considers the ability of the consumer to make the required payments under the terms of such account (12 CFR §1026, 2012).

After the amendment, card issuers were required to either (i) consider the consumer's independent means of repaying through information collected on a credit card application; or (ii) obligate the consumer to have a cosigner who has such means and can assume joint liability for the account. The original intent of this amendment was to restrict card issuers from extending credit to consumers under the age of 21 to address a growing concern at the time that young adults were being offered credit cards on the basis that their parents had enough income, without the parents' consent. However, the amendment raised an unexpected concern that it may restrict secondary earners and stay-at-home spouses who have limited income of their own but access to their spouse's income from establishing access to credit.

Growing concerns about the 2011 amendment having discriminatory effects on secondary earners and stay-at-home spouses prompted a Congressional hearing to consider reversing the amend-

[14]This amendment was first introduced in February, 2010 when the Board adopted TILA §226.51, which implements the provisions of the Credit CARD Act of 2009. §226.51 has two parts, 226.51(a) and 226.51(b), which respectively implement TILA Section 150 and Section 127(c)8 that concern the consumer's independent ability to pay and the restriction on extending credit lines to consumers under the age of 21. In March, 2011, the Board issued a clarification that these subparts would in effect become one statute that applies to all consumers, regardless of age. This amendment was effective starting October, 2011.

ment. The nature of these concerns are reflected in the opening statement of the June 2012 Congressional hearing by Senator Shelley Capito:

> This rule could be especially punitive for women who are in a failing marriage or an abusive relationship. As I think about what some of the fundamental steps somebody who is maybe in an unhappy marriage or an abusive relationship would take, one of the fundamental, I am sure, pieces of advice is to try to establish credit, try to establish a financial footprint. Similarly, stay-at-home spouses whose husband or wife dies unexpectedly or divorces them could face similar challenges if they have not maintained a credit history.... The ability to pay rule threatens to further complicate the situation by potentially limiting their access to credit. (House Hearing: 112th Congress, 2012)

The 2011 amendment was reversed in 2013, allowing card issuers to "consider income and assets to which consumers have a reasonable expectation of access" for consumers over the age of 21. The CFPB announced this change in May, 2013, and compliance with this rule was required by November, 2013. This paper examines the effect of the 2013 reversal by tracking how credit access and consumption allocation of U.S. couples changed around November, 2013. Figure A-1 illustrates the timeline visually. The portion highlighted in blue – 12 months before and 24 months after the reversal – covers the time period analyzed in this study.

1.3.2 Research design

My research design exploits the fact that the TILA reversal was superseded by state marital property laws in some states but not in others. In community property (CP) states, card issuers were allowed to consider secondary earners' "household income" because any income earned during marriage is considered to be jointly owned, regardless of who actually earned it. In equitable distribution (ED) states, however, card issuers were required to consider secondary earners' "independent income" prior to the reversal because income earned during marriage is considered to be separately owned, with the potential of being divided equitably upon divorce.[15] Since households living in CP states were not subject to either the 2011 amendment or the 2013 reversal while those living in ED states were, households in CP states act as my "control" group and those in ED states form the "treated" group. Figure 1-2 shows the map of where CP and ED states are located in my

[15] See 12 CFR §226 (2011) for details on the Board's suggested treatment of applicants residing in community property states when applying the 2011 independent ability-to-pay amendment.

sample. CP states are shaded in green and ED states in purple. States that are not well represented in my data are shaded in light gray.

I use the following difference-in-differences (DiD) regression specification:

$$Y_{h,t}^{i} = \alpha_h + \gamma_t + \beta \mathbf{1}[Treat \times Post]_{h,t} + \epsilon_{h,t} \tag{1.1}$$

where $Y_{h,t}^{i}$ is an outcome for secondary earner i in household h at month t. α_h are household fixed-effects, γ_t are month-year time fixed-effects, and $\mathbf{1}[Treat \times Post]_{h,t}$ is an interaction term between an indicator for a household living in ED states and an indicator for time t being November 2013 or after. The coefficient of interest, β, captures the differential change in the outcome for the treated group relative to the control group following the reversal.

In addition to Equation 1.1, I also use the following dynamic DiD specification to visualize the treatment effect dynamics:

$$Y_{h,t}^{i} = \alpha_h + \gamma_t + \sum_{s \neq t^{post}-1} \beta_s [Treat \times 1_{s=t}]_{h,t} + \epsilon_{h,t} \tag{1.2}$$

where I omit the month prior to the reversal, $\beta_{t^{post}-1}$, so the other β_s can be interpreted relative to this pre-reversal period. For all regressions, I cluster standard errors at the state-level.

The identifying assumption is parallel trends: the average outcomes for treated and controls would have followed parallel paths over time in the absence of treatment. While the parallel trends assumption does not require the outcomes to look similar in levels across treated and control units, this assumption may be implausible if pre-treatment characteristics that are thought to be associated with the *dynamics* of the outcome variable are unbalanced between the treated and the control group (Abadie, 2005). For example, if households in the two groups showed large differences in initial access to credit or debt-to-income (DTI) levels, these level differences can generate differential trends in future access to credit even in the absence of the reversal because they affect card issuers' underwriting decisions. Households with low initial DTI (credit) will be more likely to obtain credit in the future relative to those with high initial DTI (credit), regardless of the policy change. Thus, to the extent that selection for treatment is influenced by households' past outcomes, differences in pre-treatment characteristics can generate Ashenfelter's dip that leads to an upward biased DiD estimate (Ashenfelter and Card, 1985).

I apply the nearest neighbor propensity score matching method (Rosenbaum and Rubin, 1983) to strengthen the "parallel trends" assumption. Specifically, I match households based on their conditional probability of being treated given the covariates (the propensity score). Because propensity score has a balancing property, the treated and the control group households have the same distribution of covariates, conditional on the propensity score. In practice, this method requires first estimating the propensity score $p(X) = P(Treat = 1|X)$ using logit regression, then matching on the estimated propensity score. To estimate the propensity score, I choose pre-treatment covariates (X) that may influence the card issuer's underwriting criteria, discussed in Section 1.4.4. This method has been used widely in observational studies as a testing ground for nonexperimental methods (e.g., Dehejia and Wahba, 1999; Hirano and Imbens, 2001; Angrist, Autor, Hudson and Pallais, 2015).

1.3.3 The Reversal of the Truth-in-Lending Act in Practice

Two conditions must hold for the TILA reversal to provide a credible identification setting. First, the card issuer's treatment of treated and control states must be different prior to the reversal for using CP states as the control group to be valid. Second, the card issuer's compliance with the reversal must trigger a change in the income reporting behavior of treated secondary earners. Since card issuers use reported income to determine the amount of credit limit to extend to a card applicant, treated secondary earners should report higher income after the reversal to see a larger increase in credit limit relative to the control group.

I confirmed with JPMorgan Chase & Co. (JPMC) that the first condition holds. While the Board permitted card issuers to collect "household income" from applicants in CP states prior to the reversal, it was ultimately up to card issuers to decide whether to consider "independent income" universally or in ED states only. JPMC applied the independent ability-to-pay criteria to ED states only, thus validating the first condition of my identification strategy.

Secondary earners' reported income on their credit card applications validates the second condition. Figure 3-2b shows the average difference in reported monthly income between treated and control secondary earners before and after the reversal. Before the reversal (left bars of Figure 3-2b), secondary earners in treated states reported \$380 lower income on average relative to those in control states. This difference entirely disappears with the TILA reversal (right bars of Figure 3-2b), suggesting full compliance with the policy change. The difference in the reported income is even larger for single income households, corresponding to roughly \$500, or 14 percent of median monthly household income. Note that the average pre-reversal household income of the treated and the control group are by design similar due to propensity score matching. Therefore, if the 2011 Ability-to-Pay amendment or the 2013 reversal were not enforced, secondary earners or stay-at-home spouses' reported income would be similar in the two groups both before and after the reversal.

Overall, these two conditions suggest that the TILA reversal is likely to generate a differential increase in treated secondary earners' access to credit. Figure A-2 shows that there is no difference in primary earners' reported income between the two groups either before or after the reversal. These results confirms that the reversal provides an ideal setting for examining how mitigating credit disparities in the household through an improvement in only one of the two spouses' access to credit affects the within-household consumption gap.

1.4 Data and Descriptive Evidence

This study uses a panel dataset of monthly spending, income, and credit card borrowing of 66,200 opposite-sex couples from October 2012 to December 2015, covering a year before and two years after the TILA reversal.[16] This dataset is derived from de-identified transaction-level records of checking, debit card, and credit card accounts obtained from a U.S. financial services company that provides retail banking services. The key novelty of this data is the ability to track spending and credit use of individual spouses in the household. Section 1.4.1 describes the sample construction steps. Section 1.4.2 discusses how I construct the main outcomes.

[16] I do not analyze the 2011 amendment in this paper because the data starts from October 2012.

1.4.1 Analysis Sample

The construct my sample in three steps– (i) identifying individual spouses in each couple; obtaining information on each spouse's (ii) checking; and (iii) credit card accounts.

I identify individual spouses in each couple using a record of account linkages that links family members to a unique household identifier.[17] Since this dataset only covers account holders at JPMC, I focus on a sample of couples where both spouses have a financial account at JPMC. To allow my sample of households to have a diverse set of financial account structure, I do not require individual spouses to have separate financial accounts. So the sample includes, for example, couples that only have a shared joint account as well as those with a mix of individual and shared accounts. I restrict the sample to households that only have two family members to focus on the behavior of couples. To mitigate potential confounding effects from retirement, I restrict the sample to spouses in their prime working age (25 to 65 years old) at the timing of the reversal. Because I do not directly observe the marital status of individual members linked to the same household unit, I further restrict the sample to households with opposite-sex members with the age gap of less than 16 years to focus on the sample are most likely to represent married couples. Figure A-5 shows that 92 percent of households hold a joint checking account shared with the family member. Given that joint checking account in the household is typically shared between spouses – rather than, for example, siblings – my sample is highly likely to capture individuals who are married.[18]

Next, I obtain each spouse's checking account information to ensure that individual member's spending can be tracked both before and after the reversal. I require *both* spouses to have at least one active checking account (i.e., at least 5 transactions every month) at JPMC either as a primary or secondary account holder. This allows me to capture couples that have individual checking accounts, as well as those with only one joint checking account. For couples with joint checking accounts, I require spouses to have their own debit card associated with these shared accounts to

[17]Individuals must share personally identifiable information to be linked to the same household unit. In addition, only family members over the age of 18 can be tracked in this dataset.

[18]The distribution of the financial account structure – which serves as a proxy for marital status – looks similar between the treated and the control group. Thus, potential mis-measurement of family structure is unlikely to be correlated with the treatment assignment and lead to differential selection concerns across states.

be able to track each spouse's spending on these joint accounts. I further restrict the sample to couples that make annual labor income – measured as the sum of all payroll direct deposits – of at least $17,000 in 2013,[19] to focus on couples that primarily use JPMC checking accounts to manage their finances.

In the final step, I obtain each spouse's credit card account information and restrict the sample to marginal households that are likely to be affected by the reversal. I require (i) at least one spouse in the household to have a credit card account; and (ii) secondary earners to not have a sole credit card account at the beginning of my sample period (October 2012). The first restriction allows me to focus on couples that borrow from the credit card market. The second restriction allows me to focus on couples where secondary earners have the highest propensity to open credit cards around the policy change. This second restriction is motivated by the fact that the reversal is unlikely to affect existing card holders because they rarely update their income.[20] On the other hand, because new card openers are required to report their income on their credit card applications, focusing on the sample with the highest propensity to open new credit cards can strengthen internal validity when evaluating the impact of the reversal.

Focusing on marginal households raises a concern that treated secondary earners may be more likely than control secondary earners to open credit cards after the reversal, creating Ashenfelter's dip. In Section 1.5.1, I show that this is not a concern in my setting because there is no differential extensive margin effects. The absence of differential card opening is consistent with JPMC's – more broadly, the credit card industry's – practice of using income primarily for deciding how much credit limits to extend to an applicant rather than whether to issue credit. Overall, this implies that my sample is not subject to potential selection concerns.

1.4.2 Variable Construction

Consumption. Monthly spouse-specific consumption is proxied by spending on each spouse's financial accounts. Specifically, spouse i's consumption is defined as the sum of all spending categories on i's sole and joint credit card, debit card, and checking accounts, including cash

[19]$17,000 is the U.S. Department of Health and Services' 2013 poverty threshold for two-member household.
[20]JPMC asks existing card holders to update their income once a year, but the response rate is low.

withdrawals and electronic transfers. To track who spent what on the couples' joint checking account, I attribute spending to the respective debit card holder of the shared accounts. For any joint account transactions for which the identity of the spender cannot be identified, I assume that spouses equally shared these expenses.[21] Note that this is a conservative assumption as it pushes consumption shares to be equal.

The spending categories are aggregated from transaction records of each financial account using a combination of the Merchant Category Code, transaction counter party, and JPMCI's internal categorization variables:[22]

$$\begin{aligned}\text{Consumption}^i = \text{Dept store} + \text{Entertainment} + \text{Flights} + \text{Hotels/Rental} +\\ \text{Insurance} + \text{Medical} + \text{Transport} + \text{Food Away} + \text{Dur Retail} + \text{Nondur Retail} +\\ \text{Cash} + \text{Prof/Psnl Svcs} + \text{Auto Repair/Parts} + \text{Fuel} + \text{Utilities} + \text{Grocery} +\\ \text{Home Improvement} + \text{House Keeping/Home Repairs} + \text{Children's Exp}\end{aligned}$$

Each spouse's consumption can be further broken out into "private" and "public" consumption. Private consumption (shown in blue) refers to spending on exclusive goods that are consumed privately and only benefit the spouse who spends the money. Public consumption is shown in red and refers to spending on common goods that are consumed jointly by the household. The categorization of "private" or "public" consumption follows existing studies (Chiappori, Fortin and Lacroix, 2002; Mazzocco, 2007). These detailed consumption types can help assess the shift in the composition of consumption following the reversal. Finally, total household consumption is measured as the sum of each spouse's consumption.

I construct two within-household consumption measures: (i) the consumption shares of each spouse, which measure how consumption is allocated in the household; and (ii) the consumption gap in the household, which measures the magnitude of consumption inequality in the household. Each spouse's consumption share is measured as the spouse's consumption relative to total house-

[21] For example, if a joint checking account shows an electronic bill payment of $100, I attribute $50 to one spouse and $50 to the other spouse.

[22] Table A.4 reports examples of the expenditures included in these 19 spending categories.

hold consumption. The consumption gap in the household is measured as the difference between the consumption shares of each spouse. For example, if the secondary earner consumes 45% and the primary earner 55% of total household consumption (i.e., their respective consumption shares are 45 and 55%), the within-household consumption gap is $55 - 45 = 10\%$.

An important concern for my measurement is that spending may be a poor proxy for consumption to the extent that spouses spend individually but consume the purchased goods together.[23] I address this issue by constructing a narrower measure of consumption defined only based on gender-assignable goods, such as male or female clothing. Unlike the "broad" consumption measure, which hinges on the assumption that the person who spends the money is the one who consumes it, the advantage of the "narrow" measure is that the purchased good can exclusively be consumed by only one member in the household because of gender, regardless of who purchased it,[24] thus providing a more precise proxy for consumption.

Another advantage of using the "narrow" measure is that it can be used to validate the assumption behind the "broad" measure. Figure A-4 shows that the "spender as the consumer" is a reasonable assumption. For example, Figure A-4a plots the wife's average monthly consumption share constructed using the broad measure against their consumption share constructed using the narrow measure. The relationship between the two measures are positive and monotonically increasing. This suggests that the consumption share constructed using the broad measure is a reasonable proxy for the actual consumption share in the household. If the person who spends the money is not the one who consumes it, the relationship between the two measures would be flat. The husband's measures show a similar pattern (see Figure A-4b).

In addition to the "narrow" measure, I also construct a battery of alternative consumption measures to address potential concerns that (i) spending activity at other financial institutions are not captured; and (ii) there may be potential substitution with cash spending. Section 1.5.4 discusses these concerns in detail.

[23] For instance, couples can coordinate their shopping– i.e., one spouse always buys groceries and the other puts gas in the car – but they consume these goods together. Such "designated shopping" behavior can potentially overestimate the degree of consumption inequality.

[24] This approach has been used widely in the intra-household literature to identify spouse-specific demands from aggregate household expenditures.

Credit. I construct two credit measures – independent credit and total credit. Spouse i's *independent* credit access is proxied by the sum of credit limits on each spouse's sole credit card account; and i's *total* credit access is the sum of credit limits on any credit card account he or she has access to either as a primary account holder or as an authorized user. Independent credit captures how much each spouse can borrow independently and can continue to borrow even after divorce, and total credit captures how much each spouse can borrow either independently or jointly with the other spouse. Household credit access is measured as the sum of all credit limits available to spouses and measures couples' total borrowing capacity if they fully pooled credit together. Credit limits on joint accounts are only counted once in the household-level aggregation.

Similar to consumption inequality, I construct two measures of credit access inequality that capture how credit is allocated between spouses (credit shares) and the magnitude of credit inequality in the household (credit gap). The credit shares of each spouse are measured as each spouse's total credit access relative to household-level credit access. The credit gap in the household is measured as the difference between the consumption shares of each spouse.

Income. Monthly spouse-specific income is measured as the sum of labor income (payroll direct deposits), government transfers, and other income deposited to spouses' sole and joint checking accounts *for which they are the primary account holder*. Since it is difficult to identify who earned what on the joint accounts, I treat the primary account holder as the earner of any income deposited to joint accounts. Government transfers include unemployment insurance, veteran's benefits, and tax refunds; and other income includes business or gig income. Household income is measured as the sum of each spouse's income.

I use this income measure to determine who is the primary/secondary earner in the household and whether a household is single- or double-income. A spouse is primary earner if he or she earned higher average monthly labor income relative to the other spouse in the pre-reversal period. A household is double-income if (i) it receives more than 4 payroll direct deposits in a month; or (ii) receives more than 2 payroll deposits in a month and the difference in the amount deposited in each paycheck is larger than one standard deviation of monthly labor income that households receive on average. Given that wage earners typically receive income on a bi-weekly basis (U.S.

Bureau of Labor Statistics, 2020*b*), the frequency of payroll deposits and the payroll difference between spouses helps to identify double income households.

A potential concern with my income measure is mis-classifying which spouse is the primary earner. Since I cannot identify which income streams belong whom if both spouses deposit income into their joint checking account, this mis-classification can arise if the spouse making higher income is not the primary account holder. Given that a majority (84%) of opposite-sex married couples in the U.S. have husbands earning the same or more than their wife (Current Population Survey, 2020), the mis-classification concern is most likely to arise for couples with only one joint checking account where the wife is the primary account holder. In my sample, 16.1% of couples have this type of account structure. However, the mis-classification is unlikely to be a concern in my setting for two reasons. First, the distribution of household account structure types is similar between treated and control households (see Figure A-7 and A-8). This suggests that the mis-classification is unlikely to bias DiD estimates because it is uncorrelated with the treatment assignment. Second, the mis-classification attenuates (not overestimates) the degree of within-household consumption gaps in the descriptive analysis. Existing literature shows that wives' consumption share tends to be smaller than husbands' (Lise and Seitz, 2011), suggesting that mis-classifying wives (husbands) as primary earners likely under- (over-) estimates the consumption gap in the household. In my sample, 39 percent of couples have breadwinning wives, whereas this share is only 16 percent in a nationally representative sample, suggesting that the misclassification is likely to attenuate the consumption gap. Consistent with this, Section 1.5.4 shows that the effect on the consumption gap is larger for the subset of couples that do not have any joint accounts and are not subject to the mis-classification concern.

1.4.3 Pre-Treatment Characteristics and Sample Representativeness

The treated and the control group households have similar pre-treatment characteristics. Columns 1 through 3 of Table B.2 show that the treated group has higher baseline average income and liquidity, and is more likely to have credit cards before the matching procedure described in Section 1.3.2. Columns 4 through 6 show that matching on propensity score yields 66,196 households with similar pre-treatment characteristics. This is my main analysis sample.

My sample of households look similar to a representative sample of U.S. households. Table A.2 compares average characteristics of my sample to a representative sample of two-member households using the Consumer Expenditure Survey (CEX) and Bureau of Labor Statistics (BLS). Compared to the benchmark mean reported in Column 1, Column 2 shows that my sample tends to be slightly younger and has higher consumption and income. These differences can be driven by differences in sample and/or measurement: the CEX sample includes retirees, whereas I focus on couples in their prime working age that presumably have higher consumption and income; and the CEX may under-report consumption and income (Cantor, Mathiowetz, Schneider and Edwards, 2013; Mian and Sufi, 2016). Despite the differences in levels, I find that the ratio of consumption to income or the ratio of public (or private) consumption to household consumption match the CEX closely. The share of double-income households in my sample also match the share reported by the BLS.

My sample of households exhibit substantial heterogeneity both within and between couples. Panel A of Table 2.1 shows monthly pre-reversal household characteristics and illustrates heterogeneity between couples. Couples on average consume and earn total income of $6,319 and $9,011, respectively, while the median household consumes and earns roughly 25 percent less than the average household. Couples on average have access to some credit (i.e., at least one spouse has a credit card account) 74 percent of all months, while a median couple always has access to credit before the reversal. Panel B illustrates heterogeneity in consumption, income, and credit within the household. On average, primary earners earn 8 times more and consume 1.5 times more than secondary earners. Secondary earners are substantially less likely to be able to borrow independently before the reversal relative to primary earners. Note that the income gap between spouses is likely to be overstated because I attribute all income streams to the primary account holder when both spouses deposit income into their joint checking accounts (see Section 1.4.2).

1.4.4 Descriptive Evidence

Before discussing the causal impact of the TILA reversal, I document three novel facts that motivate understanding the link between disparities in credit access and consumption in the household.[25] These results show that there is a strong correlation between credit and consumption gaps in the household. Below I discuss these results in detail.

First, there are large gaps in credit access within the household. Figure 1-4 plots average total and independent credit shares by earner type. Primary earners have access to 92% and secondary earners 35% of total credit limits available at the household-level, indicating a credit gap in the household of 57%. The independent credit gap in the household is even larger ($0.61\% = 0.80 - 0.19$), suggesting that secondary earners are much less likely than primary earners to be able to borrow independently from credit markets. The large within-household credit gap is partly driven by the fact that I focus on households where secondary earners did not have a sole credit card account at the beginning of my sample period (see Section 1.4.1 for details). However, Figure A-9 shows that the within-household credit gap is still large even in the broader sample of households without this credit card sample restriction.

Second, there are large gaps in consumption within the household. The average consumption gap in the household is 18%, with secondary earners consuming 41% and primary earners 59% of total household consumption. This implies that secondary earners consume 69 cents for every dollar consumed by primary earners. The consumption gap in the household cannot be explained by differences in the spouses' income. If income determines consumption shares of each spouse, individuals that make similar levels of income should consume similar shares of consumption in their respective household. However, Figure 1-3 shows that individuals in the same income bin consume more than their respective spouse only if they are primary earners, suggesting that the relative (rather than nominal) financial power in the household determines consumption share.[26]

Finally, secondary earners' average credit share is positively correlated with their consumption

[25] The descriptive analysis presented in this section uses all sample period, not just the pre-reversal period.

[26] Specifically, this figure plots the average consumption share of individuals in the same income bin by their earner status (primary vs. secondary) in their respective household. Individuals in every income bin has higher consumption share relative to the other spouse if they are primary earners but not if they are secondary earners.

share in the household. Figure 1-5a shows that the consumption share of secondary earners increases monotonically with their share of total accessible household credit, suggesting that having a higher relative borrowing capacity allows one to consume relatively more in the household. However, Figure 1-5b shows that this positive correlation disappears when the average consumption share is plotted against the amount of credit access secondary earners have as an authorized user. This suggests the positive correlation between credit and consumption allocation in the household is driven by one's ability to access credit independently.

1.5 Effect of the Reversal on Inequality in the Household

The descriptive evidence presented above may partly reflect the causal link between disparities in credit and consumption in the household. Motivated by these correlations, I examine the causal effect of the reversal on credit and consumption in the household. Throughout this section, I focus on 11,686 households where secondary earners opened sole credit card accounts at some point during my sample period (i.e., "the card holder sample"),[27] and show that my results hold up to using a broader sample of 66,196 households, including those where secondary earners never opened credit cards (i.e., "the all sample").

1.5.1 Effect on Secondary Earner Credit

Panel A of Table 1.3 presents results from estimating Equation 1.1 where the outcome is secondary earner's independent credit, or the credit limit on their sole credit card account. The outcome is scaled by secondary earners' pre-reversal average monthly consumption, so the estimated coefficient can be interpreted as a percent change relative to their pre-reversal consumption. Estimates restricting the sample to "the card holder sample" are reported in Columns 1 through 3. Estimates using "the all sample" are reported in Columns 4 through 6. Within each "card holder" versus "all sample", Columns 2 and 5 report estimates using only a sample of single income households and Columns 3 and 6 report estimates using only households where secondary earners are older than

[27]Conditioning on card openings does not bias estimates because card opening is invariant to the reversal.

their spouse. Panels B and C reports pre-reversal mean of the outcome and economic significance.

The estimates in Table 1.3 confirm that the TILA reversal expanded credit access for secondary earners and reduced the credit gap in the household. Column 1 shows that secondary earners' credit access increased by 40 percent relative to their pre-reversal average monthly consumption, or \$1,024 ($0.405 \times$ \$2,532). The unconditional (conditional) average credit limit on secondary earner's sole credit card account before the reversal was \$925 (\$6,826), implying that the estimated effect is as large as (15 percent of) the average pre-reversal credit limit. This increase reduced the credit gap in the household by 52 percent, from 49 to 23.2. Columns 2 and 3 show that the estimated effects are larger for stay-at-home spouses and older secondary earners, consistent with the reversal having a larger effect on secondary earners who have limited income of their own in the pre-reversal period have a longer credit history. Columns 3 through 6 show that the results are robust to using all sample of households. These results can be thought of as the "first-stage" since the reversal would not trigger a change in couples' consumption allocation without any impact on credit access.

Table A.6 shows the estimated coefficients on credit card opening or closing are close to 0, suggesting that the reversal did not differentially increase the likelihood of opening or closing a credit card account for treated secondary earners. The reversal could – in theory – have generated differential mean reversion (i.e., Ashenfelter's dip) if treated secondary earners waited until after the reversal to apply for credit cards in anticipation of the reversal, or if card issuers advertised credit card products more aggressively in the treated states to meet pent-up demand. However, secondary earner's card opening rates were invariant to the reversal. This finding is consistent with the fact that card issuers use income information for deciding *how much* credit limit to extend to an applicant rather than whether to extend a credit line.[28]

The reversal did not differentially affect the joint account opening behavior or credit card pricing. Table A.6 shows that the reversal did not affect the joint account opening behavior of primary or secondary earners, and Table A.7 shows that treated and control secondary earners paid similar APR on their sole credit card accounts. These results suggests that my estimates are not

[28] I confirmed this with various credit card professionals and regulators who are familiar with this policy change. Card issuers primarily use applicants' credit profile for deciding whether to extend a credit line.

subject to potential selection concerns arising from changes in guarantee behavior or the reversal inducing riskier borrowers to open credit cards. Table A.7 shows that the reversal increased secondary earners' sole credit card balance but reduced utilization rates, indicating that the reversal relaxed secondary earners' borrowing constraints as the increase in credit limits more than offset the increase in spending. The reversal increased secondary earners' credit card payments to other financial institutions. This alleviates the substitution concern that secondary earners may be increasing spending on cards that I observe while reducing spending on cards that I do not observe, thus leading to overestimation of consumption effects.

Figure 1-7a plots the event-study estimates on secondary earner's credit limit, β_s, obtained from estimating Equation 1.2. The vertical red line shows the month of the reversal (November 2013), and the light gray bands around the estimates show the 90 percent confidence bands. The figure illustrates that treated secondary earners' credit limit trended in parallel with respect to that of the control group before the reversal; increased differentially following the reversal; and leveled off one year after the reversal. The gradual increase in credit limit is driven by the fact that opening of credit cards slow: since treated secondary earners get a bigger credit limit increase relative to the control group conditional on opening a credit card, the gradual increase in the treatment effect captures more secondary earners opening credit cards over time both in the treated and the control group.

The event-study results are robust to controlling for linear pre-trends. The blue shaded area in Figure 1-7a denotes the "phase-in" period in which the CFPB announced the policy change and allowed card issuers early compliance with the reversal before the law went into effect. Consistent with this "phase-in" period, the figure shows a differential upward trend in credit limit for the treated group a few months before the reversal. As discussed in Roth (2020), I parametrically control for the linear pretend in event time to ensure that my results are not driven by the pre-existing trend during this phase-in period.[29] Figure A-11a illustrates that a linear pre-trend is a reasonable functional form assumption, and Figure A-11b shows that my results are robust to

[29]This approach has been used widely in event study settings where pre-existing trend may confound the treatment effect. See, for example, Wolfers (2006); Dobkin, Finkelstein, Kluender and Notowidigdo (2018); Goodman-Bacon (2018); Gross, Notowidigdo and Wang (2020); Miller and Soo (2020).

accounting for this differential pre-trend.

1.5.2 Effect on Secondary Earner Consumption

Turning to consumption, the TILA reversal increased secondary earners' consumption. Column 1 through 3 of Table 1.4 show that secondary earners' consumption increased by 14 to 25 percent relative to their pre-reversal average monthly consumption, or \$343 to \$565. The implied marginal propensity to consume (MPC) out of credit limit increases are reported at the bottom of the table, and are in the range of 0.22 to 0.38.[30] Columns 4 through 6 show that using all sample of households – including those where secondary earners do not open a credit card – generate similar MPC estimates. This MPC range is similar in magnitude to that of liquidity constrained consumers documented in the literature. For example, Agarwal, Chomsisengphet, Mahoney and Stroebel (2018) and Gross, Notowidigdo and Wang (2020) estimate the MPC out of credit limits in the range of 0.37 to 0.59 for consumers with high credit risk or those whose bankruptcy flags were removed.[31]

The central result of this paper is that the TILA reversal reduced the consumption gap in the household by increasing secondary earners' share of consumption in the household. Column 1 of Table 1.4 shows that the share of consumption allocated to secondary earners increased by 5 percent relative to their pre-reversal average monthly consumption share, and Columns 2 and 3 show that the estimated effects are larger for stay-at-home spouses and older secondary earners. Put another way, the reversal differentially increased secondary earners' consumption share by 2.2 percentage points (45.1×1.05) more fore the treated group relative to the control group. The shift in consumption toward secondary earners reduced the consumption gap between spouses. The third row of Table 1.4 shows that the consumption gap in the household by 10 to 14 percent of its pre-reversal mean. Taken together, these results suggest that evening out credit disparities between spouses led spouses to share consumption more equally.

Figure 1-7b illustrates dynamic treatment effects on secondary earner's consumption share

[30]To enable comparison to prior work, the implied MPC is calculated by dividing the change in credit card balances (rather than consumption) by the change in credit limits.

[31]Studies that estimate the MPC out of liquidity that do not entail wealth effect document similar estimates (e.g., Johnson, Parker and Souleles, 2006).

from estimating Equation 1.2. The figure shows no detectable trend in secondary earner's consumption share before the reversal, and the share increases over time after the reversal. Figure 1-7c shows that the consumption gap in the household declined over the period when the share of consumption tilted toward secondary earners.

1.5.3 Household Credit, Consumption, and Other Financial Outcomes

I now turn to estimating the effect of the TILA reversal on household consumption and financial outcomes. Table 1.5 reports DiD effects on household credit, consumption, and borrowing. All outcomes are scaled by pre-reversal average monthly household consumption, so the estimated coefficients can be interpreted as a percent increase (or decrease) of pre-reversal consumption.

The reversal expanded credit access at the household-level. The first row of Table 1.5 shows that the reversal increased credit limit available at the household-level by 20 percent relative to pre-reversal household consumption, or \$1,158. The magnitude of the increase similar to the credit limit increase observed for secondary earners, suggesting that the reversal did not crowd out primary earner's access to credit. The increase in household credit is economically meaningful, representing about 24 (11) percent of unconditional (conditional) household credit. Single income households and households where secondary earners are older saw larger increase in household credit limit.

The estimated effect on household consumption is positive but substantially smaller than the effect on secondary earner consumption, providing additional evidence of consumption allocation in the household operating through primary earners cutting back consumption. As shown in the second row of Table 1.5, household consumption increased by 3 percent relative to pre-reversal average monthly household consumption, or roughly \$170. The effect on household consumption is only half as large as the effect on secondary earner consumption (\$340), implying a reduction in primary earner consumption. The MPC out of the household credit limit increase is 0.08, or similar in magnitude to that of high income consumers documented in the literature. For example, Gross and Souleles (2002) and Aydin (2021) estimate the MPC out of credit limits in the range of $0.07 \sim 0.16$ for consumers with high liquidity.

The TILA reversal had a small, positive effect on household credit card debt, but higher debt was not accompanied by worsened financial well-being of the household. The third row of Table 1.5 shows that total household-level credit card revolving balance increased by 0.9 percent relative to pre-reversal average monthly household consumption, or \$51. This implies that roughly a third of the increase in household consumption was financed through debt (i.e., $\frac{0.9}{3}$). Higher debt raises a concern that households may be taking on debt that they will not be able to pay back. However, Table 1.6 shows that a variety of financial solvency outcomes were not materially impacted, including delinquency or overdraft probabilities, or the likelihood of borrowing high-interest loans (e.g., payday or high-cost installment loans). This suggests that households did not fall behind on required monthly debt payments, at least over the two year post-reversal period I analyze. This result is not just a short-run effect of opening a credit card account – Table A.8 shows that the result holds up even when I limit the sample to the second year of the reversal.

I find a small increase in "optimal" credit card repayment behavior for households with multiple credit card accounts. Following Ponce, Seira and Zamarripa (2017) and Gathergood, Mahoney, Stewart and Weber (2019), I construct a debt prioritization indicator that equals one if households optimally repay debt by accruing more debt on credit cards with lower rates and reducing balances on those with higher rates. The fourth row of Table 1.6 shows that the optimal debt repayment behavior among households with multiple credit cards (roughly 55 percent of the card holder sample) increased by 1.2 percentage points. Table A.8 shows that this result is not driven by APR differences between the treated and the control group.

These results highlight the importance of analyzing household behavior through the lens of individual family members. The divergence of consumption effects estimated at the household- and the secondary earner-level, combined with the indicators of financial solvency not deteriorating, is consistent with spouses coordinating their consumption decision to satisfy the family budget constraint. Moreover, comparing secondary earner MPCs to household MPCs indicate that household averages mask the heterogeneity in consumption response of individual family members and fail to capture consumption reallocation in the household. By taking intra-household dynamics into account, I show that the narrowing of a consumption gap in the household did not come at the cost of worsened financial standing.

1.5.4 Alternative Measurement, Sample, and Specification

Measurement of Spouse-Specific Consumption. An important concern for my "broad" consumption measure is that individual spending may not be an accurate proxy for consumption. The broad measure hinges on the assumption that the person who spends the money is the one who consumes it, but this assumption may not hold if couples consume the purchased good together irrespective of who spent the money. I mitigate this concern by constructing a "narrow" consumption measure that only includes spending on gender-assignable goods, such as male or female clothing, that can exclusively be consumed by only one member in the household because of gender. Specifically, I use the Merchant Category Code and identify spending at gender-assignable counterparty to proxy for spending on gender-assignable goods. Table A.9 shows that my results are robust to using this narrower measure.

I consider a battery of alternative "narrow" measures to test the sensitivity of my results. Unlike the "broad" measure, which includes cash spending of individual spouses, one drawback of the narrow measure is that it does not account for potential cash substitution. Specifically, if secondary earners (differentially) become more likely to use credit or debit cards to purchase gender-assignable goods that they used to buy with cash after the reversal, failure to capture cash spending can lead to an upward biased estimate. To address this concern, I construct an alternative "narrow" measure that includes cash withdrawals. This measure assumes that 100 percent of spouses' cash withdrawals were used to buy gender-assignable goods, so if secondary earners reduced cash spending by more than they increased spending on gender-assignable goods, I would not find any effect on secondary earners' consumption. Column 3 of Table A.9 shows that my results still hold up to using the alternative measure. Column 4 considers a slightly broader "narrow" measure that includes spending on goods and services that are more likely to be consumed by one gender versus another (e.g., hair or nail salon for women; gambling or tobacco for men), and show that my results are robust.

Alternative "Broad" Consumption Measures. Table A.10 shows that my main results are robust to using alternative "broad" measures. Because I obtain data from only one financial services provider, a potential measurement concern is that my estimates may be upward-biased if secondary earners increased spending on financial accounts that I observe while reducing spending on accounts that I do not observe. Column 1 shows that my results hold up to including credit card payments to other financial institutions. While this measure does not fully address the concern that there may be substitution across accounts, the fact that married individuals hold significantly fewer debit cards relative to credit cards (see Table A.3) mitigates the concern about bias arising from having incomplete data. In addition, Column 2 confirms that my results are robust to excluding spending on work-related expenses, addressing a potential concern that spending may be correlated with the earner status in the household. Column 3 illustrates that the "broad" measure I use for my main analysis, in fact, accounts for potential cash substitution – the estimates are upward-biased if I do not include cash spending.

Specification robustness. Table A.11 shows that my main findings are robust to using alternative specifications. Compared to Column 1, which reports my preferred baseline specification, Columns 2 through 4 show that my estimates are not sensitive to the choice of fixed effects. This illustrates that my empirical strategy is not subject to the "negative weighting" problem that can arise in staggered DiD settings.[32] Column 5 shows that my results are also robust to controlling state-specific time trends, such as local economic trends. Column 6 shows that estimating effects in quarterly frequency by time-aggregating monthly data yields similar estimates, suggesting that my main results are not an artifact of the timing of income and expenditure commitments being misaligned.[33]

[32]Recent advances in econometric theory point to potential pitfalls associated with estimates from two-way fixed effects specifications in a staggered adoption DiD design (de Chaisemartin and D'Haultfœuille, 2020; Callaway and Sant'Anna, 2020; Sun and Abraham, 2020; Athey and Imbens, 2021; Borusyak, Jaravel and Spiess, 2021; Goodman-Bacon, 2021) Since the empirical setting considered in this paper has simultaneous absorbing treatment in which treatment happens in a single date and the never-treated group, OLS estimation does not suffer from negative weights or under-identification problem (Borusyak, Jaravel and Spiess, 2021).

[33]Baugh and Wang (2018) find that households are more likely to experience cash shortfalls if they have a greater mismatch between the timing of income and expenditure commitments.

Measurement of Income. As discussed in Section 1.4.2, I assume that any income deposited into a couple's joint checking account is income earned by the primary account holder. This assumption can misclassify who is the primary earner in the household to the extent that primary earner is not the primary account holder of the joint account. I address this issue by re-estimating the effect on secondary earners' consumption share using a subset of couples who do not have any joint accounts and thus are not subject to this mis-classification concern. Columns 1 and 2 of Table Table A.12 show that the estimated effects are *larger* restricting the sample to couples that do not have joint checking or any shared accounts, suggesting that the mis-classification likely attenuates the true effect on consumption reallocation.

Other Samples and Placebo Tests. I show that my results are generalizable to a broader population of existing card holders. As discussed in Section 1.4.1, the analysis sample used in this study is restricted to households where secondary earners did not have a sole credit card account at the beginning of my sample period. Column 4 of Table A.12 shows that my results hold up to using a broader sample that includes existing card holders, suggesting that my results are generalizable to a broader population. However, the estimated effects are smaller in the broader sample relative to my main analysis sample, confirming that existing card holders do not regularly update their income and thus less affected by the reversal.

Columns 5 thorough 7 of Table A.12 show placebo tests. Column 5 shows that secondary earners' consumption share did not change for households where neither of the spouses experienced any change in access to credit, corroborating the interpretation that the shift in consumption toward secondary earners is driven by changes in borrowing capacity. Column 6 reports DiD placebo tests and shows no detectible effect in the pre-period. Finally, Column 7 shows that the reversal did not have differential impact on primary earners' access to credit.

1.6 Mechanism for TILA Consumption Effects

In this section, I explore potential mechanisms that explain this paper's main finding that the reversal shifted consumption toward secondary earners.

1.6.1 Private and Public Consumption

The consumption reallocation in the household operated through secondary (primary) earners increasing (reducing) their private consumption. Column 1 of Table 1.7 shows that secondary earner (household) consumption of "private" goods – i.e., goods that are consumed privately such as clothing – increased by 8.5 (1.3) percent relative to secondary earners' (households') pre-reversal average monthly consumption. Interpreted in dollars, the increase in secondary earner private consumption is substantially larger ($214) than the increase in household private consumption ($72), suggesting that primary earners reduced their consumption of private goods by $142 (72-214). These results are consistent with the prediction of the collective model of the household that spouses' relative bargaining power should determine how much private consumption is allocated to them. Specifically, to the extent that the reversal increased secondary earners' marital bargaining power through higher borrowing capacity, theory predicts that secondary earners' private consumption should unambiguously increase with their bargaining power, while that of primary earners should decrease with their bargaining power (Browning, Chiappori and Weiss, 2014). See Section A.1 for derivation of this prediction in a simple case with Cobb-Douglas preferences.

How should households' demand for "public" goods – i.e., goods that are jointly consumed by primary and secondary earners such as child care– change as spouses reallocate consumption in the household? This is ultimately an empirical question because theoretical prediction is ambiguous. Specifically, theory predicts that household consumption of public goods should increase only if the marginal willingness to pay on public goods of the spouse experiencing an increase in bargaining power (i.e., secondary earner) is more sensitive to changes in credit access than that of the other spouse (Blundell, Chiappori and Meghir, 2005). The fourth row of Table 1.7 shows that household public consumption increased by 1.5 percent relative to pre-reversal average monthly household consumption, or $85. This suggests that household consumption behavior changed in a way that better reflects the preferences of secondary earners, consistent with prior work in the intra-household literature (Duflo, 2003). Columns 2 through 6 show that the results on private and public consumption hold in other samples. Table A.4 details how I categorized spending types.

I find suggestive evidence that the reversal increased couples' home production. Figure 1-

8 decomposes the change in household consumption into detailed spending categories. The top 3 spending contributors include spending on home improvement (e.g., home or garden supply stores, florists, etc), groceries, and utilities (e.g., electric, gas, water, etc).[34] Figure A-12 shows that the composition of secondary earner consumption largely mirrors the change in household consumption, suggesting that secondary earners spent more time to cook and improve their home following the reversal. This result supports a key insight from the household model under limited commitment that providing downside insurance (e.g., credit access) to the lower earning partner should lead to greater specialization in the household by improving couples' marital commitment. See, for example, Pollak (2011) for theory and Lafortune and Low (2020) for empirical evidence showing how homeownership (i.e., a commitment technology that can be divided in case of divorce) leads to greater specialization.

1.6.2 The Limited-Commitment Channel

The limited-commitment (LC) channel has the potential to explain my main findings. Under the canonical collective-household model with LC, since spouses cannot precommit to future allocations of resources, factors that improve the outside option (that is, the value of being divorced) of spouses with lower initial bargaining power should shift consumption allocation in their favor to satisfy their participation constraints in marriage. Improving outside options is relevant even for couples that are not on the verge of divorce insofar as there is some risk of divorce because outside options provide insurance against potential downside risks. That said, in the cross-section, the LC channel generates two testable predictions. First, the effect of the reversal on secondary earners' consumption share should be larger (smaller) for couples with weaker (stronger) marital commitment because they will be more (less) sensitive to changes in the outside option. Second, spouses should have incentives to maintain bargaining power by not accruing too much debt on their sole credit card accounts while borrowing more on accounts that are held by the other spouse.

Heterogeneity analysis show results consistent with the first prediction that the estimated effect on secondary earners' consumption share should be more (less) pronounced for couples that are

[34] The number indicates each category's size of the change in terms of the overall effect on consumption, and they sum to 100. For example, the increase in home improvement explains 60% of the overall consumption effect.

more (less) likely to divorce. Columns 1 and 2 of Table 1.8 show that the estimated effect is roughly three times larger for couples with high unconditional probability of divorce relative to those with low probability, where "high" ("low") refers to couples that reside in states with above the top (below the bottom) tercile of pre-reversal state-level divorce rates. Columns 3 and 4 similarly show that the estimated effect is 40 percent larger for couples that have marital problems relative those without. I classify couples with positive pre-reversal spending on counseling (including couple counseling) or dating services as those with marital problems. Columns 5 and 6 show that the estimated effect is 50 percent smaller for couples that have stronger marital commitment (i.e., have children). I classify couples with above median pre-reversal spending on children's clothing or childcare as those with children.

Analysis of credit card utilization rates show results consistent with the second prediction that spouses should have incentives to keep their credit lines available to preserve bargaining power. Table A.13 presents the effect of the reversal on revolving balance utilization rates for credit cards where secondary earners are primary account holders (Panel A), and for cards where primary earners are primary account holders (Panel B). This analysis focuses on households with multiple credit card accounts that carry the same interest rate to ensure that borrowing behavior actually captures strategic motives rather than price effects. Specifically, in the absence of strategic motives, borrowing behavior across these accounts should be similar because the borrowing cost is the same. The first row of each panel shows that spouses reduced utilization rates on their sole credit card accounts, while the second row of each panel shows that they accrued more revolving debt on their joint accounts. These results are consistent with spouses actively reducing debt on their sole accounts while accruing more debt on joint accounts to preserve better bargaining position.

Several institutional features make credit a plausible factor that increases outside options in practice. First, credit limits are "portable" in the sense that secondary earner or stay-at-home divorcees can keep high credit limits that they obtained using household income during marriage even after divorce because credit card issuers are prohibited from adjusting the account holders' credit limits based on their marital status.[35] Therefore, as long as divorcees are not financially delinquent (i.e., able to make a minimum monthly payment), their credit limits will not change after divorce.[36] Second, since having a sole credit card account helps to build one's own credit history, secondary earners' access to credit can improve after divorce. Finally, while debt obligations are divided between spouses upon divorce according to marital division-of-property laws, credit *limits* do not get contested in divorce proceedings. Given these results, it is plausible that credit shapes power dynamics between spouses. Even for couples that are not on the verge of divorce, outside options can discipline harmonious relationship by giving voice and agency to each spouse to act in his or her best interest.

1.6.3 Alternative Channels

I do not find clear cut support for alternative economic channels, such as imperfect information, financial constraint, limited attention, or self-control. Under imperfect information (Wang, 1995), primary earners' consumption should not respond to changes in secondary earners' outcomes because spouses cannot observe the realization of each other's outcomes. Moreover, secondary earners should not change their demand for public consumption if they want to hide their borrowing ability from their spouse.[37] However, this paper finds that consumption reallocation operate through primary earners cutting back consumption (see Sections 1.5.3 and 1.6.1), implying that spouses coordinate their consumption decision to satisfy the family budget constraint. Moreover, 92 percent of couples in my sample have a joint checking account. To the extent that couples pay their credit card bills using their joint checking accounts, monitoring each other's spending behavior should be easy.

My findings suggest that financial constraint is also unlikely to be main driver of my findings. In theory, the shift in consumption toward secondary earners could reflect additional credit inducing

[35] Equal Credit Opportunity Act of 1974 prohibits credit card issuers from making lending decisions based on one's marital status. Exceptions apply when an individual lives in a community property state and when one applies for a joint credit card shared with the other spouse.

[36] Figure A-13 shows changes in financial situation after divorce for divorced men and women. Women's (men's) financial situation tends to improve (deteriorate) after divorce relative to when they were married. This suggests that secondary earner divorcees (typically women) are unlikely to become financially delinquent after divorce.

[37] Ashraf (2009) finds that 21 percent of spouses are willing to pay in order to hide their income from their spouse.

secondary earners to do more shopping than primary earners if couples were at risk of maxing out their existing credit lines. Under this channel, the effect on secondary earners' consumption share should be larger for couples with higher baseline credit card utilization rates relative to those with lower rates. However, the last two columns of Table 1.8 show the opposite result. I find no evidence of consumption reallocation for financially constrained couples and large and significant effect for unconstrained couples. Financial constraint is proxied by having above or below 90th percentile of pre-reversal credit card utilization rates (i.e., initial debt levels). Finally, the sample of households this paper analyzes is not financially constrained on average – conditional having a credit card account, only 55 percent of households carry revolving debt,[38] and the average monthly revolving debt utilization rate is 15 percent.

The findings of this paper are difficult to rationalize by limited attention or self-control. Limited attention and self-control suggest that the increase in secondary earners' consumption share reflects higher borrowing capacity inducing secondary earners to over-spend, either because secondary earners do not pay attention to the family budget constraint or because they have lower self-control in overcoming the impulse to indulge (Banerjee and Mullainathan, 2010; Vissing-Jorgensen, 2021). Under these channels, the effect on secondary earners' consumption should be driven by increased spending on credit cards. However, Tables A.9 and A.10 show that secondary earners' spending on debit cards also increased following the reversal, despite the fact that they could have spent out of debit cards even before the reversal. This result is robust to using both "broad" measure and "narrow" gender-assignable measure.

Overall, the empirical patterns documented in this paper are consistent with the LC channel and do not provide clear cut support for other channels. While the LC channel appears to be main driver of my findings, reduced-form estimates do not reveal quantitative importance of this channel. In the next section, I present a model of household behavior under limited commitment to assess the quantitative importance of this channel.

[38]Revolvers are defined as having positive revolving debt for at least 25 percent of all months that a household has a credit card account.

1.7 Quantitative Analysis

How much of the increase in secondary earners' consumption share can be accounted for by the LC channel? In this section, I calibrate a model of household decision-making under limited commitment that incorporates key aspects of the reversal to evaluate the quantitative importance of my proposed mechanism. I closely follow and build on Mazzocco, Ruiz and Yamaguchi (2014) and Voena (2015) by allowing spouses to have different borrowing limits that they can keep after divorce.[39] Higher borrowing limits translate to higher outside options by relaxing borrowing constraints in case of divorce – this is the key mechanism that drives changes in the bargaining power and consumption allocation in the household.

1.7.1 Set-up

The household consists of two spouses, primary and secondary earners, indexed by $i \in (P, S)$, who live until T. In each month t, the spouses decide jointly how much to save, consume, and whether to work and divorce. The spouses have complete knowledge of all variables and preferences dated t and earlier and of probability distributions over all variables in $t' > t$.

Preferences Each spouse has preferences that are separable over time and across states, with diminishing marginal utility over consumption $u(c_t)$ and disutility, ψ, from labor market participation, P_t^i. Each spouse's period utilities take the form $u_t^{i,M} = \frac{c^{i1-\gamma}}{1-\gamma} - \psi P_t^i + \xi_t$ in marriage and $u_t^{i,D} = \frac{c^{i1-\gamma}}{1-\gamma} - \psi P_t^i$ in divorce, where ξ_t is a taste shock for marriage that follows a random walk, capturing the persistence in the taste for marriage such as the spouses' affection for one another: $\xi_t = \xi_{t-1} + \epsilon_t$ and $\epsilon_t \sim_{iid} N(0, \sigma_\xi^2)$. Primary earner always works ($P_t^P = 1$) and incurs disutility ψ, while secondary earner can choose to work. The spouses have identical discount factor, β, and beliefs.

In marriage, spouses benefit from economies of scale in consumption. Specifically, total household expenditure is given by a constant elasticity of substitution aggregator of primary and sec-

[39]This model builds on the literature on risk-sharing and household decision-making under limited commitment (Kocherlakota, 1996; Ligon, Thomas and Worrall, 2002; Mazzocco, 2007).

ondary earners' consumption: $x = [(c^P)^\rho + (c^S)^\rho]^{\frac{1}{\rho}} e(k)$. For $\rho \geq 1$, the couple gets more utility jointly from the same level of spending because there are gains from marriage.[40] The couple devotes a fraction, $e(k)$ denoting an equivalence scale, of total household expenditures on children. The economies of scale and the cost multiplier take into consideration the existence of goods that are public within the household.[41]

Income The income process, y_t^i, has two components – an endogenous component (h_t^i), and an exogenous component (z_t^i), which is correlated between spouses:

$$ln(y_t^i) = ln(h_t^i) + z_t^i \tag{1.3}$$

where the income shock follows a random walk: $z_t^i = z_{t-1}^i + \zeta_t^i$, in which $\zeta_t^i \sim_{iid} N(0, \sigma_{\zeta^i}^2)$. The law of motion for each spouse's human capital is $ln(h_t^i) = ln(h_{t-1}^i) - \delta \cdot (1 - P_{t-1}^i) + (\lambda_0^i + \lambda_1^i \cdot t) \cdot P_{t-1}^i$, such that human capital depreciates at a rate δ if a spouse does not work in the previous period or appreciates with tenure at a rate $\lambda_0^i + \lambda_1^i \cdot t$.

Budget Constraints Saving (borrowing), a_t^i, earns (pays) the market rate, $\tilde{r} > 0$. The budget constraints in marriage and divorce are:

$$\begin{aligned} A_{t+1} - (1+\tilde{r})A_t &= Y_t - x_t && \text{if married} \\ a_{t+1}^i - (1+\tilde{r})a_t^i &= y_t^i \cdot P_t^i - c_t^i \cdot e(k) && \text{if divorced} \end{aligned}$$

where $A_t = \sum_{i=P}^S a_t^i$, $Y_t = \sum_{i=P}^S y_t^i \cdot P_t^i$, and x_t denote total household savings, income, and expenditure. While married, the couple allocates A_t between one another according to their respective bargaining power (θ_t^i) in each period because divorce is possible. Therefore, in the first period after divorce, each spouse enters t with $a_t^i = \theta_{t-1}^i A_{t-1}$. After divorce, spouses live off their individual

[40] The CES consumption aggregator is a standard assumption. See, for example, Boerma and Karabarbounis (2021) for home production model and Knowles (2013) for intra-household model.

[41] This is a short-hand way to allow for public consumption for married couples. I adopt this approach for tractability because determining the relative shares of public vs. private consumption is not the primary focus of the model. An alternative way to take public consumption into account is by allowing individual spouses to have different preferences over public and private consumption Mazzocco (2007).

financial resources and contribute to the consumption of their children as a fraction of their own consumption, according to $e(k)$. Spouses pay higher interest rate when they borrow, but earn lower rate when they save:

$$\tilde{r} = \begin{cases} \overline{r}, & \text{if } a_t^i < 0 \\ \underline{\mathrm{r}}, & \text{otherwise} \end{cases}$$

Borrowing Limits A key feature in this paper is that spouses have individual borrowing limits. Couples face an exogenous borrowing limit L_t that depends on the TILA regime. Specifically, the TILA regime is modeled to capture the stylized feature of the reversal that the borrowing capacity of the secondary earner is higher after the reversal.

$$A_{t+1} \geq -L_t \qquad \text{if married} \tag{1.4}$$

$$L_t = \begin{cases} L^P + \underline{\mathrm{L}}^S, & \text{if } t < \text{TILA reversal} \\ L^P + \overline{L}^S, & \text{otherwise} \end{cases} \tag{1.5}$$

The borrowing constraint imposes limits on the couples' "net worth" (i.e., assets minus liabilities) and can be interpreted as maximum credit card debt that the couple can cumulate.[42] In case of divorce, each spouse retains the entirety of individual borrowing limit, L_t^i:

$$a_{t+1}^i \geq -L_t^i \qquad \text{if divorced} \tag{1.6}$$

This "portability" feature of borrowing limit is what makes individual borrowing capacity relevant for shaping marital bargaining power.[43] In practice, because borrowing limit is an uncontested financial resource (unlike income, assets, or debt) that belongs to the individual spouse that holds a

[42]I set the income grid ensure that preferences satisfy the Inada condition when borrowing is allowed. Specifically, I obtain the minimum required income to borrow up to $\overline{L}$ without violating the non-negativity of consumption using the "natural" borrowing constraint suggested by Aiyagari (1994): $Y_{min} = r \cdot \overline{L}$. I then set the income grid to range from values higher than this minimum.

[43]In practice, divorcees can keep credit limits that they obtained during marriage as long as they are able to make minimum monthly payment. This is because card issuers do not treat divorce any differently than other life events that could trigger financial distress, such as job loss or illness.

credit card account, borrowing capacity translates into higher outside options by allowing spouses to smooth consumption over time.

1.7.2 Decisions and Model Predictions

In each period, the household decision-making problem consists of two stages. In the first stage, each spouse computes the value of being divorced and that of staying married without taking into account participation constraints and using existing bargaining power. In the second stage, each spouse compares the value of being divorced to that of staying married and decides whether to stay married, divorce, or negotiate. If couples negotiate, they compute the value of staying married conditional on the adjusted bargaining power. Thus, the optimal value function for each spouse is determined by comparing the value functions of being divorced and staying married, $V_t^i(\omega) = \max\Big\{V_t^{i,D}(\omega), V_t^{i,M}(\omega)\Big\}$.

Stage 1.a: The Value of Being Divorced To compute the value of being divorced, the problem is solved by backward induction using the terminal condition that each spouse consumes all of his/her assets ($a_{t+1}^i = 0$) given the set of state variables, ${\omega_{\mathbf{T}}}^D = \Big\{a_T^i, h_T^i, z_T^i, \Omega_T\Big\}$:

$$V_T^{i,D}(\omega_{\mathbf{T}}^{\mathbf{D}}) = \max_{c_T^i, P_T^i} u(c_T^i)$$
$$\text{s.t.}$$
$$(1+r)a_T^i = y_T^i \cdot P_T^i - c_T^i \cdot e(k)$$

In the remaining periods $t = 1, ..., T-1$,

$$\begin{gathered} V_t^{i,D}(\omega_{\mathbf{t}}^{\mathbf{D}}) = \max_{c_t^i, a_{t+1}^i, P_t^i}\Big\{u(c_t^i) + \beta E\Big[V_{t+1}^{i,D}(\omega_{\mathbf{t+1}}^{\mathbf{D}}|\omega_{\mathbf{t}}^{\mathbf{D}})\Big]\Big\} \\ \text{s.t.} \\ a_{t+1}^i - (1+r)a_t^i = y_t^i \cdot P_t^i - c_t^i \cdot e(k) \\ a_{t+1}^i \geq -L_t^i \end{gathered} \tag{1.7}$$

given state variables ${\omega_{\mathbf{t}}}^{D} = \left\{a_t^i, h_t^i, z_t^i, \Omega_t\right\}$. I assume that spouses do not remarry after divorce. Ω_t represents the vector of the TILA regime at time t.

Stage 1.b: The Value of Staying Married To compute the value of staying married, the couple first solves the household value function, V_t^M. Then the spouses compute their individual value of staying married, $V_t^{i,M}$ using the optimal choice of consumption, labor supply, and savings decisions from this household problem.

To compute the household value function, the couple chooses the control vector in the terminal period that maximizes the weighted sum of their individual utilities, where the weights are given by the Pareto weights θ_T^i (i.e., bargaining power):

$$V_T^M(\omega_{\mathbf{t}}^{\mathbf{M}}) = \max_{c_T^P, c_T^S, P_T^S}\left\{\theta_T^P u(c_T^P) + \theta_T^S u(c_T^S)\right\}$$

$$\text{s.t.}$$

$$(1+r)A_T = Y_T - x_T$$

where $\omega_{\mathbf{T}}^{\mathbf{M}} = \left\{A_T, h_T^S, z_T^P, z_T^S, \theta_T^P, \theta_T^S, \xi_T, \Omega_T\right\}$ and requiring $A_{T+1} = 0$. The state variables capture current assets, secondary earner's human capital, income shocks, bargaining power, taste for marriage shock, and the TILA regime.

In the remaining periods $t = 1, ..., T-1$, the couple solves:

$$V_t^M(\omega_{\mathbf{t}}^{\mathbf{M}}) = \max_{c_t^P, c_t^S, P_t^S, a_{t+1}^P, a_{t+1}^S}\left\{\theta_t^P u(c_t^P) + \theta_t^S u(c_t^S) + \beta E[V_{t+1}^M(\omega_{\mathbf{t+1}}^{\mathbf{M}}|\omega_{\mathbf{t}}^{\mathbf{M}})]\right\}$$

$$\text{s.t.}$$

$$A_{t+1} - (1+r)A_t = Y_t - x_t$$

$$A_{t+1} \geq -L_t$$

given state variables $\omega_{\mathbf{t}}^{\mathbf{M}} = \left\{A_t, h_t^S, z_t^P, z_t^S, \theta_t^P, \theta_t^S, \xi_t, \Omega_t\right\}$. The initial bargaining power of each spouse, θ_0^i is determined exogenously and can be considered a bargaining structure that spouses agreed on (but did not commit to) at the time of household formation. The values of the Pareto weights may reflect factors that influence the decision process–such as relative financial resource–that are known and predicted at $t = 0$ (Chiappori and Meghir, 2015); capture values that clear the marriage market (Choo, Seitz and Siow, 2008); or can result from noncooperative threat points (Lundberg and Pollak, 1993).

Spouses consume and save jointly when computing the household value function V_t^M, but they allocate consumption and savings between one another according to θ_t^i because divorce is possible. They use this individual consumption and saving, $c_t^{i,*}$ and $a_t^{i,*}$, to compute $V_t^{i,M}$. Then given a sequence of optimal solutions $\forall\ \omega^M$, $\{c_t^{i,*}(\omega^{\mathbf{M}}), P_t^{i,*}(\omega^{\mathbf{M}}), a_{t+1}^{i,*}(\omega^{\mathbf{M}})\}_{t=1}^{T}$, the value of staying married for each spouse:

$$V_t^{i,M}(\omega_{\mathbf{t}}) = u(c_t^{i,*}(\omega_{\mathbf{t}}^{\mathbf{M}}), P_t^{i,*}(\omega_{\mathbf{t}}^{\mathbf{M}}); \xi_t^i) + \beta E[V_{t+1}^{i,M}(\omega_{\mathbf{t+1}}^{\mathbf{M}})] \tag{1.8}$$

The married couple's optimal value function is the weighted sum of each spouses' value functions, where the weights are the bargaining power from $t - 1$:

$$V_{t+1}^{M,*}(\omega_{\mathbf{t+1}}^{\mathbf{M}}) = \theta_t^P V_{t+1}^{P,*}(\omega_{\mathbf{t+1}}^{\mathbf{M}}) + \theta_t^S V_{t+1}^{S,*}(\omega_{\mathbf{t+1}}^{\mathbf{M}}) \tag{1.9}$$

Stage 2: The Divorce Choice Problem In the second stage, each spouse compares the value of being divorced ($V_t^{i,D}$) to the value of staying married ($V_t^{i,M}$). Three possible cases may arise:

1. The participation constraints are satisfied for both spouses, so it is optimal to stay married:

$$V_t^{i,M} > V_t^{i,D} \qquad \forall\ i \tag{1.10}$$

   In this case, spouse i's value function is $V_t^{i,M}$ is from the stage 1.b problem.

2. The participation constraints are binding for both spouses, so it is optimal to divorce:

$$V_t^{i,D} > V_t^{i,M} \qquad \forall\ i \tag{1.11}$$

   In this case, spouse i's value function is $V_t^{i,D}$ from the stage 1.a problem.

3. One spouse prefers to stay married but the other spouses' participation constraint binds. Suppose that only secondary earner's participation constraint binds so that it is optimal for primary earner to stay married but secondary earner prefers to divorce:

$$
\begin{aligned}
V_t^{P,M} &> V_t^{P,D} \\
V_t^{S,M} &\leq V_t^{S,D}
\end{aligned}
\tag{1.12}
$$

In this last case, the couple solves the stage 1.b. problem again *under the constraint that secondary earner's participation constraint is satisfied*. In the terminal period:

$$
V_T^M(\omega_{\mathbf{T}}^{\mathbf{M}}) = \max_{c_T^P, c_T^S, P_T^S, \theta_T^S} \left\{ \theta_T^P u(c_T^P) + \theta_T^S u(c_T^S) \right\}
$$

s.t.

$$
(1+r)A_T = Y_T - x_T \quad \text{and} \quad A_{T+1} = 0
$$

$$
u(c_T^S) = V_T^{S,D} \tag{1.13}
$$

$$
\theta_T^S = \theta_{T-1}^S + \lambda_T^S \tag{1.14}
$$

Equation 1.13 imposes secondary earner's value of staying married to be as good as the outside option. This constraint can be incorporated directly in the objective function using a standard Lagrangian multiplier method. Let λ_T^S denote the Lagrangian multiplier associated with secondary earner's participation constraint. Whenever the participation constraint binds (i.e., $\lambda_T^S > 0$), secondary earner's bargaining power increases by λ_T^S in order to make secondary earner indifferent between divorcing and staying married (Eq. 1.14).

In other periods:

$$V_t^M(\omega_{\mathbf{t}}^{\mathbf{M}}) = \max_{c_t^P, c_t^S, P_t^S, A_{t+1}, \theta_t^S}\left\{\theta_t^P u(c_t^P) + \theta_t^S u(c_t^S) + \beta E[V_{t+1}^M(\omega_{\mathbf{t+1}}^{\mathbf{M}}|\omega_{\mathbf{t}}^{\mathbf{M}})]\right\}$$

s.t.

$$A_{t+1} - (1+r)A_t = Y_t - x_t \quad \text{and} \quad A_{t+1} \geq -L_t^S$$

$$u(c_t^S) + \beta E[V_{t+1}^{S,M}(\omega_{\mathbf{t+1}}^{\mathbf{M}}|\omega_{\mathbf{t}}^{\mathbf{M}})] = V_t^{S,D} \tag{1.15}$$

$$\theta_t^S = \theta_{t-1}^S + \lambda_t^S \tag{1.16}$$

Then given a sequence of optimal solutions to this constrained Pareto problem $\forall$ ω^M, $\{c_t^{i,**}(\omega^{\mathbf{M}}), P_t^{i,**}(\omega^{\mathbf{M}}), a_{t+1}^{i,**}(\omega^{\mathbf{M}}), \theta_t^{i,**}(\omega^{\mathbf{M}})\}_{t=1}^T$, each spouse's value function is:

$$V_t^{i,M}(\omega_{\mathbf{t}}) = u(c_t^{i,**}(\omega_{\mathbf{t}}^{\mathbf{M}}), P_t^{i,**}(\omega_{\mathbf{t}}^{\mathbf{M}}); \xi_t^i) + \beta E[V_{t+1}^{i,M}(\omega_{\mathbf{t+1}}^{\mathbf{M}})] \tag{1.17}$$

The couple repeats the two stage problem again if it enters period t as married. If spouses enter t as divorcees, they solve the first stage problem for the remaining period using assets that they divided according to θ_t^i in the previous period.

Note that threat of divorce triggers a renegotiation that modifies the consumption allocation plans of the married couple – that is, in the last case, the optimal consumption allocation is such that the new plan is as good as each spouse's outside option. In equilibrium, divorce occurs when the joint surplus–the sum of the two spouses' marriage surpluses–is negative.[44]

Key Model Prediction A key prediction that the limited commitment model generates is that increasing a spouse's outside option leads to a shift in consumption allocation in his or her favor because higher outside options increase marital bargaining power to satisfy participation constraints.

[44] Divorce does not require negative surplus for both spouses because divorce can happen even when one of the spouses want to stay married but there is not enough resource to transfer to the other spouse that would make the other spouse indifferent between staying married and being divorced.

This prediction can be shown using the first-order-condition with respect to c_t^i:

$$\frac{u'\left(c_t^{P*}\right)}{u'\left(c_t^{S*}\right)} = \frac{\theta_t^S + \lambda_t^S}{\theta_t^P + \lambda_t^P} \\ = \gamma_t \tag{1.18}$$

See Appendix Section A.2 for derivation of this prediction. This condition shows that the ratio of marginal utilities of consumption has a one-to-one relationship to the relative bargaining power of the spouses, or the slope of the Pareto frontier (Kocherlakota, 1996). Thus, as in Equation 1.15, whenever the secondary earner's Lagrangian multiplier associated with her participation constraint is positive, $\lambda_t^S > 0$, primary earner's marginal utility will be relatively higher than that of secondary earner. This implies an increase in secondary earners' consumption share.

1.7.3 Quantitative Results

Using the model presented above, I attempt to answer three questions: (1) by how much did secondary earners' *relative bargaining power* increase in the household? (2) how much of the observed increase in secondary earners' consumption share can be accounted for by the *LC channel*? (3) how large are the *welfare gains* from the reversal? I use a sufficient statistics approach to answer the first and a calibration approach to answer the other two questions.

Changes in the Bargaining Power I use a sufficient statistics approach to document the size of the change in secondary earners' relative bargaining power in the household as a result of the 2013 reversal. While the spouses' relative bargaining power is not observed in the data, Equation 1.18 shows that – under certain assumptions about the spouses' preferences – the relative bargaining power (the right-hand-side) can be characterized by observable elements of household behavior: spouse-specific consumption. I obtain average monthly consumption of primary and secondary earners in the treated group before and after the reversal to quantify the size of the change in secondary earners' relative bargaining power.

Figure 1-9 illustrates that the reversal led to an economically meaningful increase in secondary earners' relative bargaining power in the household. Assuming that both spouses have the CRRA utility with relative risk aversion $\gamma = 1.5$, Panel 1-9a shows that the slope of the Pareto frontier before the reversal was -0.71. After the reversal, Panel 1-9b shows that this slope became steeper – i.e., the relative bargaining power tilted toward the secondary earner. The change in the relative bargaining power is 36 percentage points (ppt): 1.07 - 0.71. Is this change economically meaningful? In the pre-reversal period, the average monthly change in the slope among the card holder sample was 28 ppt, implying that the reversal led to a 29 percent increase in the relative bargaining power ($\frac{36-28}{28}$) over and above the shift in the bargaining power from opening a credit card account. In addition, the average monthly change in the slope is close to 0 (2 ppt) among a broader sample of households that includes secondary earners that did not open a credit card account. This illustrates that the relative bargaining power does not typically vary over-time, but an increase in secondary earners' borrowing capacity generates an economically meaningful shift in their marital bargaining power.[45]

The Limited Commitment channel I use a calibration approach to quantify the extent to which the limited commitment channel can explain the observed increase in secondary earners' consumption share in the data. Table 1.9 reports the description, source, and value for the parameters used in this exercise. I obtain parameters from existing literature, and where possible, directly from the data used in this study. I set each period to be one month and track household consumption behavior for 36 months – 12 months before and 24 months after the reversal – to match the data. Table A.14 compares the outcomes generated by the model and observed in the data, and Table A.15 shows how the model-generated outcomes change before and after the reversal. These tables show that the model generates reasonable estimates of optimal consumption and borrowing behavior compared to data.

Figure 1-10 shows that the LC channel can explain a third (32%) of the observed increase in the secondary earners' consumption share. The figure shows that secondary earners' average monthly consumption share observed in the data (red) is higher than the path of optimal consumption share

[45]To be clear, this does not imply that the spouses' *level of consumption* stays constant if secondary earners do not open a credit card account. It implies that the share of consumption does not change prior to the reversal.

generated in the model (blue).[46] Given that the model-generated increase accounts for roughly one-third of the observed increase in the data, I conclude that the LC channel is quantitatively important. I interpret this model-generated magnitude as a lower bound because it is likely attenuated by the fact that the model does not capture aspects of credit cards that are important in reality. For example, while the model only captures the *credit* aspect of credit cards[47], other aspects include the convenience value that allows individuals to use credit cards to smooth cash flows even if they pay off their debt in full at the end of the billing cycle, the ability to make minimum monthly payment, and the grace period between billing cycle and payment cycle. Therefore, the quantitative importance of credit in shaping within-household consumption allocation via the LC channel is likely to be attenuated in the model.

Welfare What is the welfare benefit of the reversal? To answer this question, I calculate the Consumption Equivalent Variation (CEV), or the percent of expected lifetime consumption that a spouse inhabiting economy without the reversal would pay *ex ante* in order to inhabit economy with the reversal. In this model, since I track household consumption behavior for only 36 months around the reversal, the CEV captures the percent of expected consumption over this 3 year period. I consider two economies, $k = \{1, 2\}$, where $k = 1$ refers to the regime without the reversal and $k = 2$ refers to the regime with the reversal. I define ex-ante welfare in economy k derived from steady state consumption and work decisions $\{c_t^{i,k}(\omega), P_t^{i,k}(\omega)\}_{t=1}^{T}$ over states $\omega_t = \{a_t^{i,k}, h_t^{i,k}, z_t^{i,k}, \theta_t^{i,k}, \xi_t^k, \Omega_t^k\}$ distributed with $\lambda_t^i(\omega)$ as:

$$S^{i,k} = U(c^{i,k}; \xi^k) - V(P^{i,k}; \xi^k) \tag{1.19}$$

[46] I obtain the red line by applying the DiD estimates shown in Figure 1-7 to the pre-reversal average monthly consumption share generated in the model to standardize the level of the consumption share.

[47] In the model, spouses cannot borrow while holding liquid assets because borrowing is modeled as negative assets. In reality, however, roughly 55% of couples in my sample have revolving credit card debt even if they have enough checking account balance to pay off the debt (i.e., they co-hold).

where ex ante utility over allocations for each of the two spouses,

$$U(c^{i,k};\xi^k) \equiv \int E_0\Big[\sum_{t=1}^{T}\beta^{t-1}u(c_t^{i,k};\xi_t^k)\Big]d\lambda^k \tag{1.20}$$

$$V(P^{i,k};\xi^k) \equiv \int E_0\Big[\sum_{t=1}^{T}\beta^{t-1}(\psi P_t^{i,k};\xi_t^k)\Big]d\lambda^k \tag{1.21}$$

Then the CEV, denoted by Δ_{CEV}, is:

$$S^i\Big((1+\Delta_{CEV}^i)c^{i,1}, P^{i,1}\Big) = S(c^{i,2}, P^{i,2}) \tag{1.22}$$

which can be expressed as

$$(1+\Delta_{CEV}^i)^{1-\gamma}U(c^{i,1}) - V(P^{i,1}) = U(c^{i,2}) - V(P^{i,2}) \tag{1.23}$$

or rewritten,

$$1+\Delta_{CEV}^i = \Big[\frac{U(c^{i,2})}{U(c^{i,1})} + \Big(\frac{V(P^{i,1})}{V(P^{i,2})} - 1\Big)\cdot\frac{V(P^{i,2})}{U(c^{i,1})}\Big]^{\frac{1}{1-\gamma}} \tag{1.24}$$

Δ_{CEV}^i captures spouse i's percent of expected 3-year period consumption that i would be willing to pay ex ante to inhibit an economy with the reversal instead of an economy without the reversal.[48] Table 1.10 shows that the TILA reversal is Pareto improving in the sense that both primary and secondary earner are willing to pay a positive share of their expected consumption to inhabit an economy with the reversal. I find that secondary earner's CEV is higher than that of primary earner, consistent with the reversal mainly benefitting secondary earners. The well-being of the couple as a

[48] I similarly calculate the household CEV by defining household ex ante social welfare criterion as the sum of the two spouses ex ante utility over allocations:

$$S^k = U(c^k;\xi^k) - V(P^k;\xi^k) \tag{1.25}$$

$$U(c^k;\xi^k) \equiv \int E_0\Big[\sum_{t=1}^{T}\beta^{t-1}u(c_t^{i,P};\xi_t^k)\Big]d\lambda^k + \int E_0\Big[\sum_{t=1}^{T}\beta^{t-1}u(c_t^{S,k};\xi_t^{S,k})\Big]d\lambda^k \tag{1.26}$$

$$V(P^k;\xi^k) \equiv \int E_0\Big[\sum_{t=1}^{T}\beta^{t-1}(\psi P_t^{P,k};\xi_t^k)\Big]d\lambda^k + \int E_0\Big[\sum_{t=1}^{T}\beta^{t-1}(\psi P_t^{S,k};\xi_t^{S,k})\Big]d\lambda^k \tag{1.27}$$

whole also increases in that the couple is willing to pay 2.7 percent of their expected consumption to inhabit an economy with the reversal. Overall, this result indicates that increasing secondary earners' borrowing capacity improves the well-being of individual members as well as the couple as a whole.

One caveat of the welfare analysis is that it does not take into account potential negative externalities arising from higher likelihood of divorce. Table A.15 shows that the reversal increased the couples' likelihood of divorce by 1 ppt. Divorce can have negative implications for the well-being of the children and can expose card issuers to losses to the extent that divorce is associated with subsequent financial distress. The CEV analysis does not fully take into account such negative effects.

1.8 Conclusion

The provision of the Truth-in-Lending Act (TILA) concerning independent ability to pay was reversed in 2013 to facilitate access to credit for secondary earners and stay-at-home spouses who have limited income of their own but have access to household income. I exploit the fact that the 2013 reversal was superseded by state-level marital division-of-property laws in some states but not others to gain identification and leverage administrative financial-transaction data that measure credit and consumption of each spouse. This allows me to examine whether reducing disparities in credit between spouses reduces consumption disparities in the household.

My central finding is that the reversal – which increased secondary earners' credit limits by \$1,024 – reduced the pre-reversal consumption gap in the household by 10 percent. Consumption shifted toward secondary earners, whose private consumption crowded out primary earners' private consumption. Household spending on "public" consumption increased following the reversal, suggesting that primary earners indirectly benefited from changes in household consumption patterns that reflect the preferences of secondary earners. Household credit card borrowing increased moderately, but a variety of financial-solvency outcomes were not materially impacted. The limited-commitment channel, which posits that higher borrowing capacity strengthens marital bargaining power, appears to best explain the empirical patterns documented in this paper. Because financial policies can have an uneven impact on individual family members in the household, a key implication of this paper is that policies aimed at reducing financial disparities between spouses can reduce consumption inequality.

I highlight three caveats and corresponding directions in which my work can be extended. First, this paper examined relatively short-run effects of the TILA reversal. Therefore, whether consumption-reallocation and financial-solvency patterns persist in the long run is an open question. Second, this paper took a step toward constructing consumption measures of individual family members, but clearly more can be done to improve the measures' accuracy. More generally, leveraging detailed financial transaction data or administrative tax records to construct within-household economic outcomes can make a meaningful contribution to the family economics and household finance literatures.[49] As Chiappori and Meghir (2015) note, accurate measurement of within-household economic outcomes is crucial for policy-making: "Understanding intrahousehold inequality and, more broadly, intrahousehold allocations is crucial for understanding the effects of policy and for targeting programs designed to alleviate poverty. The implications are far reaching and they span simple questions of who will benefit from certain programs to deeper questions about child poverty and even child development." Finally, this paper analyzed the behavior of couples that represent the traditional family structure of a married man and woman. However, American family structures have changed dramatically over the last few decades, with the rise of same-sex marriage and co-habitation. Analysis of how trends in family structure are associated with within-household inequality would be a fruitful direction of research.

[49]For example, Olafsson and Gathergood (2020), Vihriälä (2020), and Choukhmane, Goodman and O'Dea (2021) use administrative microdata to study intra-household economic outcomes. Baker and Kueng (2021) provide an overview of the types of financial transaction data used in the household finance literature.

Bibliography

12 CFR §1026. 2012. “Federal Register.” Bureau of Consumer Financial Protection.

12 CFR §226. 2011. “Federal Register.” Federal Reserve System.

Abadie, Alberto. 2005. “Semiparametric Difference-in-Differences Estimators.” *Review of Economic Studies*, 72: 1–19.

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. 2018. “Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?” *Quarterly Journal of Economics*, 133(1): 129–190.

Aguiar, Mark, and Mark Bils. 2015. “Has Consumption Inequality Mirrored Income Inequality?” *American Economic Review*, 105(9): 2725–2756.

Angrist, Joshua. 2002. “How Do Sex Ratios Affect Marriage and Labor Markets? Evidence from America’s Second Generation.” *Quarterly Journal of Economics*, 117(3): 997–1038.

Angrist, Joshua, David Autor, Sally Hudson, and Amanada Pallais. 2015. “Evaluating Econometric Evaluations of Post-Secondary Aid.” *American Economic Review: Papers and Proceedings*, 105(5): 502–507.

Argyle, Bronson, Sasha Indarte, Ben Iverson, and Christopher Palmer. 2021. “Explaining Racial Disparities in Personal Bankruptcy Outcomes.” *Working paper*.

Ashenfelter, Orley, and David Card. 1985. “Using the Longitudinal Structure of Earnings to Esimtate the Effect of Training Programs.” *Review of Economics and Statistics*, 67(4): 648–660.

Ashraf, Nava. 2009. “Spousal Control and Intra-household Decision Making: An Experimental Study in the Philippines.” *American Economic Review*, 99(4): 1245–1277.

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

Ashraf, Nava, Dean Karlan, and Wesley Yin. 2010. “Female Empowerment: Impact of a Commitment Savings Product in the Philippines.” *World Development*, 38(3): 333–344.

Athey, Susan, and Guido W. Imbens. 2021. "Design-based analysis in Difference-In-Differences settings with staggered adoption." *Journal of Econometrics*.

Attanasio, Orazio P., and Valérie Lechene. 2014. "Efficient Responses to Targeted Cash Transfers." *Journal of Political Economy*, 122(1): 178–222.

Aydin, Deniz. 2021. "Consumption Response to Credit Expansions: Evidence from Experimental Assignment of 45,307 Credit Lines." *Working paper*.

Baker, Scott, and Lorenz Kueng. 2021. "Household Financial Transaction Data." *NBER working paper*, 29027.

Banerjee, Abhijit, and Sendhil Mullainathan. 2010. "The Shape of Temptation: Implications for the Economic Lives of the Poor." *NBER working paper*, 15973.

Banerjee, Abhijit, Esther Duflo, Rachel Glennerster, and Cynthia Kinnan. 2015. "The Miracle of Microfinance? Evidence from a Randomized Evaluation." *American Economic Journal: Applied Economics*, 7(1): 22–53.

Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace. 2021. "Consumer-Lending Discrimination in the FinTech Era." *Journal of Financial Economics (Forthcoming)*.

Baugh, Brian, and Jialan Wang. 2018. "When Is It Hard to Make Ends Meet?" *Working paper*.

Bau, Natalie, and Raquel Fernández. 2021. "The Family as a Social Institution." *NBER working paper*, 28918.

Begley, Taylor A., and Amiyatosh Purnanandam. 2021. "Color and credit: Race, regulation, and the quality of financial services." *Journal of Financial Economics*.

Bertaut, Carol C., Michael Haliassos, and Michael Reiter. 2009. "Credit Card Debt Puzzles and Debt Revolvers for Self Control." *Review of Finance*, 13(4): 657–692.

Blattner, Laura, and Scott Nelson. 2021. "How Costly is Noise? Data and Disparities in Consumer Credit." *Working paper*.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review*, 98(5): 1887–1921.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. "Consumption Inequality and Family Labor Supply." *American Economic Review*, 106(2).

Blundell, Richard, Magnac Thierry Chiappori, Pierre-André, and Costas Meghir. 2007. "Collective Labor Supply: Heterogeneity and Non-Participation." *Review of Economic Studies*, 74(2): 414–455.

Blundell, Richard, Pierre-André Chiappori, and Costas Meghir. 2005. "Collective Labor Supply with Children." *Journal of Political Economy*, 113(6): 1277–1306.

Bobonis, Gustavo J. 2009. "Is the Allocation of Resource within the Household Efficient? New Evidence from a Randomized Experiment." *Journal of Political Economy*, 117(3): 453–503.

Boerma, Job, and Loukas Karabarbounis. 2021. "Inferring Inequality with Home Production." *Econometrica*, 89(5): 2517–2556.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2021. "Revisiting Event Study Designs: Robust and Efficient Estimation." *Working paper*.

Browning, Martin, Francois Bourguignon, Pierre-André Chiappori, and Valérie Lechene. 1994. "Income and Outcomes: A Structural Model of Intrahousehold Allocation." *Journal of Political Economy*, 102(6): 1067–1096.

Browning, Martin, Pierre-André Chiappori, and Yoram Weiss. 2014. *Economics of the Family.* Cambridge University Press.

Butler, Alexander W., Erik J. Mayer, and James Weston. 2021. "Racial Discrimination in the Auto Loan Market." *Working paper*.

Callaway, Brantly, and Pedro H.C. Sant'Anna. 2020. "Difference-in-Differences with multiple time periods." *Journal of Econometrics*.

Cantor, David, Nancy Mathiowetz, Sid Schneider, and Brad Edwards. 2013. "Redesign Options for the Consumer Expenditure Survey." *Westat*.

Catherine, Sylvain, and Constantine Yannelis. 2021. "The Distributional Effects of Student Loan Forgiveness." *Working paper*.

Centers for Disease Control and Prevention. 2021. "National marriage and divorce rate trends for 2000-2019." *National Center for Health Statistics*.

CFPB. 2019. "The Consumer Credit Card Market." *Bureau of Consumer Financial Protection Report*.

CFPB. 2021. "Code of Federal Regulations, 12 CFR §1026.24." Consumer Financial Protection Bureau.

Chiappori, Pierre-Andre, and Costas Meghir. 2015. "Handbook of Income Distribution: Intrahousehold Inequality." , ed. Anthony B. Atkinson and Francois Bourguignon Vol. 2B, Chapter 16, 1369–1418. Elsevier B.V.

Chiappori, Pierre-André, and Maurizio Mazzocco. 2017. "Static and Intertemporal Household Decisions." *Journal of Economic Literature*, 55(3): 985–1045.

Chiappori, Pierre-Andre, Bernard Fortin, and Guy Lacroix. 2002. "Marriage Market, Divorce Legislation, and Household Labor Supply." *Journal of Political Economy*, 110(1): 37–72.

Choo, Eugene, Shannon Seitz, and Aloysius Siow. 2008. "Marriage Matching, Risk Sharing, and Spousal Labor Supplies." *Working paper*.

Choukhmane, Taha, Lucas Goodman, and Cormac O'Dea. 2021. "Efficiency in Household Decision Making: Evidence from the Retirement Savings of US Couples." *Working paper*.

Conlon, John J., Malavika Mani, Gautam Rao, Matthew Ridley, and Frank Schilbach. 2021. "Learning in the Household." *NBER working paper*, 28844.

Cox, Natalie. 2019. "The Impact of Risk-Based Pricing in the Student Loan Market: Evidence from Borrower Repayment Decisions." *Working paper*.

Current Population Survey. 2020. "Married-Couple Families with Wives' Earnings Greater than Husbands' Earnings: 1981 to 2020." *U.S. Census Bureau Table F-22*.

de Chaisemartin, Clément, and Xavier D'Haultfœuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review*, 110(9): 2964–2996.

Dehejia, Rajeev H., and Sadek Wahba. 1999. "Causal Effects in Nonexperimental Studies: Reevaluating the Evaluation of Training Programs." *Journal of the American Statistical Association*, 94(448): 1053–1062.

Dew, Jeffrey, Sonya Britt, and Sandra Huston. 2012. "Examining the Relationship Between Financial Issues and Divorce." *Family Relations Interdisciplinary Journal of Applied Family Science*, 61(4): 615–628.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo. 2018. "The Economic Consequences of Hospital Admissions." *American Economic Review*, 108(2): 308–352.

Dubois, Pierre, and Ethan Ligon. 2011. "Incentives and Nutrition for Rotten Kids: Intrahousehold Food Allocation in the Philippines." *California University Department of Agricutural and Resource Economics Working Paper*, 1114.

Duflo, Esther. 2003. "Grandmothers and granddaughters: Old age pension and intra-household allocation in South Africa." *World Bank Economic Review*, 17(1): 1–25.

Duflo, Esther, and Christopher Udry. 2004. "Intrahousehold Resource Allocation in Côte D'Ivore: Social Norms, Separate Accounts, and Consumption Choices." *NBER Working Paper*, , (10498).

Fletschner, Diana. 2009. "Rural Women's Access to Credit: Market Imperfections in a Intrahousehold Dynamics." *World Development*, 37(3): 618–631.

Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. 2021. "Predictably Unequal? Effects of Machine Learning on Credit Markets." *Journal of Finance (Forthcoming)*.

Gathergood, John, Neale Mahoney, Neil Stewart, and Jorg Weber. 2019. "How Do Individuals Repay Their Debt? The Balance-Matching Heuristic." *American Economic Review*, 109(3): 844–875.

Goldsmith-Pinkham, Paul, and Kelly Shue. 2020. "The Gender Gap in Housing Returns." *Journal of Finance (Forthcoming)*.

Goodman-Bacon, Andrew. 2018. "Public Insurance and Mortality: Evidence from Medicaid Implementation." *Journal of Political Economy*, 126(1): 216–262.

Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics*.

Gross, David B., and Nicholas S. Souleles. 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Quarterly Journal of Economics*, 117(1): 149–185.

Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang. 2020. "The Marginal Propensity to Consume over the Business Cycle." *American Economic Journal: Macroeconomics*, 12(2): 351–384.

Hertzberg, Andrew. 2016. "Time-Consistent Individuals, Time-Inconsistent Households." *Journal of Finance (Forthcoming)*.

Hirano, Keisuke, and Guido W. Imbens. 2001. "Estimation of Causal Effects using Propensity Score Weighting: An Application to Data on Right Heart Catheterization." *Health Services and Outcomes Research Methodology*, 2: 259–278.

House Hearing: 112th Congress. 2012. "An Examinination of the Federal Reserve's Final Rule on the Card Act's "Ability to Repay" Requirement." Serial No. 112-133, Committee on Financial Services U.S. House of Representatives.

Johnson, David S., Jonathan A. Parker, and Nicholas Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review*, 96(5): 1589–1610.

Ke, Da. 2021. "Who Wears the Pants? Gender Identity Norms and Intrahousehold Financial Decision-Making." *Journal of Finance*.

Kermani, Amir, and Francis Wong. 2021. "Racial Disparities in Housin Returns." *NBER working paper*, 29306.

Knowles, John A. 2013. "Why are Married Men Working So Much? An Aggregate Analysis of Intra-Household Bargaining and Labour Supply." *Review of Economic Studies*, 80: 1055–1085.

Kocherlakota, Narayana. 1996. "Implications of Efficient Risk Sharing without Commitment." *Review of Economic Studies*, 53(4): 595–609.

Krueger, Dirk, and Fabrizio Perri. 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory." *Review of Economic Studies*, 73(1): 163–193.

Lafortune, Jeanne, and Corinne Low. 2020. "Collateralized Marriage." *Working paper*.

Lanning, Jonathan. 2021. "Testing Models of Economic Discrimination Using the Discretionary Markup of Indirect Auto Loans." *Working paper*.

Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall. 2002. "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies." *Review of Economic Studies*, 69: 209–244.

Lise, Jeremy, and Shannon Seitz. 2011. "Consumption Inequality and Intra-household Allocations." *Review of Economic Studies*, 78: 328–355.

Lundberg, Shelly, and Robert A. Pollak. 1993. "Separate Spheres Bargaining and the Marriage Market." *Journal of Political Economy*, 101(6): 988–1010.

Lundberg, Shelly, Richard Startz, and Steven Stillman. 2003. "The retirement-consumption puzzle: a marital bargaining approach." *Journal of Public Economics*, 87: 1199–1218.

Lundberg, Shelly, Robert A. Pollak, and Terence J. Wales. 1997. "Do Husbands and Wives Pool Their Resources? Evidence from the U.K. Child Benefit." *Journal of Human Resources*, 32(3): 463–480.

Malmendier, Ulrike, and Nirupama Kulkarni. 2021. "Mortgage Policies and their Effects on Racial Segregation and Upward Mobility." *Working paper*.

Mazzocco, Maurizio. 2007. "Household Intertemporal Behavior: A Collective Characterization and a Test of Commitment." *Review of Economic Studies*, 74(3): 857–895.

Mazzocco, Maurizio, Claudia Ruiz, and Shintaro Yamaguchi. 2014. "Labor Supply and Household Dynamics." *American Economic Review: Papers and Proceedings*, 104(5): 354–359.

Mian, Atif, and Amir Sufi. 2016. "Who Bears the Cost of Recessions? The Role of House Prices and Household Debt." *Handbook of Macroeconomics*, 2: 255–296.

Miller, Sarah, and Cindy K. Soo. 2020. "Does Increasing Access to Formal Credit Reduce Payday Borrowing?" *NBER working paper*, 27783.

Morse, Adair, and Karen Pence. 2020. "Technological Innovation and Discrimination in Household Finance." *NBER working paper*, , (26739).

Olafsson, Arna, and John Gathergood. 2020. "The Co-Holding Puzzle: New Evidence from Transaction-Level Data." *Centre for Economic Policy Research Discussion Paper*, 14799.

Olafsson, Arna, and Michaela Pagel. 2017. "Family Finances: Intra-Household Bargaining, Spending, and Capital Structure." *Working paper*.

Pew Research Center. 2017. "Americans see men as the financial providers, even as women's contributions grow."

Pollak, Robert A. 2011. "Allocating Time: Individuals' Technologies, Household Technology, Perfect Substitutes, and Specialization." *NBER working paper*, 17529.

Ponce, Alejandro, Enrique Seira, and Guillermo Zamarripa. 2017. "Borrowing on the Wrong Credit Card? Evidence from Mexico." *American Economic Review*, 107(4): 1335–1361.

Rosenbaum, Paul R., and Donald B. Rubin. 1983. "The central role of the propensity score in observational studies for causal effects." *Biometrika*, 70(1): 41–55.

Roth, Jonathan. 2020. "Pre-test with Caution: Event-study Estimates After Testing for Parallel Trends." *Working paper*.

Schultz, T. Paul. 1990. "Testing the Neoclassical Model of Family Labor Supply and Fertility." *Journal of Human Resources*, 25(4): 599–634.

Sun, Liyang, and Sarah Abraham. 2020. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics*.

Thomas, Duncan. 1990. "Intra-Household Resource Allocation: An Inferential Approach." *Journal of Human Resources*, 25(4): 635–664.

U.S. Bureau of Labor Statistics. 2020*a*. "Comparing characteristics and selected expenditures of dual- and single-income households with children." *Monthly Labor Review*.

U.S. Bureau of Labor Statistics. 2020*b*. "Length of pay periods in the Current Employment Statistics survey." *Current Employment Statistics*.

Vihriälä, Erkki. 2020. "Untangling the Credit Card Debt Puzzle." *Working paper*.

Vissing-Jorgensen, Annette. 2021. "Consumer Credit: Learning Your Customer's Default Risk from What (S)he Buys." *Working paper*.

Voena, Alessandra. 2015. "Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?" *American Economic Review*, 105(8): 2295–2332.

Wang, Cheng. 1995. "Dynamic Insurance with Private Information and Balanced Budgets." *Review of Economic Studies*, 62(4): 577–595.

Wolfers, Justin. 2006. "Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results." *American Economic Review*, 96(5): 1802–1820.

Figures and Tables

Figure 1-1: Average Difference in Secondary Earners' Reported Monthly Income on Credit Card Applications Between Treated and Control

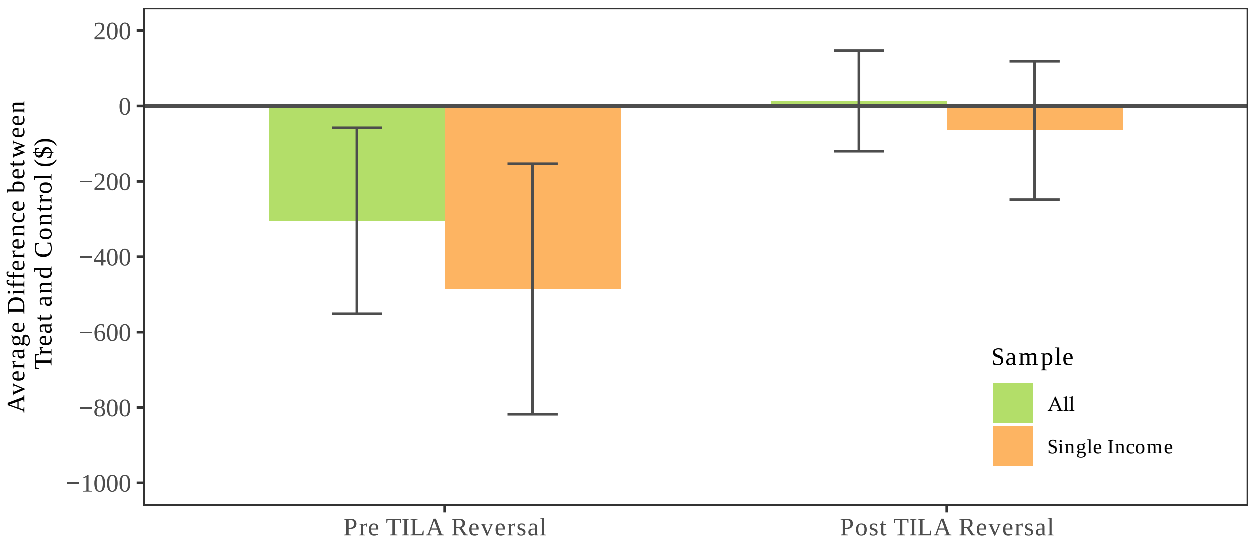

Notes: This figure reports the average difference in the monthly income reported on secondary earners' credit card applications between the treated and the control group. The difference is obtained by regressing reported monthly income on the treatment dummy. The whiskers denote 90 percent confidence intervals.

Figure 1-2: Community Property vs. Equitable Distribution States

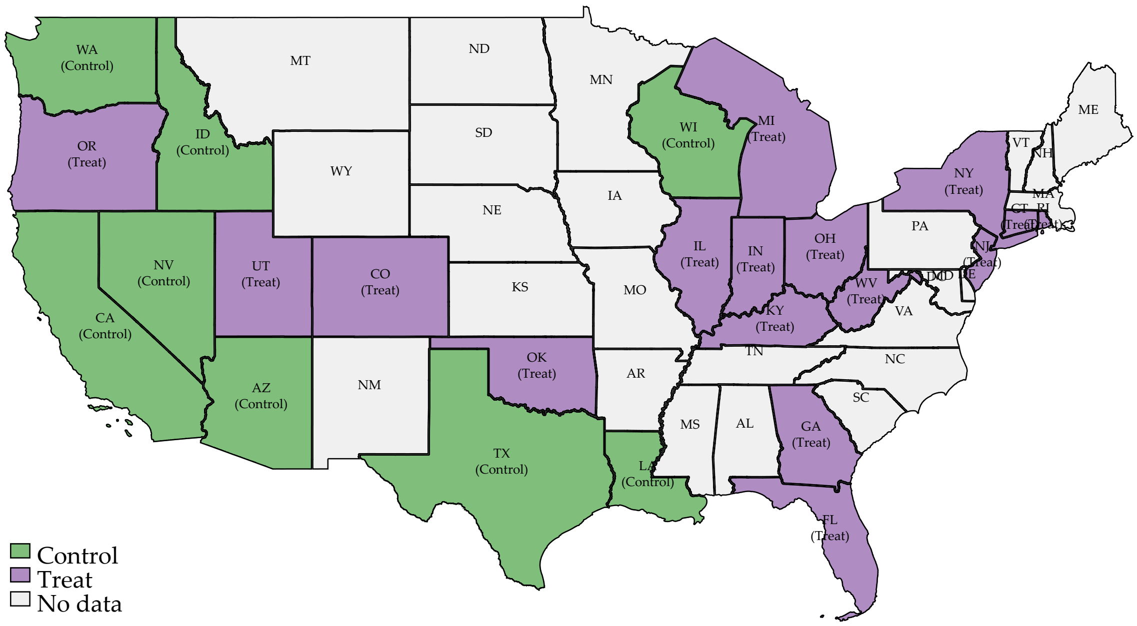

Notes: This figure shows the states in my data, color-coded by the doctrine that govern the disposition of marital property in divorce. Equitable distribution states shown in purple and community property states are shown in green. States in gray are not well represented in my data.

Figure 1-3: Consumption Shares Across the Income Distribution

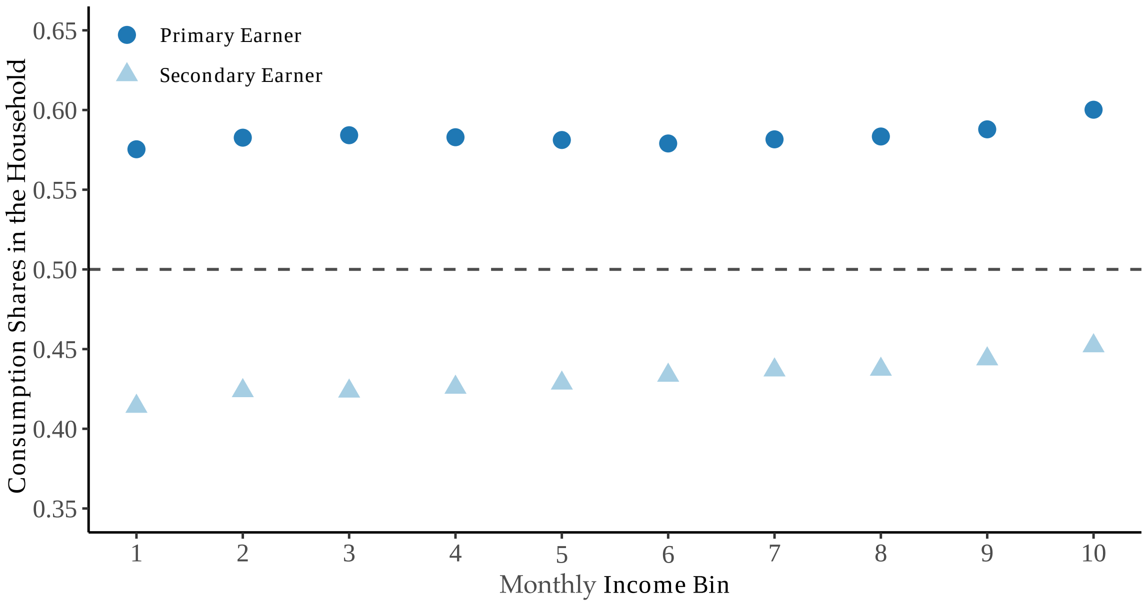

Notes: This figure compares average monthly consumption shares of individuals in the same income bin by their earner status in the household. For example, the first income bin shows that individuals in the same (lowest) income bin on average consumes 58% of total household consumption if they are primary earners but only 42% if they are secondary earners in the household.

Figure 1-4: Share of Accessible Credit in the Household

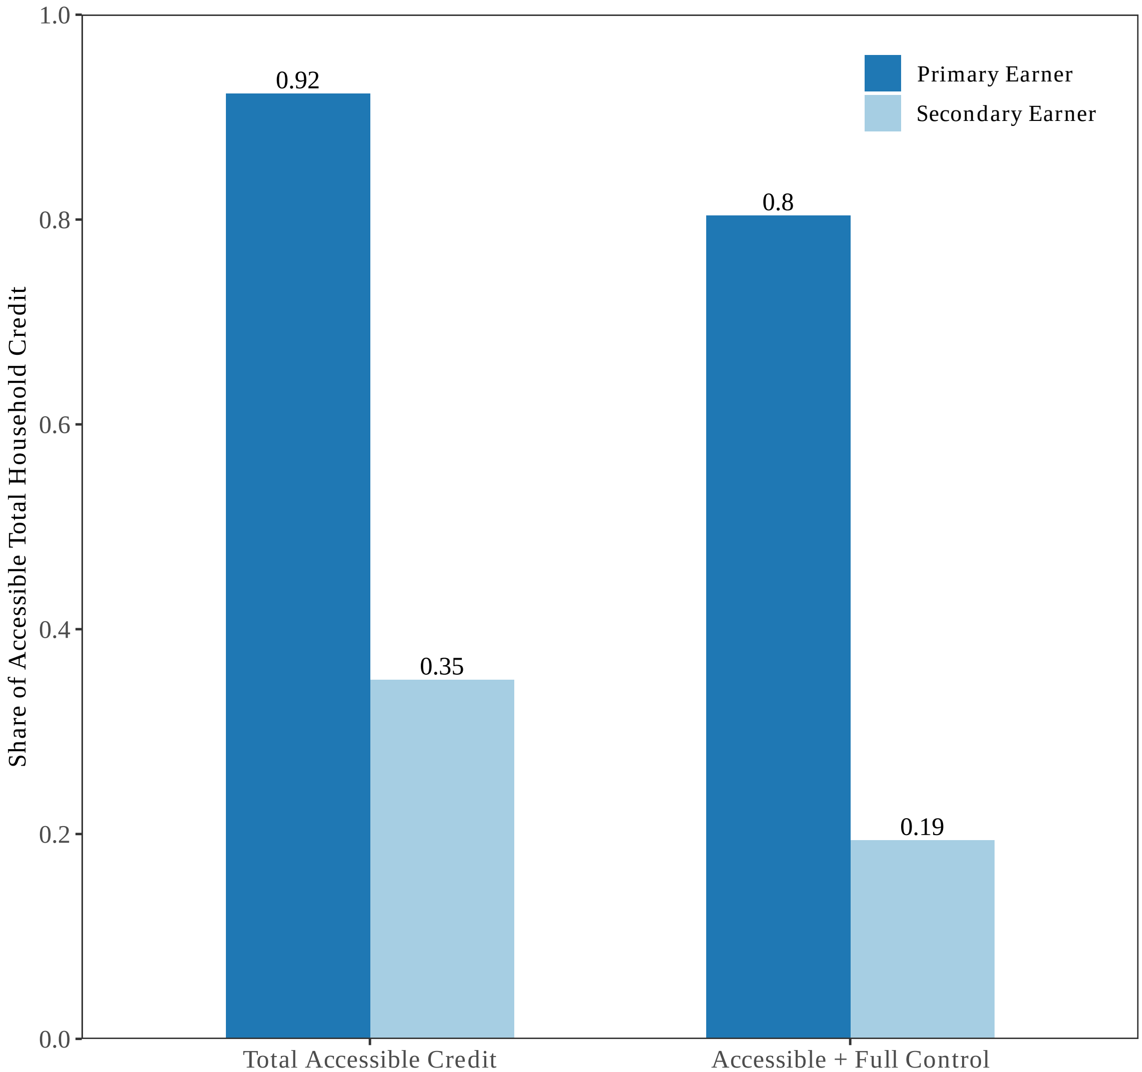

Notes: This figure shows the average monthly share of total household credit that each spouse can access during my sample period. "Total Accessible Credit" shows the average monthly credit limit that each spouse can access either as a primary account holder or as an authorized user as a share of total credit limit available at the household level. For example, the light blue bar shows that secondary earners can, on average, access 35% of total household credit limit. "Accessible + Full Control" shows the average monthly credit limit on accounts held by each spouse as a share of total household credit limit. For example, the light blue bar shows that secondary earners, on average, have full control over 19% of total household credit limit.

Figure 1-5: The Link between Within-Household Consumption and Credit Shares

(a) Secondary Earners' Within-Household Consumption Share by Credit Share Bin

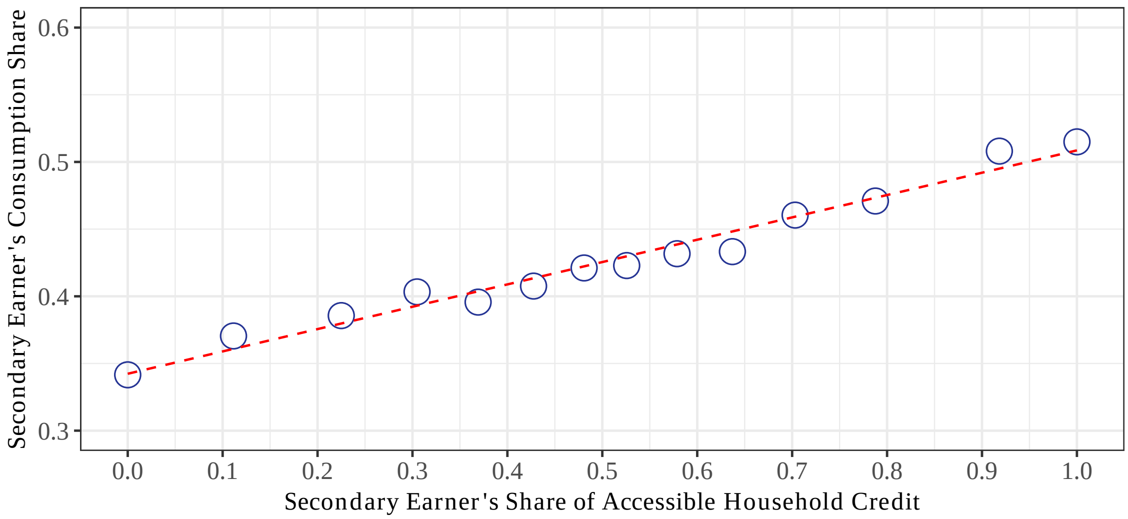

(b) Secondary Earners' Within-Household Consumption Share by
Accessible Credit as an Authorized User

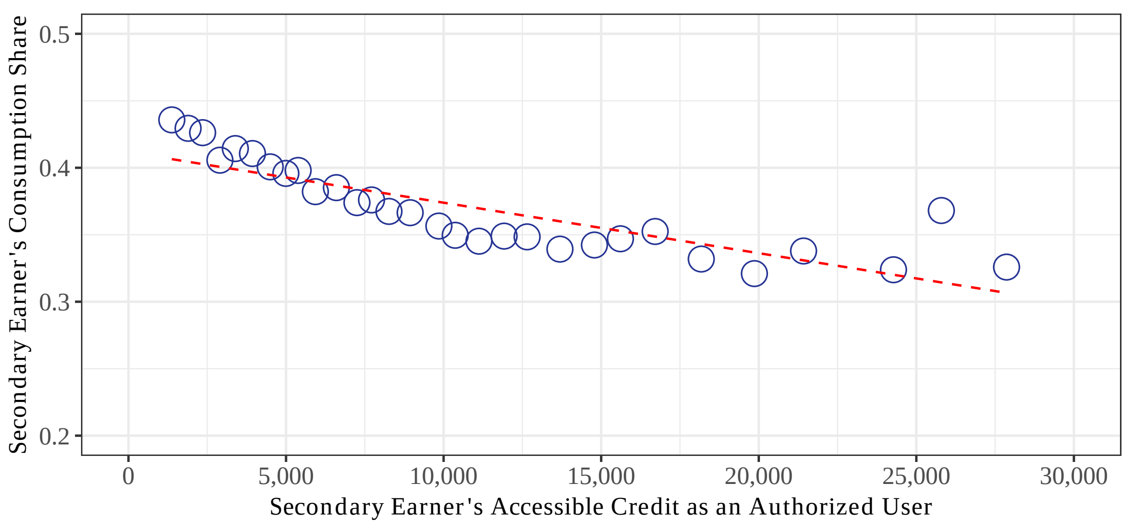

Notes: Figure a plots secondary earners' average monthly consumption share (y-axis) against the share accessible credit (x-axis) in the household. Figure b plots secondary earners' average monthly consumption share (y-axis) against the amount of average monthly credit limit they can access as an authorized user (x-axis). The red dashed line in each figure shows a linear fitted line.

Figure 1-6: The Link between Within-Household Consumption and Credit Shares by the Gender of the Secondary Earner

(a) Secondary Earners' Within-Household Consumption Share by Credit Share Bin

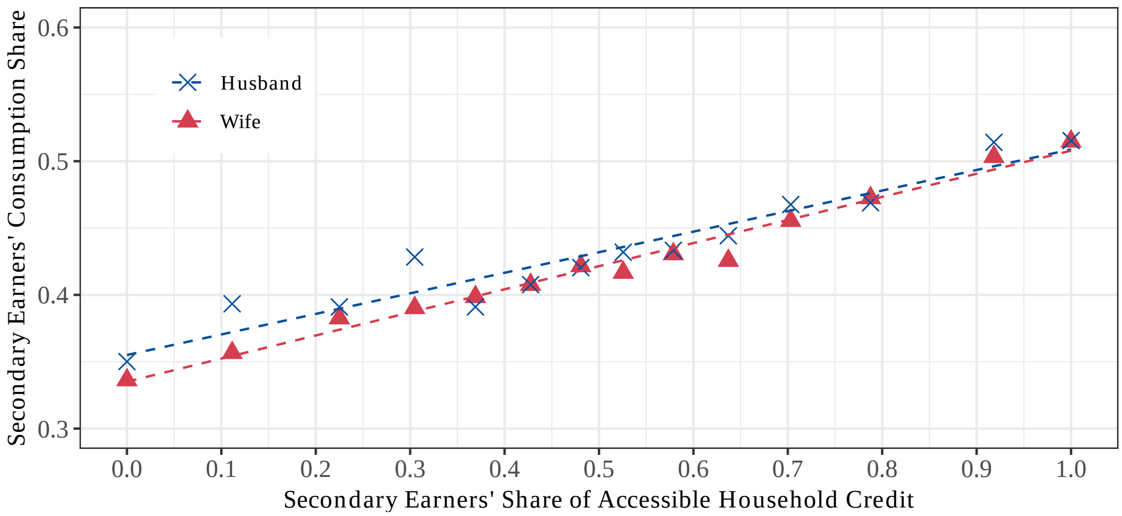

(b) Secondary Earners' Within-Household Consumption Share by Accessible Credit as an Authorized User

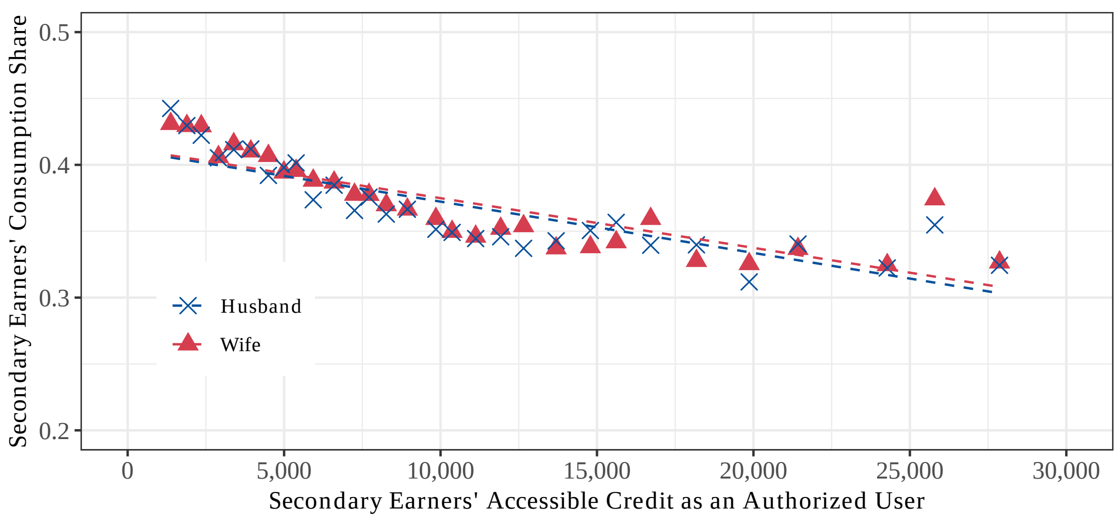

Notes: Figure a plots secondary earners' average monthly consumption share (y-axis) against the share accessible credit (x-axis) in the household by the gender of the secondary earner. Figure b plots secondary earners' average monthly consumption share (y-axis) against the amount of average monthly credit limit they can access as an authorized user (x-axis). The dashed lines show linear fitted lines.

Figure 1-7: Effect of the Reversal on Secondary Earners' Credit Limit and Consumption Share, and Consumption Gap in the Household

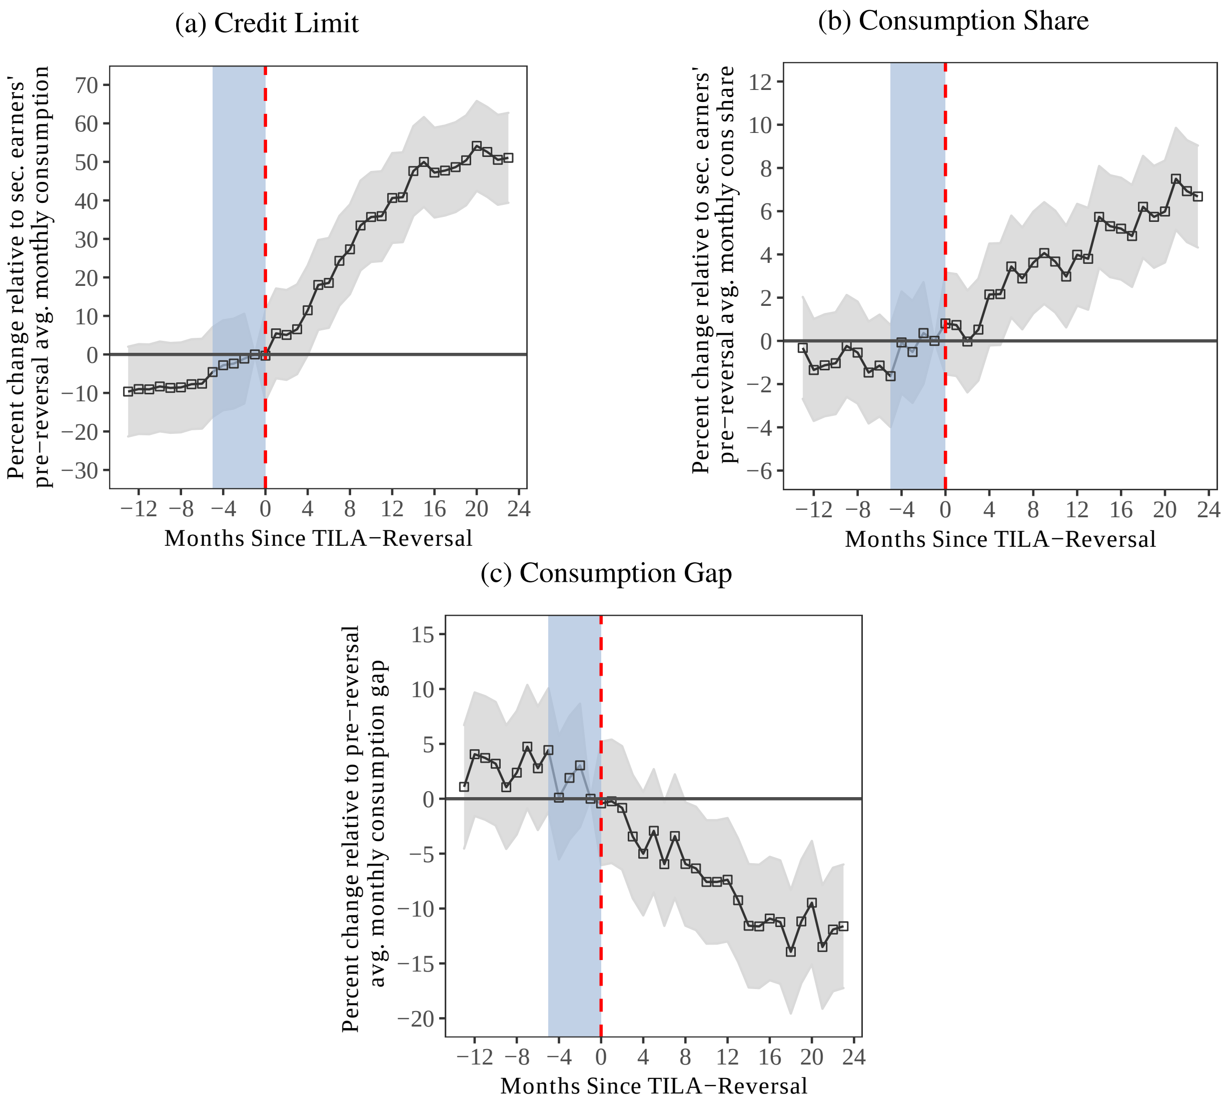

Notes: This figure plots the event-study estimates from the following specification:

$$Y_{h,t} = \alpha_h + \gamma_t + \sum_{s \neq t^{post}-1} \beta_s (Treat_h \times 1_{s=t}) + \epsilon_{h,t} \tag{1.28}$$

The outcome variables are secondary earners' sole credit card limit scaled by their pre-reversal average monthly consumption (top left); secondary earners' consumption share scaled by their pre-reversal average monthly consumption share (top right); and the consumption gap in the household scaled by pre-reversal average monthly consumption gap (bottom center). The consumption gap in the household is defined as the difference between the consumption shares of each spouse (i.e., consumption share of primary earner minus that of secondary earner). The month prior to the reversal is omitted, so β_s can be interpreted relative to this pre-reversal baseline period. Red dashed lines denote the month of the reversal. 90 percent confidence intervals are shown in gray. The shaded blue area denotes the phase-in period when the CFPB first announced the reversal and allowed credit card issuers to start adopting the new income collection standard.

Figure 1-8: Decomposition of the Change in Household Consumption

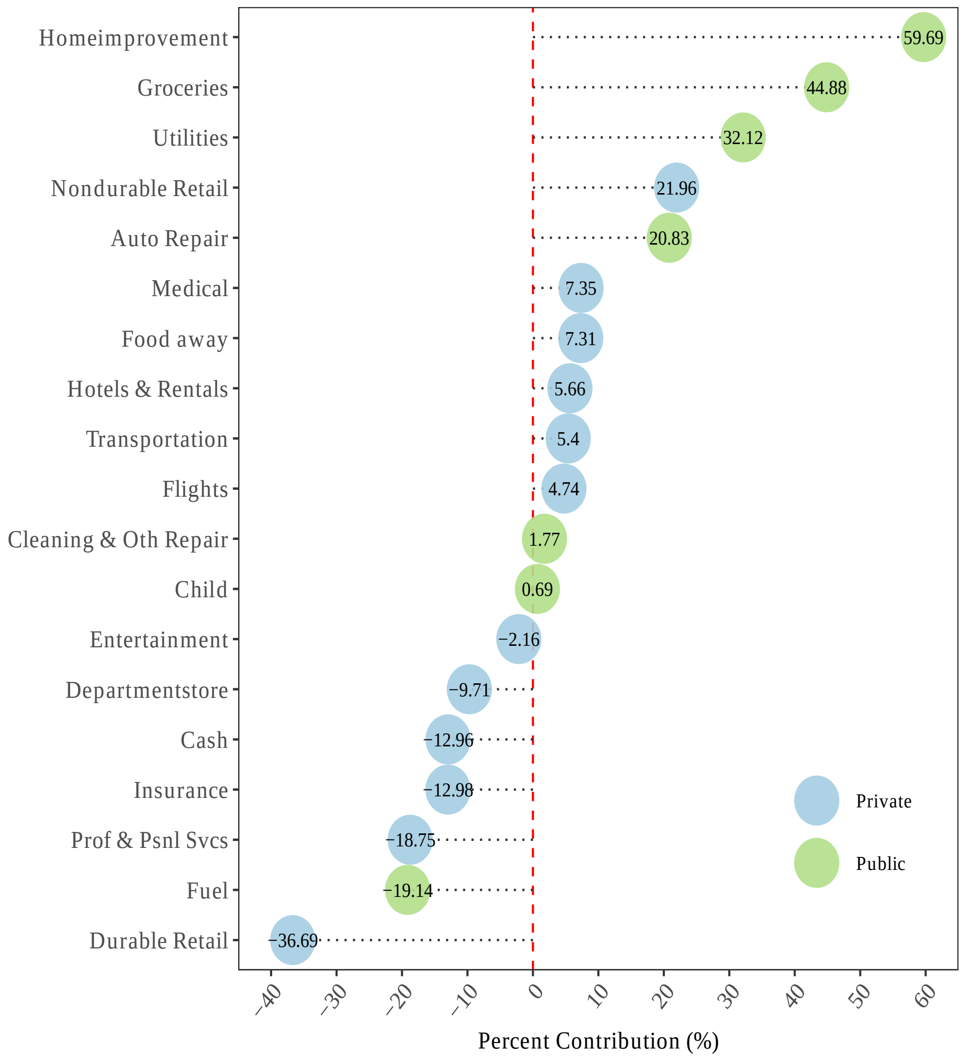

Notes: This figure decomposes the change in household consumption into detailed spending categories. The number shown in each bubble denotes how much each category contributes to the overall consumption effect, and the numbers sum to 100. For example, spending on home improvement explains 60% of the total increase in household consumption.

Figure 1-9: Changes in Secondary Earners' Bargaining Power:
A Sufficient Statistics Approach

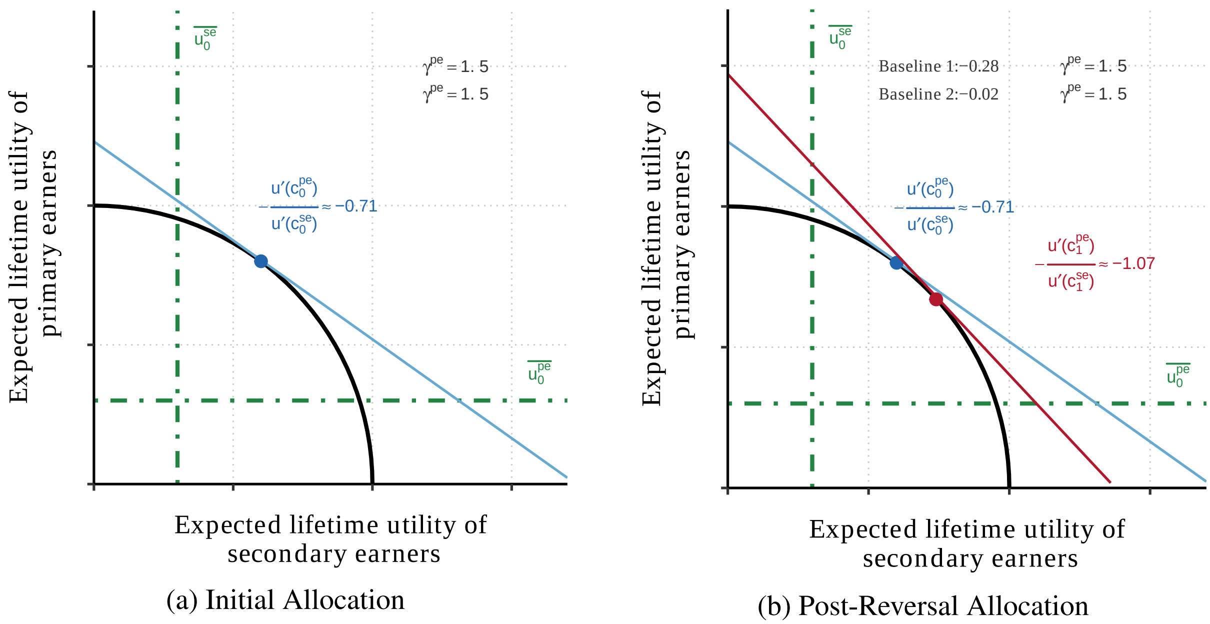

Notes: This figure illustrates the change in spouses' marital bargaining power in using a sufficient statistics approach. Equation 1.18 shows the model's key prediction that the ratio of spouses' marginal utilities of consumption has a one-to-one mapping to the ratio of spouses' bargaining power (i.e., the slope of the Pareto frontier). Using this equation and reduced form statistics on secondary and primary earners' average monthly consumption for the treated group, I quantify the change in secondary earners' relative bargaining power. I assume a risk-aversion parameter of 1.5 for both spouses. Panel a shows the location of couples' consumption sharing plan before the reversal (blue dot) and Panel b shows how this consumption sharing plan changed after the reversal (red dot). In each figure, the y-axis plots the primary earner's expected utility and the x-axis plot the secondary earner's expected utility. Vertical (horizontal) dot-dash line shows secondary (primary) earners' outside options, or their expected lifetime utility in case of divorce. Curved black lines show the Pareto frontier and the tangency points on the curve indicate the location of efficient intra-household allocation of resources. Comparing the two panels show that secondary earners' relative bargaining power increased by 36 percentage points after the reversal, from 0.71 to 1.07. I benchmark this increase to two baseline numbers annotated in Panel b. "Baseline 1" shows the average monthly change in secondary earners' relative bargaining power in the pre-reversal period among the card holder sample where secondary earners eventually open a credit card account. "Baseline 2" shows the same statistics among the all sample. "Baseline 2" shows that the typical monthly variation in secondary earners' bargaining power is only 2 percentage points, but can be as large as 28 percentage points among card openers. Then the change can be as high as 36 percentage points after the reversal. This figure builds on Chiappori and Mazzocco (2017).

Figure 1-10: Quantitative Importance of the Limited-Commitment Channel

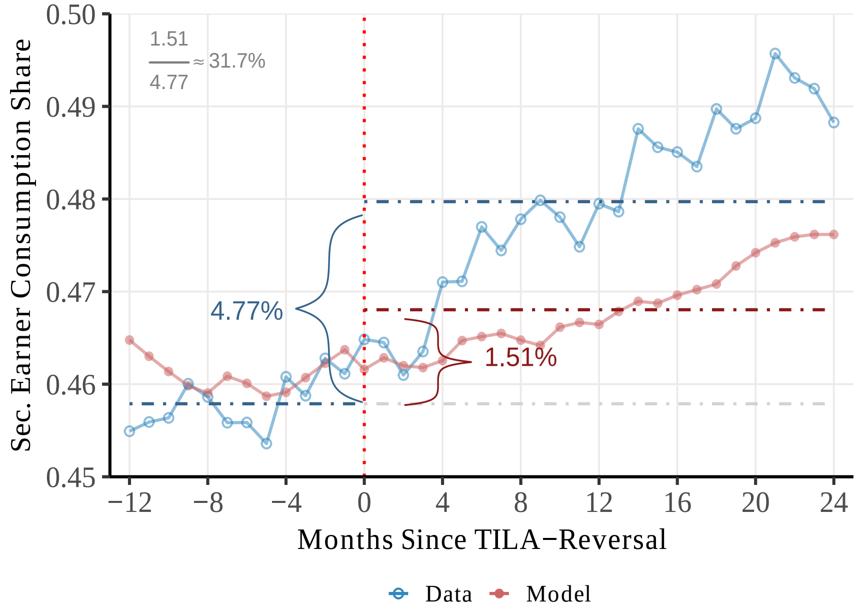

Notes: This figure compares secondary earners' average monthly consumption shares observed in the data and the model. The red (blue) line shows secondary earners' average monthly consumption share observed in the model described in Section 3.3 (data). The blue line is obtained by applying the dynamic DiD estimates shown in Figure 1-7 to the model-generated pre-reversal consumption share mean. The dot-dash lines around the 0 x-intercept shows the pre- and post-reversal mean of each line. The annotation at the top left corner shows that the model explains roughly 32% of the observed increase in secondary earners' consumption shares in the data.

Table 1.1: Covariate Balancing between Control and Treated Households (Pre-Reversal Characteristics)

	Raw			Matched		
	Control Mean (1)	Treated Mean (2)	Treat - Control (3)	Control Mean (4)	Treated Mean (5)	Treat - Control (6)
Age gap	1.58	1.52	-0.06	1.58	1.61	0.03
Wife's age	40.28	39.72	-0.56	40.27	40.29	0.02
Husband's age	41.86	41.24	-0.62	41.85	41.90	0.05
Wife is Secondary Earner	0.61	0.61	-0.01	0.61	0.61	0.00
Debt-to-Income	0.19	0.20	0.01	0.19	0.19	0.00
Household Income ($)	9,037	9,417	381	9,013	9,010	-3
Cash on hand ($)	7,023	7,559	536	6,985	6,938	-48
Has a credit card	0.74	0.79	0.05	0.74	0.74	0.00
Total credit limit ($)	9,217	10,697	1,479	9,115	9,276	161
Total card balance ($)	2,572	2,833	260	2,566	2,593	27
Revolving balance ($)	2,152	2,321	170	2,085	2,113	28
Number of Households	33,140	47,995	14,855	33,098	33,098	0

Notes: This table reports average pre-reversal characteristics for treated and control households before and after the propensity score matching (PSM) procedure described in Section 1.3.2. The PSM method strengthens the "parallel trends" assumption by balancing pre-treatment characteristics that could influence the dynamics of outcome variables. I calculate the conditional probability of being treated (i.e., the propensity score) using covariates listed in this table. I chose covariates that may be associated with the dynamics of credit card opening. Treated households are those that reside in equitable distribution (ED) states and control households are those that reside in community property (CP) states. The first three columns report average characteristics for the pre-match sample and the last three columns report those for the matched sample. Columns 1 and 4 report the control group mean, Columns 2 and 5 report the treated group mean, and Columns 3 and 6 report the differences in means. All variables are reported in monthly frequency. Age related variables are reported in years. Debt-to-Income reports total monthly debt payments (e.g., auto, credit card, mortgage, student, and other) to household income. Household income is the sum of labor income (payroll direct deposits), government transfers, business, and gig income. Cash on hand reports the sum of month-end checking account balances of spouses' checking accounts. Has a credit card is an indicator for at least one member in a household having a credit card account at JPMC. Total credit limit reports the sum of all credit card limits available at the household-level (joint credit card limits are counted only once). Total card balance and revolving balance respectively refer to the end-of-billing-cycle credit card balance and interest-accruing revolving balance. The matched sample of 66,196 households is my main analysis sample.

Table 1.2: Pre-TILA reversal Descriptive Statistics
(Matched Sample)

	A. Household-level Characteristics				
	Mean	SD	p25	p50	p75
	(1)	(2)	(3)	(4)	(5)
Age gap	1.59	3.93	0.00	1.00	4.00
Wife's age	40.28	11.09	31.00	38.00	49.00
Husband's age	41.87	11.13	32.00	40.00	51.00
Consumption ($)	6,159	5,267	2,780	4,849	7,859
Total Income ($)	9,011	14,172	4,616	6,752	9,974
Cash on hand ($)	6,962	30,280	1,182	2,771	6,271
Has credit access	0.74	0.44	0.00	1.00	1.00
	B. Within-Household Characteristics				
	Secondary Earner		Primary Earner		Mean
	Mean	SD	Mean	SD	Difference
Female	0.61	–	0.39	–	-0.23
Age	40.8	11.1	41.3	11.1	0.5
Total Income ($)	1,004	4,689	8,007	13,363	7,002
Cash on hand ($)	977	6,022	5,985	29,420	5,008
Consumption share	0.41	0.21	0.59	0.21	0.17
Consumption ($)	2,464	2,485	3,695	4,923	1,231
Public consumption ($)	866	1,807	1,249	1,983	383
Private consumption ($)	1,598	939	2,446	4,002	848
Has a sole credit card	0.02	0.15	0.55	0.50	0.53
Credit limit ($)	155	1,442	6,055	9,053	5,900
Card balance ($)	35	405	1,778	3,875	1,743
Revolving balance ($)	33	397	1,525	3,815	1,493
Number of Households	66,196	66,196	66,196	66,196	66,196

Notes: This table reports summary statistics for my main analysis sample. Panel A reports household-level and Panel B reports within-household characteristics. To preserve anonymity, "percentiles" are presented as means of ten observations in the p^{th} percentiles. All variables are reported in monthly frequency. Age is in years. Household-level consumption is defined as the sum of spending on all financial accounts (debit, credit, and checking) associated with a household, and spouse-level consumption is defined as the sum of spending on financial accounts of individual spouses in the household. Consumption shares of each spouse is each spouse's spending as a share of total household spending. Public consumption denotes spending on goods that are jointly consumed by the household and private consumption denotes spending on goods that are consumed individually. Total income is defined as the sum of labor income, government transfers, and other income. Cash on hand refers to the end-of-month checking account balance and "has credit access" is an indicator for whether a household has at least one credit card account. Spouse-level credit limit measures credit card limits on each spouse's sole credit card account, and credit card balance refers to the end of billing cycle credit card balance. Revolving balance denotes interest-accruing revolving credit card balance. Detailed spending categories are reported in Table A.4.

Table 1.3: Effect of the TILA Reversal on Secondary Earner Credit Limits

Secondary Earner Outcomes	Card Holders						All Sample					
	Baseline		Single Income		Sec. Earner Older		Baseline		Single Income		Sec. Earner Older	
	(1)		(2)		(3)		(4)		(5)		(6)	
	A. Difference-in-Differences Estimates											
Credit Limit	40.45	***	56.92	***	50.93	***	7.73	***	11.01	***	10.26	***
	(1.69)		(2.6)		(2.73)		(0.35)		(0.55)		(0.58)	
Number of Observations	455,157		211,916		171,117		2,577,970		1,198,437		944,585	
	B. Pre-Reversal Mean											
Credit Limit ($)	877.0		825.8		850.9		154.8		146.0		154.1	
Conditional Credit Limit ($)	6,826		6,537		7,001		6,826		6,537		7,001	
Consumption ($)	2,532		2,282		2,581		2,464		2,236		2,459	
	C. Economic Significance											
Dollar Effects ($)	1,024		1,299		1,314		191		246		252	

Notes: Panel A reports estimates from the following difference-in-differences regression:

$$Y_{h,t} = \alpha_h + \gamma_t + \beta \mathbf{1}[Treat \times Post]_{h,t} + \epsilon_{h,t}$$

where the dependent variable is secondary earners' sole credit card limits, scaled by their average pre-reversal monthly consumption. α_h and γ_t denote household and time (month-year) fixed effects, and $\mathbf{1}[Treat \times Post]_{h,t}$ is an interaction term between a treatment indicator (i.e., h in equitable distribution states) and a post-reversal indicator (i.e., $t \geq$ November 2013). Reported coefficients are multiplied by 100 for readability. β can be interpreted as a percent change in secondary earner credit limits relative to their pre-reversal monthly consumption. Standard errors are clustered at the state-level and reported in parentheses. The first three columns restrict the sample to households where secondary earners eventually opened sole credit card accounts during my sample period (i.e., The "Card Holders" sample), and the last three columns use all sample, including households where secondary earners never opened credit card accounts (i.e., The "All Sample"). Within each sample, Columns 2 and 5 further restrict the sample to single-income households where primary earners are the only breadwinners, and Columns 3 and 6 restrict the sample to those where secondary earners are older than primary earners. Panel B reports secondary earners' pre-reversal average monthly credit card limits (unconditional and conditional) and consumption. Panel C reports estimated effects in dollars, computed as $\beta\times$ pre-reversal average monthly consumption reported in Panel B. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 1.4: Effect of the TILA Reversal on Secondary Earners' Consumption and the Consumption Gap in the Household

	Card Holders			All Sample		
Secondary Earner Outcomes	Baseline (1)	Single Income (2)	Sec. Earner Older (3)	Baseline (4)	Single Income (5)	Sec. Earner Older (6)
	A. Difference-in-Differences Estimates					
Consumption	13.56 ***	24.77 ***	15.63 ***	2.97 ***	7.67 ***	3.6 ***
	(1.04)	(1.65)	(1.7)	(0.33)	(0.52)	(0.55)
Consumption Share	4.97 ***	8.09 ***	7.13 ***	1.49 ***	2.75 ***	1.93 ***
	(0.34)	(0.54)	(0.55)	(0.12)	(0.19)	(0.2)
Consumption Gap	-10.22 ***	-14.26 ***	-12.67 ***	-4.41 ***	-5.71 ***	-5.66 ***
	(0.80)	(1.22)	(1.31)	(0.31)	(0.46)	(0.51)
Number of Observations	455,157	211,916	171,117	2,577,970	1,198,437	944,585
	B. Pre-Reversal Mean					
Consumption ($)	2,532	2,282	2,581	2,464	2,236	2,459
Consumption Share (%)	45.1	44.3	45.4	41.2	40.6	41.5
Consumption Gap (%)	9.62	11.37	9.10	17.47	18.76	16.91
	C. Economic Significance					
Dollar Effects ($)	343	565	403	73	172	89
MPC out of Credit Increases	0.22	0.30	0.28	0.46	0.54	0.37

Notes: This table reports estimates from the following difference-in-differences regression:

$$Y_{h,t} = \alpha_h + \gamma_t + \beta \mathbf{1}[Treat \times Post]_{h,t} + \epsilon_{h,t}$$

where the dependent variables are secondary earners' consumption and consumption shares scaled, respectively, by their average pre-reversal monthly consumption and consumption shares. α_h and γ_t denote household and time (month-year) fixed effects, and $\mathbf{1}[Treat \times Post]_{h,t}$ is an interaction term between a treatment indicator (i.e., h in equitable distribution states) and a post-reversal indicator (i.e., $t \geq$ November 2013). Reported coefficients are multiplied by 100 for readability. β can be interpreted as a percent change in secondary earners' consumption (share) relative to their pre-reversal monthly consumption (share). Standard errors are clustered at the state-level and reported in parentheses. The first three columns restrict the sample to households where secondary earners eventually opened sole credit card accounts during my sample period (i.e., The "Card Holders" sample), and the last three columns use all sample, including households where secondary earners never opened credit card accounts (i.e., The "All Sample"). Within each sample, Columns 2 and 5 further restrict the sample to single-income households where primary earners are breadwinners, and Columns 3 and 6 restrict the sample to those where secondary earners are older than primary earners. Panel B reports secondary earners' pre-reversal average monthly consumption and consumption shares. Panel C reports estimated consumption effects in dollars. MPC out of credit limit increases is computed as the estimated effects on credit card balance divided by the effects on credit limits. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 1.5: Effect of the TILA Reversal on Household Outcomes

Household Outcomes	Card Holders			All Sample		
	Baseline	Single Income	Sec. Earner Older	Baseline	Single Income	Sec. Earner Older
	(1)	(2)	(3)	(4)	(5)	(6)
	A. Difference-in-Differences Estimates					
Credit Limit	20.48***	29.19***	25.68***	6.97***	9.418***	7.77***
	(1.49)	(2.33)	(2.41)	(0.47)	(0.73)	(0.76)
Consumption	3.00***	5.15***	1.83**	0.75***	2.17***	0.90***
	(0.55)	(0.84)	(0.89)	(0.21)	(0.31)	(0.34)
Revolving Debt	0.90**	1.37**	0.72	0.67***	0.09	0.68***
	(0.37)	(0.58)	(0.61)	(0.14)	(0.22)	(0.24)
Number of Observations	455,118	211,916	171,078	2,577,591	1,198,242	944,429
	B. Pre-Reversal Mean					
Credit Limits	4,827	4,684	4,723	9,195	8,957	8,712
Consumption	5,658	5,178	5,702	6,159	5,653	6,055
Revolving Debt	1,148	1,104	1,142	2,099	2,035	2,075
	C. Economic Significance					
Dollar Effects: Credit Limits	1,159	1,512	1,465	429	532	471
Dollar Effects: Consumption	170	267	104	46	123	55
Dollar Effects: Revolving Debt	51	71	41	41	5	41
MPC out of Credit Increases	0.08	0.08	0.06	0.17	0.08	0.17

Notes: This table reports estimates from the following difference-in-differences regression:

$$Y_{h,t} = \alpha_h + \gamma_t + \beta \mathbf{1}[Treat \times Post]_{h,t} + \epsilon_{h,t}$$

where the dependent variables include total household credit limits, consumption, and revolving debt, measured as the sum of the two spouses' credit card limits (joint accounts are aggregated only once), expenditures, and interest-bearing credit card balances. All outcomes are scaled by households' average pre-reversal monthly consumption. α_h and γ_t denote household and time (month-year) fixed effects, and $\mathbf{1}[Treat \times Post]_{h,t}$ is an interaction term between a treatment indicator (i.e., h in equitable distribution states) and a post-reversal indicator (i.e., $t \geq$ November 2013). Reported coefficients are multiplied by 100 for readability. β can be interpreted as a percent change in household outcome relative to pre-reversal consumption. Standard errors are clustered at the state-level and reported in parentheses. The first three columns restrict the sample to households where secondary earners eventually opened sole credit card accounts during my sample period (i.e., The "Card Holders" sample), and the last three columns use all sample, including households where secondary earners never opened credit card accounts (i.e., The "All Sample"). Within each sample, Columns 2 and 5 further restrict the sample to single-income households where primary earners are breadwinners, and Columns 3 and 6 restrict the sample to those where secondary earners are older than primary earners. Panel B reports pre-reversal average of the outcome variables. Panel C reports estimated coefficients in dollar terms, computed as $\beta\times$ pre-reversal average consumption. MPC out of credit limit increases is computed as the estimated effects on household consumption divided by effects on household credit limits. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 1.6: Effect of the TILA Reversal on Household Financial Outcomes

Household Outcomes	Card Holders: Baseline (1)	Card Holders: Single Income (2)	Card Holders: Sec. Earner Older (3)	All Sample: Baseline (4)	All Sample: Single Income (5)	All Sample: Sec. Earner Older (6)
	A. Difference-in-Differences Estimates					
Delinquency	0.06	-0.06	0.03	-0.01	-0.06 **	0.03
	(0.04)	(0.06)	(0.07)	(0.02)	(0.03)	(0.03)
Overdraft	0.09	0.04	0.05	0.01	-0.01	-0.02
	(0.06)	(0.07)	(0.08)	(0.02)	(0.03)	(0.03)
High-interest loans	-0.02	0.10	0.10	0.03	0.05	0.02
	(0.05)	(0.07)	(0.09)	(0.02)	(0.03)	(0.03)
Debt Prioritization	1.16 **	1.13	0.75	0.54	0.82 *	0.99 *
	(0.56)	(0.81)	(0.93)	(0.30)	(0.45)	(0.50)
Number of Observations	455,157	211,916	171,117	2,577,970	1,198,437	944,585
	B. Pre-Reversal Mean					
Delinquency (%)	0.3	0.3	0.3	0.6	0.6	0.7
Overdraft (%)	0.7	0.8	0.8	0.5	0.5	0.5
High-interest loans (%)	1.0	1.1	1.2	0.8	0.7	0.9
Debt Prioritization (%)	80.6	80.4	80.0	81.7	81.2	80.7

Notes: This table reports estimates from the following difference-in-differences regression:

$$Y_{h,t} = \alpha_h + \gamma_t + \beta \mathbf{1}[Treat \times Post]_{h,t} + \epsilon_{h,t}$$

where the dependent variables include indicators for whether (i) at least one credit card account in the household is falling behind on making required monthly payments for at least 30 days (i.e., delinquent); (ii) at least one checking account in the household incurred overdraft fees; (iii) households make any payments to payday or subprime personal loan lenders; and (iv) whether households optimally pay debt in a way that they pay down more expensive debt first while borrowing more on lower-interest cards. Debt prioritization analysis is limited to households with at least two credit card accounts. Roughly 55% (20%) of households sampled in the first (last) three columns have multiple credit cards. α_h and γ_t denote household and time (month-year) fixed effects, and $\mathbf{1}[Treat \times Post]_{h,t}$ is an interaction term between a treatment indicator (i.e., h in equitable distribution states) and a post-reversal indicator (i.e., $t \geq$ November 2013). Reported coefficients are multiplied by 100 for readability. β can be interpreted as a percentage point change. Standard errors are clustered at the state-level and reported in parentheses. The first three columns restrict the sample to households where secondary earners eventually opened sole credit card accounts during my sample period (i.e., The "Card Holders" sample), and the last three columns use all sample, including households where secondary earners never opened credit card accounts (i.e., The "All Sample"). Within each sample, Columns 2 and 5 further restrict the sample to single-income households where primary earners are breadwinners, and Columns 3 and 6 restrict the sample to those where secondary earners are older than primary earners. Panel B reports pre-reversal average of the outcome variables. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 1.7: Private vs. Public Consumption

	Card Holders						All Sample					
	Baseline		Single Income		Sec. Earner Older		Baseline		Single Income		Sec. Earner Older	
	(1)		(2)		(3)		(4)		(5)		(6)	
					A. Secondary Earner Outcomes							
Private Consumption	8.19	***	16.04	***	9.20	***	1.00	***	4.07	***	1.92	***
	(0.84)		(1.33)		(1.35)		(0.27)		(0.42)		(0.45)	
Public Consumption	5.37	***	8.73	***	6.43	***	1.96	***	3.60	***	1.69	***
	(0.42)		(0.67)		(0.73)		(0.13)		(0.21)		(0.23)	
					B. Household Outcomes							
Private Consumption	1.24	***	2.89	***	0.64		-0.09		1.11	***	0.27	
	(0.45)		(0.68)		(0.72)		(0.17)		(0.25)		(0.27)	
Public Consumption	1.77	***	2.27	***	1.19	***	0.84	***	1.07	***	0.64	***
	(0.24)		(0.36)		(0.39)		(0.09)		(0.14)		(0.15)	
Number of Observations	455,157		211,916		171,117		2,577,970		1,198,437		944,585	
					C. Pre-reversal Mean							
Sec. Earner: Private	1,678		1,516		1,701		1,598		1,449		1,586	
Sec. Earner: Public	854		766		880		866		786		872	
Household: Private	3,776		3,461		3,798		4,044		3,707		3,966	
Household: Public	1,882		1,718		1,905		2,115		1,945		2,089	
					D. Economic Significance							
Sec. Earner: Private	207		366		237		25		91		47	
Sec. Earner: Public	136		199		166		48		81		42	
Household: Private	70		150		36		-5		63		16	
Household: Public	100		117		68		52		60		39	

Notes: Panel A and B report difference-in-differences estimates. The dependent variables in Panel A (B) are secondary earners' (households') private and public consumption scaled by their pre-reversal average monthly consumption. β can be interpreted as a percent change in secondary earner (household) private or public consumption relative to their pre-reversal monthly consumption. Private consumption refers to spending on goods and services that are consumed privately, such as clothing. Public consumption refers to spending on goods and services that are consumed jointly by the household, such as childcare. The categorization of "private" or "public" consumption follows existing studies (Chiappori, Fortin and Lacroix, 2002; Mazzocco, 2007). Table A.4 reports categorization details. All specifications include household and time (month-year) fixed effects. Standard errors are clustered at the state-level and reported in parentheses. The first three columns restrict the sample to households where secondary earners eventually opened sole credit card accounts during my sample period (i.e., The "Card Holders" sample), and the last three columns use all sample, including households where secondary earners never opened credit card accounts (i.e., The "All Sample"). Within each sample, Columns 2 and 5 further restrict the sample to single-income households where primary earners are breadwinners, and Columns 3 and 6 restrict the sample to those where secondary earners are older than primary earners. Panel C reports pre-reversal average of the outcome variables. Panel D reports estimated coefficients in dollar terms for each outcome, computed as $\beta\times$ pre-reversal average (secondary earner or household) consumption. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 1.8: Economic Mechanism

	Divorce Rates				Troubled Marriage				Has Children				Financially Constrained			
Secondary Earner Outcomes	Low (1)		High (2)		Yes (3)		No (4)		Yes (5)		No (6)		Yes (7)		No (8)	
							A. Difference-in-Differences Estimates									
Credit Limit	48.79	***	48.79	***	52.93	***	40.10	***	40.92	***	41.85	***	10.58	*	42.23	***
	(4.82)		(3.28)		(4.46)		(1.8)		(2.31)		(2.25)		(5.66)		(1.75)	
Consumption Share	3.32	***	9.40	***	6.77	***	4.86	***	2.65	***	6.28	***	-0.40		5.26	***
	(0.98)		(0.68)		(1.07)		(0.36)		(0.5)		(0.45)		(1.43)		(0.35)	
Number of Observations	163,589		117,925		36,601		418,556		154,050		301,107		21,850		433,307	
							B. Pre-Reversal Mean									
Sec. Earner Credit Limit ($)	942		841		749.0		888.2		957.0		836.0		267.9		907.7	
Sec. Earner Consumption Share (%)	44.1		45.6		44.9		45.2		46.1		44.6		43.2		45.2	

Notes: This table reports difference-in-differences estimates using various subsamples among the "Card Holder Sample," or households where secondary earners eventually opened sole credit card accounts during my sample period. The dependent variables include secondary earners' sole credit card limits and consumption shares. The outcomes are scaled the same way as in Tables 1.3 and 1.4. Thus, β can be interpreted as a percent change in secondary earners' credit limits relative to their pre-reversal monthly consumption (first row) or their consumption share relative to pre-reversal monthly consumption share (second row). All specifications include household and time (month-year) fixed effects. Standard errors are clustered at the state-level and reported in parentheses. States with high versus low divorce rates are those with above the top tercile (4.31) or below the bottom tercile (3.12) of the annual state-level divorce rates per capita between 1990 and 2012. The bottom (top) tercile represents 45.5 (55.4) percent of divorce rates per marriage rates in each state. I obtain state-level divorce and marriage rates data from the Centers for Disease Control and Prevention National Center for Health Statistics. I classify couples with above median ($0) pre-reversal monthly spending on counseling (incl. couple counseling) or dating services as being in troubled marriages. I classify couples with above median ($0) pre-reversal monthly spending on children's clothing or childcare as those with children. I classify couples with average monthly pre-reversal credit card utilization rates above $p90^{th}$ (0.83) as being financially constrained. Panel B reports average pre-reversal monthly outcome variables. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 1.9: Parameters of the Model

Parameter	**Value**	**Reference**
Relative risk aversion (γ)	1.5	Attanasio et al (2008)
Discount factor (β)	0.97	Attanasio et al (2008)
Rate of return on assets ($\underline{\mathrm{r}}$)	0.0016	Bayot and Voena (2014)
Cost of borrowing ($\overline{r}$)	0.013	Data
Economies of scale for children ($e(k)$)	1.2	Voena (2015)
Economies of scale in couple (ρ)	1.4	Voena (2015)
Disutility from labor market participation (ψ)	0.05	match BLS LFP rate
Standard deviation of preference shocks (σ_ξ)	0.05	match CDC divorce rate
Gains from experience (λ_0, λ_1)	$0.03, -0.0003$	Attanasio et al (2008)
Depreciation rate (δ)	0.08	Voena (2015)
Standard deviation of PE's permanent shock (σ_{ζ^P})	0.14	Attanasio et al (2008)
Standard deviation of SE's permanent shock (σ_{ζ^S})	0.14	Attanasio et al (2008)
Wage covariance of PE and SE ($\sigma_{\zeta^P \zeta^S}$)	0.19	Attanasio et al (2008)
Primary earners' credit limit (L^P)	$10,410$	Data
Secondary earners' credit limit ($\underline{\mathrm{L}}^S$, $\overline{L}^S$)	$(6,826, 8,225)$	Data

Notes: This table reports parameters used in the dynamic model presented in Section 3.3.

Table 1.10: Welfare Gains

	Primary Earner (1)	Secondary Earner (2)	Household (3)
Consumption Equivalent	1.12	4.39	2.69

Notes: This table reports the welfare gain from the 2013 reversal. Section 1.7.3 details how I compute the consumption equivalent variation.

Chapter 2

The Economic Impact of Education Spending: Evidence from Self-Employed Households

2.1 Introduction

Paying for college has become a growing financial burden for American households, amid rising tuition costs and student debt. There is an active policy debate on easing this financial burden by making college more affordable and forgiving student debt (Farrell, Greig and Deadman, 2019). However, despite recent work on the effect of student debt, there has been relatively little research on the broader economic consequences of out-of-pocket education spending on households.[1] Understanding the impact of education spending on households that may have limited downside insurance against income risks is a first-order question given the well-established evidence that households lack adequate financial buffer and exhibit "excess sensitivity" of consumption to income shocks (see, for example, Parker, Souleles, Johnson and McClelland (2013)). Are education costs large enough to affect the standard of living of families? And which economic margins are impacted?

[1] A notable exception is Souleles (2000) who finds that households do a good job at smoothing consumption. I discuss this paper's contribution relative to Souleles (2000) at the end of this section.

Self-employed households provide a unique opportunity to understand the implications of education spending for several reasons. First, self-employed households have a choice between continue running a business and transitioning to wage-earning households. Thus, tracing the response of self-employed households provides an opportunity to examine the job transition channel behind consumption smoothing patterns, an important aspect that has been under-researched in the household finance literature despite the critical role that job mobility plays in life-cycle earnings dynamics (Altonji, Smith JR. and Vidangos, 2013). Second, these households face higher income volatility and have limited insurance against downside risks relative to wage-earning households (Hombert, Schoar, Sraer and Thesmar, 2016) as most small business owners are not eligible to collect unemployment benefits upon exiting self-employment. Focusing on this population can show how families that are most predisposed to being impacted by high education costs due to uncertain income streams smooth consumption. Lastly, anecdotal evidence points to small business owner households with student debt being less likely to hire workers and apply for business loans (Headd, 2014). This study can shed light on the broader implication of education spending on business dynamism when a child enters college.

In this paper, I examine the effect of household financial burden from education spending on consumption and labor decisions of 150,000 self-employed households. Using unique financial transaction data linking small business checking and credit card accounts to personal checking and credit card accounts of self-employed households, I investigate whether households adjust their labor decisions in response to education spending. Specifically, I document novel stylized facts on how business size and exit rates evolve over a child's age profile, and estimate the elasticity of business spending with respect to the changes in college expenditure to capture the magnitude of this response. I find that education spending has meaningful impact on households along several economically important margins. In particular, households downsize business production, exit self-employment, and change household consumption patterns when a child enters college. Small business owners transition from self-employment to wage-earning or gig economy jobs after exiting self-employment, indicating that households may incur loss in non-pecuniary benefits of small business ownership.

Testing the economic impact of education-born financial burden is challenging due to data limitations that complicate measurement. To address this challenge, I collaborate with a large U.S. financial institution – henceforth referred to as *my financial services company*– to construct a de-identified panel dataset of self-employed households with children. My financial services company provides numerous retail products to small businesses and consumers, including checking accounts and credit cards. I exploit my financial services company's large network of retail clients to identify small business owners who hold both business and personal checking accounts. I link business owners to their family members to recover demographic and financial information, and restrict the sample to relatively comparable self-employed households *with children* aged between 14 and 25. This data provide a granular view of the spending patterns of the business owners and their family members.

I document several novel facts using this data. First, the average quarterly spending on education jumps sharply when a child becomes 18 years old, reaches its peak at 21, and declines past the 21 year-old age mark. The conditional mean of education spending constitutes over 20% of households' non-durable consumption at its peak. Second, business expenses, revenues, and investment in machinery decline dramatically when a child turns 18 years old, and exit probability from self-employment rises over a child's age profile. These results indicate that self-employed households' business performance and career choice are tightly linked to the timing of when households are most likely to incur high education spending. One interpretation of this pattern is that households may be smoothing education expenses by downsizing their business, or by transitioning to more stable wage-earning jobs. I test this hypothesis by estimating the elasticity of business performance and household consumption with respect to the changes in education expenditure.

My identification strategy exploits two institutional features about the U.S. schooling system. First, I exploit the norm that the typical college entering age is 18 and 19; and second, that households are billed their first tuition payments in the second to third quarter transition in the year that a child turns 18-19 because academic terms begin in the fall. I combine these features and use a child turning 18-19 interacted with quarter-transition dummies as instruments for a sudden increase in the propensity to spend on a child's education. I employ two-stage least squares (2SLS) to estimate the elasticity of business spending and household consumption (outcome) with respect to the changes in education spending (endogenous variable). To make comparison across similar households, I limit the sample to households with college-entering dependents (18-19 year olds) and to those with near college-entering dependents (15-17 year olds). Given that my instruments are pre-determined institutional features that are orthogonal to potential business outcomes, and to the extent that households with near college-entering dependents serve as a valid control group as they likely undergo similar business and life-cycle dynamics as those with college-entering dependents, my instruments can estimate the local average treatment effect of education spending. I run several 2SLS diagnostic tests to confirm the validity of my instruments.

I show evidence that households scale down business production and exit self-employment in response to increased education spending. I find that education spending of treated households with 18 or 19 year-old dependents increases by 41 log points relative to control households with 15-17 year-old dependents. Treated households cut back on business expenses and generate lower revenues by 4 log points, and have 0.2 percentage points higher probability of exiting in a given quarter. The elasticity of business spending with respect to education spending is -8 log points, which implies that a standard deviation increase in the instrumented education spending leads to 4 percent decline in business expenses and revenues annually. I find that businesses cut spending on machinery and office supplies the most and utilities the least, indicating that businesses are more likely to cut back on variable expenses relative to fixed costs of operating a business. Overall, self-employed households exhibit economically meaningful intensive (expenses and revenues) and extensive margin (exit) responses to increased spending on a child's education.

I find that households do a good job at smoothing consumption when a child becomes college-going age, and they do so by adjusting the composition of household spending. Household consumption *net* of education spending increases by 4 log points for treated relative to control group households. Analyzing detailed consumption categories, I find that households increase medical and restaurant expenses, but they reduce spending on groceries and mortgage payments relative to control households. This intratemporal substitution patterns imply that households incur high non-tuition related discretionary spending associated with a child entering college. While households appear to target a fixed consumption budget by reducing some expense categories while increasing others, the net consumption increases during a child's enrollment spell.

I analyze whether business and consumption effects vary by baseline business growth propensities. Given that the timing of a child's college attendance is predictable, self-employed households may strategically adjust firm growth in advance of a child entering college in order to smooth college expenses. I find evidence that self-employed households downsize their business production during a child's college enrollment spell regardless of their baseline growth rates. Therefore, while households may plan for a child's college expense to some degree, I find limited evidence of sufficient pre-planning to smooth education expenses. I analyze heterogeneity in business and consumption effects by comparing households that remain in self-employment relative to those that eventually exit self-employment. I find that consumption smoothing is largely driven by households that remain in self-employment. Both types of households downsize business production, but the magnitude is more pronounced for exiting households.

These findings raise the question of whether the economic well-being of households that are induced to exit self-employment worsen after exiting. Households can experience financial difficulty if they're unable to transition to wage-earning jobs quickly, or they may actually be better off if they can earn stable income. I track the financial accounts of households after they exit self-employment and find that household consumption and labor income increase dramatically as soon as they exit. Households also earn side income by participating in the gig economy. Thus, while households respond to the increase in education spending by exiting self-employment, it does not translate into negative consumption impact as they can quickly switch jobs and earn side income.

Overall, these results have several implications. First, inferring economic well-being of the self-employed solely based on consumption patterns may mistakenly lead to a conclusion that the financial burden from education spending is not high enough to affect households because (non-tuition related) consumption can mechanically increase when a child enters college. Given that self-employed households adjust labor margins when a child enters college, it is important to take into account other economic margins that are impacted by education spending. Second, despite the muted impact that education spending has on consumption, downsizing a business or switching jobs can generate loss in non-pecuniary benefits of small business ownership (Hurst and Pugsley, 2011) or disutility from job search. Moreover, whether self-employed households can transition to utility-maximizing wage-earning jobs is an open question. Given that most small business owners are not eligible to collect unemployment benefits, they may not have the financial flexibility to search for an optimal outside option.

This study contributes to two main strands of the household finance literature. First, a growing literature documents that household credit access and financial wealth matter for students' educational outcomes (see for example, Stephen Teng Sun and Constantine Yannelis (2016); Vyacheslav Fos, Andres Liberman and Constantine Yannelis (2017); Sarena Goodman, Adam Isen and Constantine Yannelis (2018))[2]. Recent studies capture the cost of education that students face more directly by analyzing the implications of risk-based vs. uniform student loan pricing (Bachas, 2017), or student debt and debt repayment burden (Marco Di Maggio, Ankit Kalda and Vincent Yao, 2019; Daniel Herbst, 2019; Holger M. Mueller and Constantine Yannelis, 2019). Relative to the existing studies that focus on the impact of education costs on student outcomes, this paper examines broader labor market, business outcomes, and consumption behavior at the household-level for the families that bear the cost of education. My results are complementary to those of existing literature, as they indicate that the economic impact of education spending goes beyond the individuals that attend college.[3]

Second, existing studies in household finance examine households' ability to smooth their consumption past transitory income shocks to test the Life-Cycle Permanent Income Hypothesis (LC-PIH) theory. These studies exploit randomized timing of disbursement of economic stimulus (Parker, Souleles, Johnson and McClelland, 2013; Broda and Parker, 2014), tax refunds or rebates (Johnson, Parker and Souleles, 2006; Baugh, Ben-David, Park and Parker, 2018; Caldwell, Nelson

[2]See also the following studies in the economics of education literature on the link between credit constraints and students' educational, financial, household formation, marriage market, and occupational outcomes – Keane and Wolpin (2001); Carneiro and Heckman (2002); Belley and Lochner (2007); Lochner and Monge-Naranjo (2011); Lovenheim (2011); Lovenheim and Reynolds (2013); Cooper and Luengo-Prado (2015); Gicheva (2013); Rothstein and Rouse (2011).

[3]Recent papers that examine the link between student debt burden and entrepreneurship document a negative correlation between student debt burden and entry into entrepreneurship (Baum, 2015; Ambrose, Cordell and Ma, 2015; Krishnan and Wang, 2018). I complement these studies by showing that the family structure of the self-employed – specifically, having a college-going child – explains business growth and exit.

and Waldinger, 2018), household income or liquidity shocks (Gross and Souleles, 2002; Blundell, Pistaferri and Preston, 2008; Blundell, Pistaferri and Saporta-Eksten, 2016; Baker, 2018), or unemployment insurance (Ganong and Noel, 2019), and find that households exhibit *excess sensitivity* (Hall and Mishkin, 1982) of consumption to transitory income shocks. I contribute to this literature by focusing on the consumption smoothing behavior of self-employed households. Despite the fact that nearly 20% of the families in the U.S. derive income from self-employment[4], research on their consumption behavior is limited. Understanding the behavior of these households is crucially important because they are predisposed to being affected by transitory income shocks due to more uncertain and irregular income streams than typical wage-earning households. To my knowledge, this is the first paper to shed light on the consumption smoothing behavior of the self-employed.

The most related study to this paper is Souleles (2000), who uses the Consumer Expenditure Survey and document that households are able to maintain their standard of living (i.e, consumption) as they pay for college. This finding is at odds with a large body of studies that find violation of the LC-PIH. One possible explanation behind this perfect consumption smoothing pattern is that education spending may not have been as burdensome for households as it is now in the time period that Souleles (2000) examines (1980-1993).[5] Alternatively, families may be adjusting other non-consumption margins to meet the financial obligations for their child's human capital investment. This paper tests these hypotheses by analyzing data from more recent time period when the cost of education rose dramatically, and by examining economic margins beyond consumption. By exploiting the unique feature of financial transactions data from the linked accounts of small businesses and their owners, my findings provide a novel perspective that despite the relatively muted consumption response, there is an economically important impact of educational spending on households via labor margins.

The remainder of this paper is organized as follows. Section 2.2 discusses data and sample construction steps for this study. Section 2.3 describes the identification strategy for estimating

[4]According to the FRB (2019), 16-24 percent of households received income from self-employment and occasional side jobs in 2018.

[5]The perfect consumption smoothing can also be explained by the fact that the federal student loan program does not ration borrowers based on the students' credit worthiness. Therefore, easier access to credit might allow households to smooth consumption. However, student loans did not constitute a significant share of household balance sheet in the sample period that Souleles (2000) examines, and thus it is unlikely to be the main explanation for this result.

the impact of education spending on business spending and household consumption. Section 3.5 presents the elasticity estimates and explores heterogeneous response by baseline firm growth rates and households' decision to exit self-employment. Section 2.5 explores the economic well-being of households that exit self-employment. Section 3.7 concludes.

2.2 Data and Stylized Facts

I use de-identified financial accounts data provided by a large U.S. financial institution to construct a panel dataset of self-employed households with children. The final dataset provides a granular view of the business and personal checking account transactions of the business owners and their family members, along with some basic demographic information about the businesses and households. Section 3.2.2 describes the data construction and the sample selection steps. Section 2.2.2 discusses how I measure education spending and business performance.

2.2.1 Data and Sample Construction

The starting point of the sample construction is a universe of 1.3 million small businesses with *active* checking accounts at my financial services company with at least $500 in outflows and 10 transactions for 3 out of 12 consecutive months between October 2012 and April 2018. From this universe, I identify businesses whose owners also have a personal checking account at my financial services company. I use my financial services company's record of de-identified account linkages to link businesses to their owners, where a link is established when an individual has multiple accounts with my financial services company. This process reduces the sample to roughly 550,000 small businesses where both business and its owner's personal account transactions can be tracked. Given that my sampled business owners bank both business and personal accounts at my financial services company, my dataset likely provides a comprehensive view of their financial activity.

Next, I identify other members in the household using my financial services company's record of personal account linkages. This de-identified record connects all household members who also have an account at my financial services company and assigns a unique household identifier. Once the members are identified, I obtain checking and credit card account information of all family members in the household. This allows me to capture total education spending that a given household incurs regardless of which member in the household makes the payment. I aggregate the account-transaction level data into a quarterly firm-household level data that captures the total cash flows into and out of the household's combined accounts. I also obtain basic firm and household characteristics, such as the industry in which a business operates, state of residence, the age of the business, and the age of each family member.

A critical step is obtaining the age of the dependents. One challenge with this is that my financial services company does not provide any information on minors who are younger than 19 years old, which makes identifying households with children difficult. I address this challenge by applying the following rule: when a 19 year old's account is linked to a household for the first time in the data, I recover the dependent's birth year to calculate her age for all years that a business is in operation. This method allows me to capture households with children younger than 19 years old without obtaining additional information on the minors. I restrict the sample to households with the oldest dependents' age between 14 to 25 years, and to those with no more than 5 family members. This leads to a sample of roughly 150,000 self-employed households with children, which serves as the primary sample for analysis.

Table 2.1 reports the descriptive statistics of household and business characteristics for self-employed households with children. The top panel reports statistics for all sample and the bottom panel restricts the sample to households with children aged between 15 and 19 years old. The former sample is used for most of the analysis in this paper, and the latter sample is used for estimating the elasticities. Self-employed households in my sample has on average 3 family members and 1 dependent. The average age of the oldest member of the household is 52 years old, and the oldest dependent's age is 21. Businesses in professional services, other services, and construction industries represent 38% of all businesses, and 50% of all businesses operate in California, New York, and Texas. The descriptive statistics for self-employed households with near college-entering dependents look very similar to the overall sample.

2.2.2 Measurement

I use the transaction-level data to construct measures of household and business spending. I construct three main business performance measures: operating expenses, revenues, and exit. To construct operating expenses and revenues, I first calculate the total cash flows out of and into business checking accounts for each firm and quarter. If a business has multiple checking accounts, they are rolled-up to the firm-level. From these totals, I subtract any *financial* transactions that are unlikely to represent the actual costs or revenues from running a business. These transactions include transfers between accounts, interest, or fee payments. I identify these transactions using my financial services company's categorization of transaction channels (e.g., "Fees" or "Transfers") and confirm its validity using the identity of the counter party in a transaction. I consider the remaining inflows and outflows to be my operating expenses and revenues.

My financial services company provides several categorization variables that tag transactions based on the counter party. I use these variables to categorize operating expenses into finer spending groups: auto maintenance, office supplies/tools, machinery, or utilities. Auto maintenance captures spending at automobile services, repair, or body shops. Office supplies capture spending at home improvement or office supply shops, which likely represents variable input costs for operating a business. Machinery expenses include spending on electronic appliances or industrial equipment. Utility expenses include cable, electric, gas, water, sewer, and other utility services. Finally, I infer exit from the closure of a business checking account or an account's inactivity. Accounts with less than 10 transactions and $500 in outflows for three consecutive months (i.e., inactive) are automatically dropped from the sample. Thus, if a firm drops out from my sample before the end of the period, I consider the last quarter that a firm was active as the exiting quarter.[6]

Household consumption aggregates any durable and non-durable spending from personal checking and credit card accounts of all members in a household. This consumption measure can be broken out into goods (e.g., groceries, fuel, home improvement, etc), services (e.g., restaurants, medical spending, air fares, etc), uncategorizable bill payments using PayPal or wire transfers, utilities (e.g., phone bills, internet, and cable, etc), rent, housing debt (e.g., HELOC, mortgages) or

[6] I do not artificially impose business expenses and revenues to be 0 in the exiting quarter, but cash flows of exiting firms tend to be very small (median expenses and revenues of less than $100).

non-housing debt payments (e.g., auto, personal, and student loans, etc), or credit card payments. For all of the consumption analysis in this paper, I use consumption *net* of education spending by subtracting any expenses categorized as education spending in order to track changes in non-tuition related household expenditures.

To construct total spending on education, I validate both personal and business account transactions that my financial services company pre-categorized as education spending based on the transaction descriptions by verifying the channels (e.g., wire, ACH, etc) and counter parties of transactions. The final measure captures any payments to post-secondary institutions (tuition, fees, room and board), testing service agencies such as the ETS, transfers to the 529 plans since tuitions are typically paid directly from the 529 accounts, and student loan payments to student loan servicing and lending institutions such as Sallie Mae, Navient, Nelnet, etc.[7] My measure does *not* capture room and board expenses if a student lives off-campus, spending on a child's health or dental insurance, and other discretionary financial support that a family provides when a child enters college. It also does not capture any off-the-book borrowing from friends or family, and education spending made using paper checks.

My education spending measure likely underestimates the actual spending that a household incurs when sending kids to college because it captures narrowly defined out-of-pocket education costs associated with tuitions, and because it does not include any payments made in checks even if a payment satisfies my definition of education spending. However, I argue that my measure is nevertheless a useful proxy for capturing the magnitude of financial burden that arises from sending kids to college for two reasons. First, any results using this conservative measure can be interpreted as a lower-bound of the economic impact of education spending. Second, majority of tuition payments occur through wire transfers in the time period I analyze, and thus, it is unlikely that the exclusion of paper checks will alter the narrative of my findings. To verify this, I contacted a mid-sized research university in the Northeast, and the institution shared data on the breakdown of tuition payment types. Table B.1 shows the breakdown of payment types for the fiscal years

[7]Since my financial services company heavily redacts checking account transaction descriptions to protect anonymity, finer categorization into tuition, fees, room and board is not feasible. Thus, I aggregate spending as a broader "education spending".

2013 and 2018. This table verifies that roughly 80% of all payments in terms of dollar amounts are made through wire transfers, which assuages the concern that check payments capture a significant fraction of total education spending.

Table 2.2 reports descriptive statistics of quarterly education spending for self-employed households with children. Panel A and B compare the average education spending incurred by self-employed households with children aged between 15-17 years old (e.g., Near College-Entering Sample) to those with 18-19 years old (e.g., College-Entering Sample). Panels C and D compare the statistics for households with children aged between 18-22 years old (e.g., College-Going Sample) to those with 23-25 years old (College Graduating Sample). In each panel, I report unconditional and conditional statistics of education spending. Top panel shows that the average out-of-pocket quarterly spending on education for households with near-college dependents is around \$321, but it increases by 85% for households with college-entering age. Average spending on education remains high while a child is enrolled in school, but it decreases to \$356 when a child becomes graduating age. A detailed breakdown of education spending shows that the average 529 drawdowns increase when a child become college-entering age[8], and families make student loan payments even during a child's college-enrollment period. Note that the positive student outflows for near college-entering households capture the parents paying for their own student debt.

2.3 Empirical Strategy and Stylized Facts

Section 2.3.1 presents the identification strategy for estimating the elasticity of business and household spending responses to education spending. Section 2.3.2 presents stylized facts on the self-employed households' education spending and business outcomes over a child's age profile to motivate the labor margin consideration when examining the economic impact of education spending.

[8]Roughly 4 percent of my sampled households have a 529 account, slightly higher than the national average of 3 percent (Sager, 2012).

2.3.1 Identification

A simple reduced-form model of the economic impact of education spending is:

$$Y_{h,t} = \alpha + \sum_{t=1}^{4} \gamma_t \mathbb{1}(Q = t) + \beta E_{h,t} + \mathbf{X}'_{h,t}\Theta + \eta_{h,t} \tag{2.1}$$

where I regress the business outcome $Y_{h,t}$ of a self-employed household h in quarter t on total education-related spending $E_{h,t}$ that the same self-employed household incurs, controlling for time trends using quarter dummies $\mathbb{1}(Q = t)$ with $t = 1, 2, 3, 4$ and a vector of baseline business and household covariates, $\mathbf{X}_{h,t}$. These covariates include the age of the business and its owner, the number of dependents in a household, employer status of the business, and a vector of business industry and state of residence indicators. The estimated coefficient β then captures the average effect of education spending on business performance.

A concern with this specification is that β may be subject to selection bias if potential outcomes and education spending are correlated. Let $\{Y^1_{h,t}, Y^0_{h,t}\}$ denote potential outcomes– $Y^1_{h,t}$ is the business outcome that a self-employed household h would obtain after spending on a child's college education, and $Y^0_{h,t}$ is the outcome that would have prevailed in the absence of a child attending college. $D_{h,t}$ is a binary variable that indicates whether a child attends college, and the observed outcome – $Y_{h,t} = Y^0_{h,t} + (Y^1_{h,t} - Y^0_{h,t})D_{h,t}$ – additively captures the causal impact of education spending on potential outcomes, $(Y^1_{h,t} - Y^0_{h,t})D_{h,t}$.[9] Since potential outcomes for any one household is not observed, a naive regression 2.1 that estimates the average difference between households that do or don't spend on education may be subject to selection bias if, for example, high ability self-employed households are more likely than the low ability households to spend more on education. Such ex-ante sorting may lead to over-estimation of β because high ability self-employed households have better potential outcomes, $Y^0_{h,t}$.

An ideal strategy to address this concern is the one where college spending is not correlated with potential business outcomes, such that the conditional independence assumption holds – i.e., $\{Y^1_{h,t}, Y^0_{h,t}\} \perp\!\!\!\perp D_{h,t} | X_h$. While it is unlikely to find a setting where households are randomly as-

[9]Since I do not directly observe a child's college enrollment, I proxy for college attendance with college spending.

signed to send their kids to college, I exploit two institutional features about the U.S. schooling system that provides a setting akin to this ideal experiment. First, I exploit the fact that the typical college entering age in the U.S. is between 18 and 19 years old. Panel A of figure B-1 plots the age distribution of first-year students who enrolled in the 2015-2016 academic year. 18-19 year olds make up over 96% of the total first-year enrollment while younger individuals make up less than 2%. While there are no laws requiring students to reach a specific age at the timing of college enrollment, this pattern may in part be explained by the U.S. compulsory schooling laws that prevents students from dropping out of school until they reach 16-18.[10] Therefore, if a student enrolls in a college within a year or two of her high school graduation, she would most likely be 18 or 19 years old. Second, I exploit the fact that the academic calendar runs from early fall. Since the academic billing cycle starts a few months before the beginning of the academic year, this implies that households are billed their first tuition payment between the second and third calendar quarters.

I combine these features and use a child turning 18-19 interacted with quarter-transition dummies as instruments for a sudden increase in the propensity to spend on a child's education. Specifically, this strategy compares self-employed households with college-entering aged dependents (18-19) to otherwise similar self-employed households with near college-entering aged dependents (15-17) in each quarter-to-quarter transition cell.[11] This strategy allows for examining the impact of education spending on business outcomes for several reasons. First, the academic billing cycle and the norm of entering college at age 18-19 are pre-determined institutional features that are orthogonal to business cycles, industry trends, or other factors that may be correlated with firm performance. Thus, it is unlikely that my instruments are correlated with *potential* business outcomes. Second, the instruments have clear and monotonic impact on education spending. Households either increase or do not spend on education when a child becomes college-going age, but it is unlikely that they would reduce spending on education. Thus, by comparing similar households with dependents that are close in age, but with one group of households having exogenously higher propensity to spend on education due to U.S. schooling norms, I can estimate the effect of

[10]While the legal dropout age varies by state, the compulsory schooling age is 18 in roughly half of the U.S. states.
[11]This strategy is akin to Angrist and Krueger (1991) who used quarter-of-birth instruments to estimate the returns to schooling.

education spending on business performance that is not confounded by selection bias.

Despite the fact that the conditional independence assumption $\{Y_{h,t}^1, Y_{h,t}^0\} \perp\!\!\!\perp D_{h,t} | X_h$ likely holds in this setting, households may be able to predict the cost and timing of their dependents' college enrollment well in advance. Given that sending kids to college is a major financial event for many households, it is plausible that self-employed parents prepare for their child's college education in various ways through saving, adjusting labor supply, or reducing consumption in advance. While some degree of ex-ante college planning is expected, whether households can perfectly plan for this highly anticipated negative income shock is unclear given that a child's college admission decisions and expenditures are not known until a few months before a child enters college when the acceptance letters are sent out and the financial aid packages are revealed. And even if households have perfect foresight, the extent to which they are able to smooth this financial shock is ultimately an empirical question. If households do not respond to short-term transitory shocks (LC-PIH) or can perfectly plan for college costs, this will bias against finding an impact.[12] On the other hand, if households cannot perfectly plan for college costs, a sudden increase in the propensity to spend on education can lead to meaningful impact on households' economic well-being.

Figure 2-4 provides a graphical motivation of this identification strategy. This figure plots the average education spending against calendar time for self-employed households with college-entering dependents – i.e., treated households with 18-19 year olds– and that for otherwise similar households with near college-entering dependents – i.e., control households with 15-17 year olds. Three features are notable. First, treated households spend more on education relative to control households at any given time. Second, treated households have large spending spikes between the second and third quarters, consistent with the fact that treated households experience a large education spending burden just before the fall enrollment due to academic billing cycle. Lastly, control households have relatively flat education spending spending path over time with small spending spikes around fourth quarter each year, which may reflect the costs associated with college applications. Overall, the figure highlights that a child becoming college-entering age coupled with

[12]Pre-planning can bias my estimates upward if a self-employed household grows its business more aggressively prior to a child turning college-entering age. However, figure 2-2 shows that there is no systematic business expansion before a child turns college-entering age. Thus, it is unlikely that anticipation leads to an upward bias.

academic billing calendar can generate a strong first-stage response.

I apply this logic to the following Two-Stage Least Squares (2SLS) model:

$$E_{h,t} = \alpha_h + \sum_{t=1}^{4} \gamma_t^{FS} \mathbb{1}(Q=t) + \sum_{t=1}^{4} \beta_t \mathbb{1}(Q=t) \times \mathbb{1}(Age_{h,t} \in 18, 19) + \mathbf{X}'_{h,t}\Theta^{FS} + \eta_{h,t} \quad (2.2)$$

$$Y_{h,t} = \alpha_h + \sum_{t=1}^{4} \gamma_t^{SS} \mathbb{1}(Q=t) + \rho \widehat{E_{h,t}} + \mathbf{X}'_{h,t}\Theta^{SS} + \varepsilon_{h,t} \quad (2.3)$$

where the first stage and second stage (reduced form) outcomes are logarithmic of total education spending and business performance measure of self-employed household h, respectively. $\mathbb{1}(Q=t)$ and $\mathbf{X}_{h,t}$ are the same as in specification 2.1, and $\mathbb{1}(Q=t)\times\mathbb{1}(Age_{h,t} \in 18, 19)$ are the instruments that interact quarter dummies with a binary indicator for whether the oldest child in a household h is 18 or 19 years old. I include household fixed-effects, α_h, to control for the time-invariant differences across households. This 2SLS-IV model allows me to interpret the coefficient ρ as the average impact of education spending on business performance that is not confounded by selection bias. Standard errors are clustered at the household-level.

I apply two sample restrictions for these regressions. First, as mentioned above, the regression sample restricts self-employed households to those with college-entering aged dependents and with near college-entering aged dependents. This restriction is applied because it allows for apples-to-apples comparison across households. The identifying assumption is that households with younger dependents that are close in age to college-entering dependents serve as a valid control group as they likely undergo similar business and household life-cycle dynamics as treated households, conditional on covariates. Table B.2 reports the average characteristics of the self-employed households with near-college entering dependents (15-17 year olds) and those with college-entering dependents (18-19 year olds). Consistent with the identifying assumption, this table shows that the two groups are very similar in terms of their family and business characteristics, with the main noticeable difference being the age of the dependents and the head of the household being older for the latter group. Second, I exclude any self-employed households that operated a business for less than 1 year. This restriction is made because including new entrants may lead to selection bias if a child being college-age is correlated with the parents' decision to become self-employed.

Therefore, I only consider self-employed households that have operated a business for at least one year in my regression analysis.

Several features of the 2SLS equation 2.2 are noteworthy. First, using logarithmic of outcome variables allows me to interpret ρ as the elasticity of household outcome $Y_{h,t}$ with respect to the predicted education spending $\widehat{E_{h,t}}$. It also has the additional benefit of attenuating the influence of outliers.[13] Second, using firms fixed-effects allows me to interpret β_t as the change in education spending between quarter-to-quarter transitions. For example, β_3 captures the household response between second to third quarter rather than the average effect of the third quarter. Thus, flexibly fitting quarter-to-quarter transitions reveal differential magnitude of education spending induced by the academic billing cycle. Finally, the excluded instruments from the second stage equation 2.3 are the four quarter-transition dummies interacted with an indicator that equals 1 if a child is college-entry age. Since quarter-transition dummies are also included in this equation, the effect of education spending is identified by variation in spending across similar households with varying children's age within each quarter-transition cell. In other words, the estimated effect picks up the differential business response for households with 18-19 year olds relative to other households with younger dependents (i.e., $15 \leq$ age ≤ 17) within each t. In summary, the proposed 2SLS specification considers a child being college-entering age as the *intention-to-treat* instrument based on the institutional feature that the typical college entering age in the U.S. is between 18 and 19.

2.3.2 Stylized Facts

I provide several novel stylized facts on how education spending and business outcomes evolve over a child's age profile. Throughout the rest of the paper, only the age of the oldest dependent is considered when a household has multiple dependents. I focus on the age of the oldest dependents because their college entrance likely constitutes more salient life and financial events for families compared to sending younger dependents to college. Moreover, using the age of the younger child may bias the results upward if households with multiple college-going children have higher education spending.

[13]Measuring outcomes in logarithmic of scaled measure is a conventional approach used to control for baseline differences in the outcomes across units of observations. See, for example, Autor, Dorn and Hanson (2013).

Figure 2-1 provides a visual presentation of how self-employed households' education spending evolves over a child's age profile. Panel A plots conditional (dashed) and unconditional (solid) average quarterly spending on education for all self-employed households with children. The shaded gray area represents the age profile when a dependent is most likely to be enrolled in college. The figure shows a striking pattern of a sudden and large increase in spending around dependents' 18 year-old age mark and persistence in high levels of spending. The conditional averages are higher across all ages cells compared to unconditional averages, and this gap widens particularly during the dependent's college-enrollment period.

Panel B plots the education spending as a share of household's non-durable consumption, where non-durable consumption includes household spending on non-durable goods, services, and utilities net of education spending. Both panels A and B show that education spending– either in levels or as a share of consumption – rises steeply when a child becomes 18, continues to rise, and reaches its peak when a child becomes 21 years old. This is consistent with the feature cost of college attendance tends to rise as students persist through college. The conditional average of education spending represents around 20% of non-durable consumption at its peak, implying that out-of-pocket spending on education is economically meaningful.

To motivate the self-employed households' labor margin consideration for examining the economic impact of education spending, I document how business outcomes evolve over a child's age profile. If a child's college-going behavior does not affect how self-employed parents operate their businesses or their decision to switch jobs, there would be no detectible pattern of business outcomes over a child's age profile. Figure 2-2 plots the average business expenses, revenues, investments, and exit rates over a child's age profile. This figure shows that business operating expenses, revenue, and investment spending decline monotonically, whereas exit rates steeply rise when a child becomes college-going age.

To better account for differences in household and business characteristics in examining the relationship between business outcomes over a child's age profile, I run the following regressions:

$$Y_{h,t} = \alpha + \beta_a \sum_{a=18}^{25} \mathbb{1}(\text{Child's Age} = a) + \mathbf{X}'_{h,t}\Theta + \eta_{h,t} \tag{2.4}$$

where $Y_{h,t}$ denotes quarterly business outcomes scaled by the household (h)-specific baseline average of the outcome, where the baseline period is before a child becomes 18 years old.[14] The household-specific scaling factor allows for comparison across different business size. $\mathbf{X}_{h,t}$ includes household and business characteristics, such as business industries, state of residence, and the age of the oldest member in a household. Therefore, the estimate vector β_a captures the average effect of self-employed households having a-year old dependents, and it's identified by both cross-sectional and within variations of households by comparing business outcomes when a child is a years old relative to when a child is younger than 18.

Figure 2-3 plots the estimated β_a against the dependent's age profile. The estimates are also reported in table B.3. Similar to figure 2-2, self-employed households' operating expenses, revenues, and investment spending declines, whereas the probability of exit rises when a child becomes college-entering age relative to when a child is younger. Specifically, self-employed households spend 1 cent less in operating expense per dollar of baseline average when a child is 18 years old, and this number declines to roughly 4 cents per baseline dollar when a child is 25 years old. Given that the baseline mean of expense is around \$66,000, these estimates correspond to roughly \$6,600 to \$26,400 decline in quarterly expenses in levels, which is as large as a full quartile of the distribution of business expenses. Households also earn 1 cent less per dollar of baseline average and exit probabilities rise by 0.1 -0.5% when a child is college-entering age relative to pre-18 averages.

Overall, these stylized facts provide descriptive evidence that self-employed households' business performance is tightly linked to their dependent's age profile. Over the same age profile during which households incur large education spending for sending kids to college, business performance declines and exit rates from self-employment rises. One interpretation of this evidence is that self-employed households downsize their business production either willingly by choice or unwillingly due to education spending burden that arises from sending kids to college. I examine this labor margin response of the self-employed households in the remaining sections.

[14]For households that do not have pre- (i.e., child is always older than 18) or post-period (i.e., child is always younger than 18 in the sample), I use the household-specific sample mean of the outcome as a scaling factor.

2.4 Main Results

This section presents the estimation results. Sections 2.4.1 and 2.4.2 presents the business spending and household consumption elasticity estimates from the 2SLS model described in section 2.3.1. Section 2.4.3 explores heterogenous business and consumption effects by baseline firm growth rates to test whether self-employed households plan for a child's education spending by adjusting firm growth rates prior to a child a child entering college. 2.4.4 highlights the link between career choice and consumption more directly by comparing business spending and consumption patterns of households that continue to run business to those that exit from self-employment.

2.4.1 The Effect of Education Spending on Business Outcomes

Table 2.3 reports the estimates of the 2SLS-IV model described in section 2.3. Column 1 reports the control mean of each outcome in levels. Columns 2 to 5 report β_t from the first stage and reduced form equations that regress log transformed education spending (endogenous variable) and business outcomes (reduced form outcome) on a set of quarter transition dummies interacted with a college-entry indicator that equals 1 if a child is 18 or 19 years old. Column 6 reports the 2SLS estimate ρ from the second stage equation 2.3. All regressions control for the age of the business and its owner, employer status of the business, industries and state of residence, number of dependents in a household, and employer status of the business. Standard errors are clustered at the household-level.

The first stage estimate in column 2 shows that education spending of self-employed households with college-entering dependents increases by 24 log points compared to other similar households with younger dependents in the same first to second quarter transition cell. Columns 3 and 4 show that the estimated effects increase to 41 and 33 log points in the second to third and the third to fourth quarter transition cells, respectively. The increase in spending in the second half of the year reflects the norm that post-secondary institutions bill incoming students just before the semester starts. Converting the first stage effect into an implied dollar magnitude, I find that a 41 log point increase translate into \$162, or 46%$\approx\frac{\$162}{\$354}$ of the sample mean.[15] Therefore, the the first stage impact is economically meaningful.

The reduced form estimates show that households with college-entering dependents cut back on business expenses, generate lower revenues, and have higher probability of exit. Columns 2 to 5 report that business operating expenses and revenues decline by 2 to 4 log points for the treated relative to control. These estimates correspond to reductions of \$2,732 in expenses and \$2,598 in revenues, or reductions of 15-20% of median expense and revenue, respectively. Since the control mean of log transformed expenses (revenues) is -0.10 (-0.14), the treatment effect in a given quarter is one fifth in size relative to the sample means. The probability of exit increases by 0.2 percentage points for treated relative to control group in each quarter, or roughly 22% of the sample mean. Figure 2-5 illustrates these results graphically. This figure plots β_t from the first stage and reduced form regressions for each quarter-transition cell.

The detailed breakdown of operating expenses explores heterogeneous intratemporal substitution margins. Specifically, the magnitude of the treatment effect indicates whether businesses cut expenses equally across various spending categories or whether there are margins they cut by more relative to others. For households with college-entering dependents, business spending on office supplies/tools and machineries decline by roughly 6 log points. However, there is no differential spending response on utilities for treated relative to control households, implying that businesses differentially cut back on margins that are easier to adjust.[16]

Column 6 reports the 2SLS elasticity of business spending to education spending. The results indicate that an increase in *instrumented* spending on education by a log point leads to a reduction in business expenses and revenues by 7-8 log points. Since one standard deviation of instrumented education spending is 0.13, these estimates imply that a one standard deviation increase in the instrumented education spending leads to a one percentage point ($.13 \times -.08 \approx -.0104$) decline in

[15]Note that the estimated effect β_t picks up the average difference in the log transformed outcome between treated and control. Thus, treatment effects in log points can be transformed into the same unit as the outcome by raising e to the β_t power and subtracting 1 (i.e., e^{β_t}-1). To obtain the implied dollar terms, I multiply this number with the sample average of each outcome.

[16]In my sample, only 15 percent of small businesses in my sample are employer firms that have regular payroll expenses. Thus, I use the employer status of the firm as a baseline covariate instead of examining change in payroll expense as an outcome.

business expenses and revenues per quarter, or roughly 4 percentage points annually. The extensive margin estimate for exit probability is 0.3 percentage points, which is about as large as 33% of the sample mean. Figure 2-6 illustrates these results graphically. Overall, education spending induced by a child's college-entry age generates economically meaningful intensive and extensive margin responses for small businesses.

2.4.2 The Effect of Education Spending on Household Consumption

Table 2.4 reports the elasticity of consumption to changes in education spending. To capture how non-education related consumption responds to increased education spending, my consumption measure aggregates spending incurred by each member of a household through their personal checking and/or credit card accounts *net* of education spending. I estimates the elasticity using this *net* consumption measure as well as several sub-categories of consumption, such as spending on non-durable goods, durable goods, services, utility, and mortgage. The variable description table B.6 reports details on the spending types that each sub-category captures. Note that the total consumption measure includes more categories than those reported in table 2.4.[17]

Columns 2 to 5 show that consumption net of education expenditure increases by 2 to 4 log points in a given quarter for the self-employed households with college-entering dependents relative to those with younger dependents. Examining the detailed consumption sub-categories, I find that spending on non-durables – particularly on groceries– and mortgage payments decline by 4 to 8 log points for treated relative to the control group. On the other hand, spending on durable goods and services – particularly medical and restaurant expenses – rise by 2 and 8 log points, respectively, for treated relative to the control households. I find that there is no differential spending patterns for treated relative to control households on utilities. Column 6 reports that the elasticity of consumption to education spending is 8.5.

Figures 2-7 and 2-8 illustrates these results graphically. In both figures, panel A reports reduced form estimates that regress household consumption outcomes on a set of quarter transition dummies interacted with treatment status that equals 1 if a child is 18 or 19 years old, and panel B

[17]For example, other categories include credit card payments, non-housing debt payments, HELOC, taxes, and other miscellaneous online payments.

reports the 2SLS-IV estimates. These figures show that while the estimated effect on net consumption is positive in every quarter-transition cells, there is substantial heterogeneity in the spending response by detailed categories.

There are several ways to interpret the consumption results. Households may be cutting back some margins in order to increase the discretionary spending associated with a child entering college that is not directly captured by tuition costs, or self-employed parents may change spending patterns due to adjustments in consumption-leisure tradeoff after sending kids to school. For example, the increase in medical spending may capture the costs associated with increasing doctor's visits in order to provide a proof of physical examination and vaccine receipts, as mandated by many colleges in the U.S. It is also reasonable to interpret this as self-employed parents increasing doctor's visits because they have more time in hand. Similarly, the increased restaurant spending and decreased spending on groceries may reflect reduction in home production due to a child leaving home or parents dining out more. While both of these hypotheses likely contribute to the changes in consumption patterns, the next section presents results that show that the former hypothesis – households adjust consumption in order to invest in a child's human capital – is likely to be the primary driver of the changes in household consumption.

2.4.3 Strategic Planning

Do self-employed households plan for a child's college entry by adjusting their firm size prior to a child entering college? If households strategically downsize the firm before a child turns 18 in order to smooth college expenses, such strategic planning will reveal that businesses with lower baseline growth rates fare better with college expenses. On the other hand, it is plausible that self-employed households with high baseline firm growth rates smooth education spending shocks better relative to households that operate low growth firms. To explore these hypotheses, I calculate firm-specific year-over-year average sales growth *before* a child turns 18 among self-employed households with college-entering dependents. I then group firms into quartile bins, where bin 1 includes firms with the lowest pre-18 average growth rates and bin 4 the highest. For each growth bin, I estimate equation 2.4 to examine how β_a evolves over a child's college enrollment spell.

Figure 2-9 plots β_a against a child's age for each subgroup of firms with varying baseline growth propensities. The estimated coefficients capture the effect of self-employed households having an a-year old dependent on business expenses and revenues. This figure shows that the lowest growth firms show the largest business response when a child turns 18 while the highest growth firms show the smallest response. However, business response in all growth groups converge over a child's enrollment spell– expenses and revenues decline when a child enters college, but they bounce back once a child reaches graduation age for all growth bins. This spending pattern is consistent with self-employed households reducing spending during a child's college enrollment. Appendix figure B-2 reports the results for spending on machinery and exit probabilities, and show that firms in the lowest growth bin cut back on machinery the most and are most likely to exit when a child enters college.

Figure 2-10 explores household consumption response net of education spending by baseline growth bins. For all growth bins, net consumption increases during a child's college enrollment spell and declines once a child reaches 21. This spending pattern is consistent with households temporarily increasing non-tuition related discretionary spending associated with sending kids to college rather than self-employed parents permanently changing their consumption behavior due to a child leaving home. If the adjustment in consumption-leisure associated with a child leaving home is the primary driver behind the changes in consumption patterns, these patterns would remain the same after a child reaches graduation age. Consistent with figure 2-9, the lowest growth bin increases net consumption the least while they cut back on non-durable goods spending the most. Appendix figure B-3 reports the results for spending on durable goods and services.

Overall, these figures highlight several takeaways. First, it is unlikely that self-employed households can perfectly plan for a child's college spending by strategically growing or downsizing the firm in advance because businesses in all growth bins cut back expenses when a child enters college. Second, there is substantial heterogeneity in the magnitude and timing of business effects when a child enters college by baseline growth propensities. High growth firms tend to cut back on business spending less and later relative to low growth firms, and they are more likely to increase net consumption. Lastly, there is less heterogeneity in household consumption response relative to business effects across growth bins. This implies that examining household consumption may not be sufficient to understand the overall welfare implication of self-employed households. Even if self-employed households of all business growth types increase net consumption, there is substantial variation in the degree to which education spending affects labor decisions.

2.4.4 The Link between Career Choice and Consumption

So far, the results indicate that while education spending leads to increase in non-education related household consumption, it induces self-employed households to downsize their business production and exit from self-employment. One interpretation of this pattern is that self-employed households are adjusting their labor margins in order to meet the financial obligation of investing in a child's human capital. This section explores this link between career choice and consumption more directly by comparing business and consumption response of households that continue to be self-employed relative to those that ultimately exit from self-employment. In order to track the outcome paths by subgroups, I return to employing the reduced-form model 2.4, which estimates the treatment effect at each age $a \geq 18$ relative to the pre-18 baseline average.

Table 2.5 reports education and business spending responses for self-employed households that continue to remain in self-employment ("stay") to those that switch their career to wage-earning employment ("exit"). To be specific, the "exit" sample contains self-employed households that eventually exit at some point during my sample period (2012Q4 to 2018Q2), whereas the "stay" sample contain households that never exit. Therefore, the outcomes for the "exit" sample reflect the spending dynamics *before* the owners exit self-employment. Comparing columns 2 through 4 to 6 through 8, I find that the business response is larger for exit relative to stayer sample at each of dependent's age profile a, whereas columns 1 and 5 show that stayer households spend more on education than those that exit. This result not surprising given that businesses that ultimately exit would be expected to downsize their production before they exit. However, this result is not a simple mechanical artifact purely driven by exiting households because the stayer households also downsize their business production. Figure B-4 provides a visual guidance of these results.

Table 2.6 compares the consumption response for households that stay and exit. As before, the consumption outcomes for the "exit" sample reflect the consumption path before the business owner exits self-employment. Comparing columns 1 to 5, I find that the increase in total consumption net of education expenditure is larger for households that remain in self-employment relative to those that eventually exit. Exiting households cut back on groceries by more than those that remain in self-employment, whereas the increase in restaurant and medical spending is mainly driven by households that remain in self-employment. Overall, both exiting and stayer households do a good job at smoothing consumption, but households that remain in self-employment appear to smooth consumption better relative to those that exit eventually. Figure B-5 provides a visual guidance of these results.

To recap, the main goal of this section is to explore whether self-employed households adjust their labor margins by downsizing business production or switching careers in order to smooth consumption. I find that households that remain in self-employment increase both education spending and non-education related consumption more than those that eventually exit from self-employment, even though they scale back on business production less. One interpretation of this result is that exiting households are being forced to return to wage-earning employment in order to smooth consumption from transitory education spending shocks. Otherwise – if exit decisions are voluntary and unrelated to a child's education spending – we would not expect to see any heterogeneous response in household consumption by the self-employed parents' career choice decisions.

To better understand whether exit rates are linked to the child's college-going behavior, figure B-6 presents the share of households that exit from self-employment over a child's age profile (top panel) and the cumulative exit rates over a child's age profile conditional on exiting (bottom panel). Two features are notable. First, the share of businesses that exit increases as the child gets older. Second, conditional on exiting, most households exit when the child is likely to be enrolled in college. The cumulative exit rates indicate that less than 20% of firms have exited when the dependent is younger than 18 years old, but more than 80% of the firms have exited by the time the dependent reaches the graduating age. Therefore, the descriptive evidence is consistent with the hypothesis that a child's college-going behavior is linked to the self-employed households' career choice decisions.

These findings raise the question of whether the economic well-being of households that are induced to exit from self-employment worsen after exiting. These households can experience economic difficulty if they're unable to transition to wage-earning jobs quickly, or they may actually be better off if they can quickly earn stable income. The next section addresses these questions by tracking the financial accounts of households *after* they exit self-employment.

2.5 Post-Exit Response

This section explores the household consumption and income paths *after* self-employed households exit from self-employment. Subsection 2.5.1 provides reduced form evidence of how consumption and income paths evolve after exiting. Subsection 2.5.2 investigates variation in these effects by the timing of exit.

2.5.1 Consumption and Income Paths

I track consumption and income paths of households that exit from self-employment by estimating the following equation:

$$Y_{h,t} = \sum_{s=1}^{8} \beta_s \mathbb{1}(\text{Exit}_{t+s}) + \mathbf{X}'_{h,t}\Theta + \eta_{h,t} \tag{2.5}$$

where $Y_{h,t}$ denotes quarterly consumption or income scaled by its pre-exit average for self-employed household, h. The household-specific scaling factor allows for comparison across different households with varying levels of pre-exit non-business income. $\mathbf{X}_{h,t}$ includes the same household and business covariates used in the previous section. The estimate vector β_a captures the average effect of post-exit response on the outcome s periods after exiting from self-employment. Standard errors are clustered at the household level.

Table 2.7 reports the consumption and income paths after a household exits self-employment. I track the results for 2 years (or 8 quarters) after they exit from self-employment. Column 1 shows that households that exit from self-employment immediately increase consumption relative to the pre-exit mean. In the first quarter after exiting, households increase consumption by 11 cents per dollar of pre-exit consumption mean, and this number increases to 19 cents 2 years after exiting. Column 2 shows that households experience more than 100% increase in the labor income as soon as they exit from self-employment as they earn $1.12 per dollar of pre-exit mean of labor income, and this number increases to $2.52 after 2 years. Column 3 shows that households also earn side income by participating in the gig economy after they exit, and continue to earn more gig income in the periods following exit even though labor income also goes up. It is worth noting that the share of households that receive unemployment benefits does not change after exiting self-employment. This result is expected given that non-employer small business owners are not typically eligible to collect unemployment benefits.

Overall, I find that households that exit from self-employment do a good job at smoothing consumption even after exiting and that they transition to wage-earning jobs quickly. A natural interpretation of these results is that the self-employed parents become wage-earners in order to smooth consumption, or to meet the financial obligations of investing in a child's human capital. Under this interpretation, high education costs may have implications for welfare even if households do a good job at smoothing consumption because households may incur significant loss of non-pecuniary benefits of small business ownership (Hurst and Pugsley, 2011) or derive disutility from switching careers. Moreover, the fact that the business owners have limited downside insurance may play a role in self-employed owners quickly transitioning to wage-earning jobs or entering the gig economy. However, whether this transition is optimal is unclear given that the limited downside insurance may force people to take worse jobs than they would otherwise take if they are eligible for unemployment benefits and have the flexibility to search for jobs longer.

2.5.2 Heterogeneity by Exit Timing

One implication of the result that households can quickly transition to wage-earning jobs is that they may have timed their exit from self-employment. In particular, it is plausible that households that exit before a child reaches college-entering age may be forward-looking households that exited preemptively in order to smooth consumption, whereas those that exit after a child becomes 18 may be those that are induced to exit due to high spending burden. Thus, to the extent that exit timing captures financial constraint that a household experiences, we would expect to see that households that are forced to exit earn more labor and gig economy incomes to smooth the financial shock.

Figure 2-11 investigates post-exit consumption and income responses for households that exit before a child becomes 18 (blue) and those that exit after (red). Consistent with the results in table 2.7, I find that both consumption and income increase immediately after households exit self-employment. However, the consumption response is lower while the labor income and gig-economy income responses are higher for households that exit after a child is 18 relative to households that exit before a child is 18. This result is consistent with the hypothesis that exit timing may capture the financial burden that households may experience from sending a child to college.

2.6 Conclusion

This paper studies the economic implications of education spending for self-employed households. When a child goes to college, households downsize business production and exit self-employment. While they do a good job at smoothing consumption, they change the composition of spending by reducing debt payments and groceries and increasing medical and restaurant spending. This result is consistent with households adjusting intratemporal consumption margins to smooth costs associated with sending kids to college. Despite the fact that high college costs induce households to exit self-employment, this does not lead to a permanent decline in consumption even after they exit because they can quickly transition to wage-earning employment and make side income from entering the gig economy. These results suggest that when a child enters college, consumption smoothing may come at the expense non-pecuniary benefits of small business ownership.

This paper provides several policy implications. First, I find that education-related financial burden extends beyond the individuals who attend college and has broader impact on the families that share the education spending burden. Policy-makers must take the financial health of families into account to better assess the implications of rising education costs. Second, this paper finds that the ability to participate in the gig economy plays a crucial role for household consumption smoothing. Thus, policies promoting alternative labor market opportunities can help families meet the financial obligations of investing in a child's human capital. Lastly, policies aimed at fostering entrepreneurial activity must consider the family structure of the business owner. To the extent that paying for college constitutes a financial hindrance for operating a business, policies that provide subsidies or tax breaks when business owners have to pay for a child's education may revitalize business dynamism in the U.S.

There are many promising directions for future research. One is examining the household-level response beyond student outcomes to study the broader economic implications of rising education costs. Another direction is to examine how wage-earning households adjust their labor supply in order to smooth consumption when a child attends college. Finally, there is a dearth of credible evidence on the financial health of self-employed households. To the extent that these households face different income risks relative to typical wage-earning households, studying this population can broaden our understanding of their financial behavior and decision making.

Bibliography

Altonji, Joseph G., Anthony A. Smith JR., and Ivan Vidangos. 2013. "Modeling Earnings Dynamics." *Econometrica*, 81(4): 1395–1454.

Ambrose, Brent W., Lawrence Cordell, and Shuwei Ma. 2015. "The Impact of Student Loan Debt on Small Business Formation." *Working paper*.

Angrist, Josh D., and Alan B. Krueger. 1991. "Does Compulsory Schooling Attendance Affect Schooling and Earnings?" *Quarterly Journal of Economics*, 106(4): 979–1014.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review*, 103(6): 2121–2168.

Bachas, Natalie. 2017. "Pricing, Selection, and Welfare in the Student Loan Market: Evidence from Borrower Repayment Decisions." *Working paper*.

Baker, Scott. 2018. "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data." *Journal of Political Economy*, 126(4): 1504–1557.

Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A. Parker. 2018. "Asymmetric Consumption Response of Households to Positive and Negative Anticipated Cash Flows." *Working paper*.

Baum, Sandy. 2015. "Does Increasing Reliance on Student Debt Explain Declines in Entrepreneurial Activity?" *Research Report. Washington: Urban Institute*.

Belley, Philippe, and Lance J. Lochner. 2007. "The Changing Role of Family Income and Ability in Determining Educational Achievement." *Journal of Human Capital*, 1(1): 37–89.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review*, 98(5): 1887–1921.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. "Consumption Inequality and Family Labor Supply." *American Economic Review*, 106(2).

Broda, Christian, and Jonathan A. Parker. 2014. "The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption." *Journal of Monetary Economics*, 68.

Caldwell, Sydnee, Scott Nelson, and Daniel Waldinger. 2018. "Tax Refund Expectations and Financial Behavior." *working paper*.

Carneiro, Pedro, and James Heckman. 2002. "The Evidence on Credit Constraints in Post-Secondary Schooling." *The Economic Journal*, 112: 989–1018.

Cooper, Daniel, and Maria José Luengo-Prado. 2015. "House Price Growth when Children are Teenagers: A Path to Higher Earnings?" *Journal of Urban Economics*, 86: 54–72.

Di Maggio, Marco, Ankit Kalda, and Vincent Yao. 2019. "Second Chance: Life without Student Debt." *NBEr Working Paper 25810*.

Farrell, Diana, Fiona Greig, and Erica Deadman. 2019. "Student Loan Payments: Evidence from 4 Million Families." *JP Morgan Chase Institute*.

Fos, Vyacheslav, Andres Liberman, and Constantine Yannelis. 2017. "Debt and Human Capital: Evidence from Student Loans." *Working paper*.

FRB, Federal Reserve Board. 2019. "Report on the Economic Well-Being of U.S. Households in 2017 - May 2018." *Survey of Household Economics and Decision Making (SHED)*.

Ganong, Peter, and Pascal Noel. 2019. "Consumer Spending During Unemployment: Positive and Normative Implications." *American Economic Review*, 109(7): 2383–2424.

Gicheva, Dora. 2013. "In Debt and Alone? The Causal Link between Student Loans and Marriage." *working paper*.

Goodman, Sarena, Adam Isen, and Constantine Yannelis. 2018. "A Day Late and a Dollar Short: Liquidity and Household Formation amond Student Borrowers." *working paper*.

Gross, David B., and Nicholas S. Souleles. 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Quarterly Journal of Economics*, 117(1): 149–185.

Hall, Robert, and Frederic Mishkin. 1982. "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." *Econometrica*, 50(2): 461–81.

Headd, Brian. 2014. "Small Business Facts: Student Debt Among Young Entrepreneurs." *SBA Office of Advocacy*.

Herbst, Daniel. 2019. "Liquidity and Insurance in Student Loan Contracts: The Effects of Income-Driven Repayment on Borrower Outcomes." *Working paper*.

Hombert, Johan, Antoinette Schoar, David Sraer, and David Thesmar. 2016. "Can Unemployment Insurance Spur Entrepreneurial Activity? Evidence from France." *working paper*.

Hurst, Erik, and Benjamin Wild Pugsley. 2011. "What Do Small Businesses Do?" *Brookings Papers on Economic Activity*, , (2): 73–142.

Johnson, David S., Jonathan A. Parker, and Nicholas Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review*, 96(5): 1589–1610.

Keane, Michael P., and Kenneth I. Wolpin. 2001. "The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment." *International Economic Review*, 42(4): 1051–1103.

Krishnan, Karthik, and Pinshuo Wang. 2018. "The Cost of Financing Education: Can Student Debt Hinder Entrepreneurship?" *Management Science*, Forthcoming.

Lochner, Lance J., and Alexander Monge-Naranjo. 2011. "The Nature of Credit Constraints and Human Capital." *American Economic Review*, 101(6): 2487–2529.

Lovenheim, Michael F. 2011. "The Effect of Liquid Housing Wealth on College Enrollment." *Journal of Labor Economics*, 29(4): 741–771.

Lovenheim, Michael F., and C. Lockwood Reynolds. 2013. "The Effect of Housing Wealth on College Choice: Evidence from the Housing Boom." *Journal of Human Resources*, 48(1): 1–35.

Mueller, Holger M., and Constantine Yannelis. 2019. "Reducing Barriers to Enrollment in Federal Student Loan Repayment Plans: Evidence from the Navient Field Experiment." *Working paper*.

Parker, Jonathan A., Nicholas Souleles, David S. Johnson, and Robert McClelland. 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review*, 103(6): 2530–53.

Rothstein, Jesse, and Cecilia Elena Rouse. 2011. "Constrained after college: Student loans and early-career occupational choices." *Journal of Public Economics*, 95(1-2): 149–163.

Sager, Michelle. 2012. "Higher Education: A Small Percentage of Families Save in 529 Plans." General Accountability Office GAO-13-64, Report to the Chairman, Committee on Finance, U.S. Senate.

Souleles, Nicholas. 2000. "College Tuition and Household Savings and Consumption." *Journal of Public Economics*, 77: 185–207.

Sun, Stephen Teng, and Constantine Yannelis. 2016. "Credit Constraints and Demand for Higher Education: Evidence from Financial Deregulation." *Review of Economics and Statistics*, 98(1): 12–24.

Figures and Tables

Figure 2-1: Average Spending on Education over a Child's Age Profile

Notes: This figure plots the average quarterly education spending by the age profile of the oldest dependent in a household. Panel A reports the unconditional (solid) and unconditional (dotted) averages in spending, and panel B reports the average spending as a share of self-employed households' non-durable consumption. The shaded area in grey indicates the period in which the dependents are most likely to be enrolled in college (18-23). Total education spending captures payments to post-secondary institutions (tuition, fees, room and board), testing service agencies (e.g., ETS), and student loan payments to student loan servicing and lending institutions (e.g., Sallie Mae, Navient, Nelnet, etc). Non-durable consumption includes household spending on non-durable goods and services excluding a child's education spending.

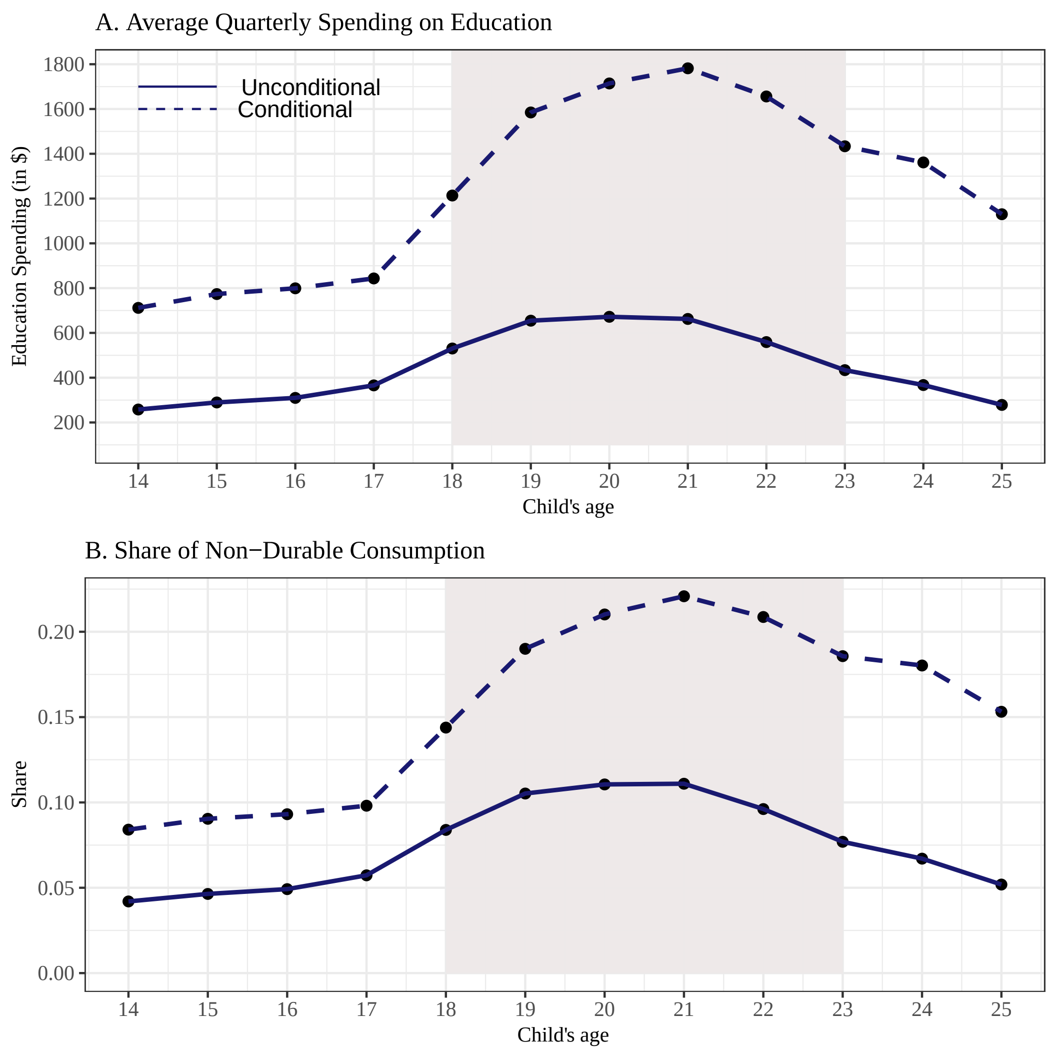

Figure 2-2: Average Business Outcomes over a Child's Age Profile

Notes: This figure plots the unconditional average of quarterly business outcomes by the age profile of the oldest dependent in a household. The shaded area in grey indicates the period in which the dependents are most likely to be enrolled in college (18- 23).

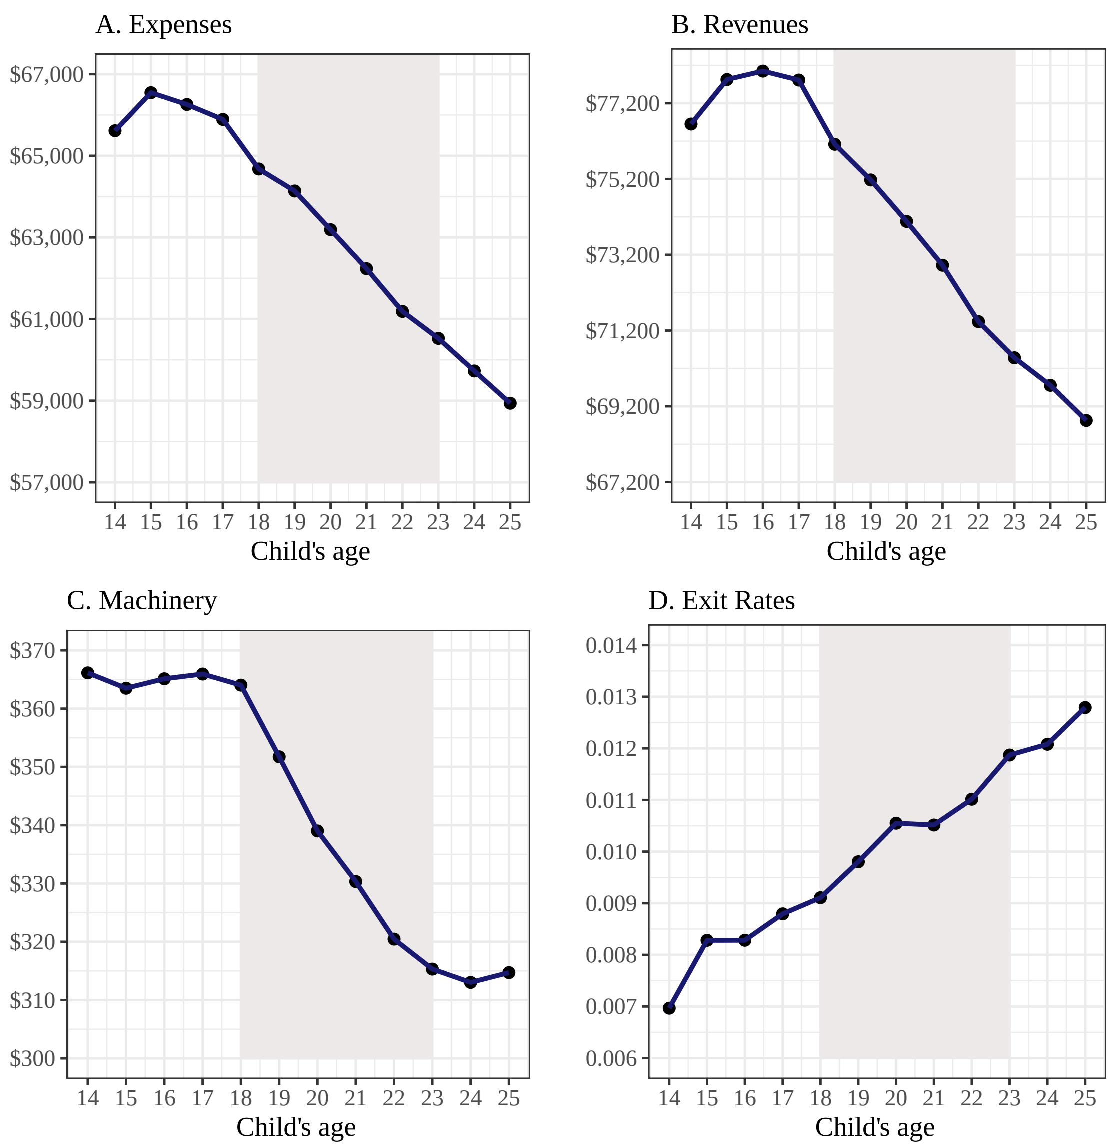

Figure 2-3: Reduced Form Effects on Business Outcomes

Notes: This figure plots β_a of equation 2.4, which captures the average effect of self-employed households having an a-year old dependent. The outcomes are scaled by household-specific baseline average, where the baseline period denotes the period before a child becomes 18 years old. The baseline period is the entire sample range for households that have dependents but do not have 18-19 year olds during the sample period I analyze. All regressions control for the age of the business owner, the number of dependents in a household, business industry and location, and employer status of a business. Whiskers show 95% confidence intervals. Solid fitted lines are estimated from local regressions.

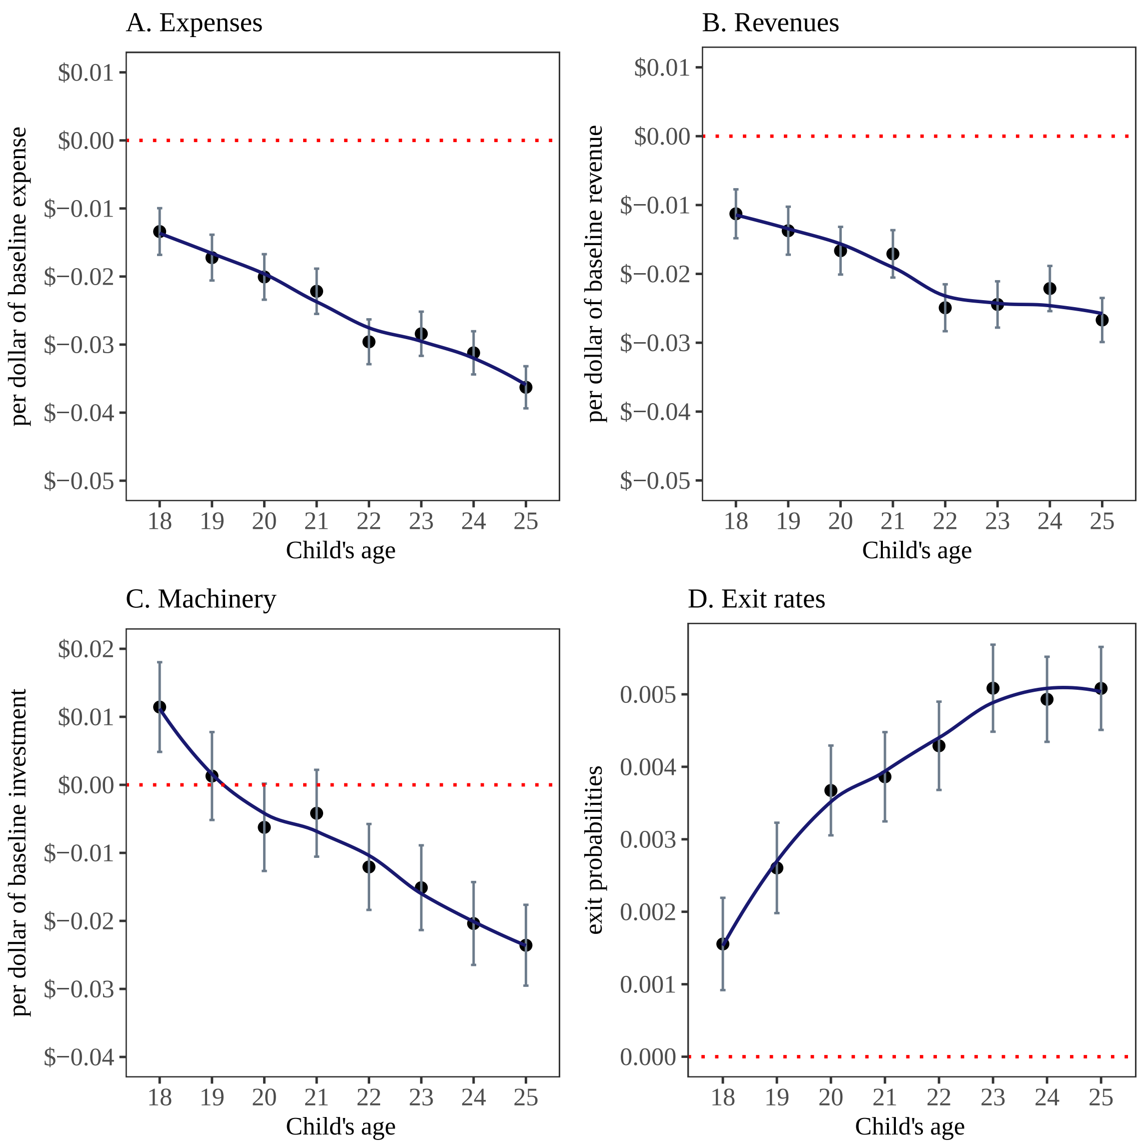

Figure 2-4: Education Spending Patterns of Treated and Control Households

Notes: This figure plots quarterly education spending patterns against calendar time for self-employed households with college-entering dependents ("treated" group with 18-19 year olds) and those with near college-entering dependents ("control" group with 15-17 year olds). Dotted vertical lines mark the third quarter of each calendar year to show that seasonal spikes in education spending coincides with the academic billing cycle. This figure illustrates the motivation behind using the age of a child interacted with quarter-dummies as instruments for education spending.

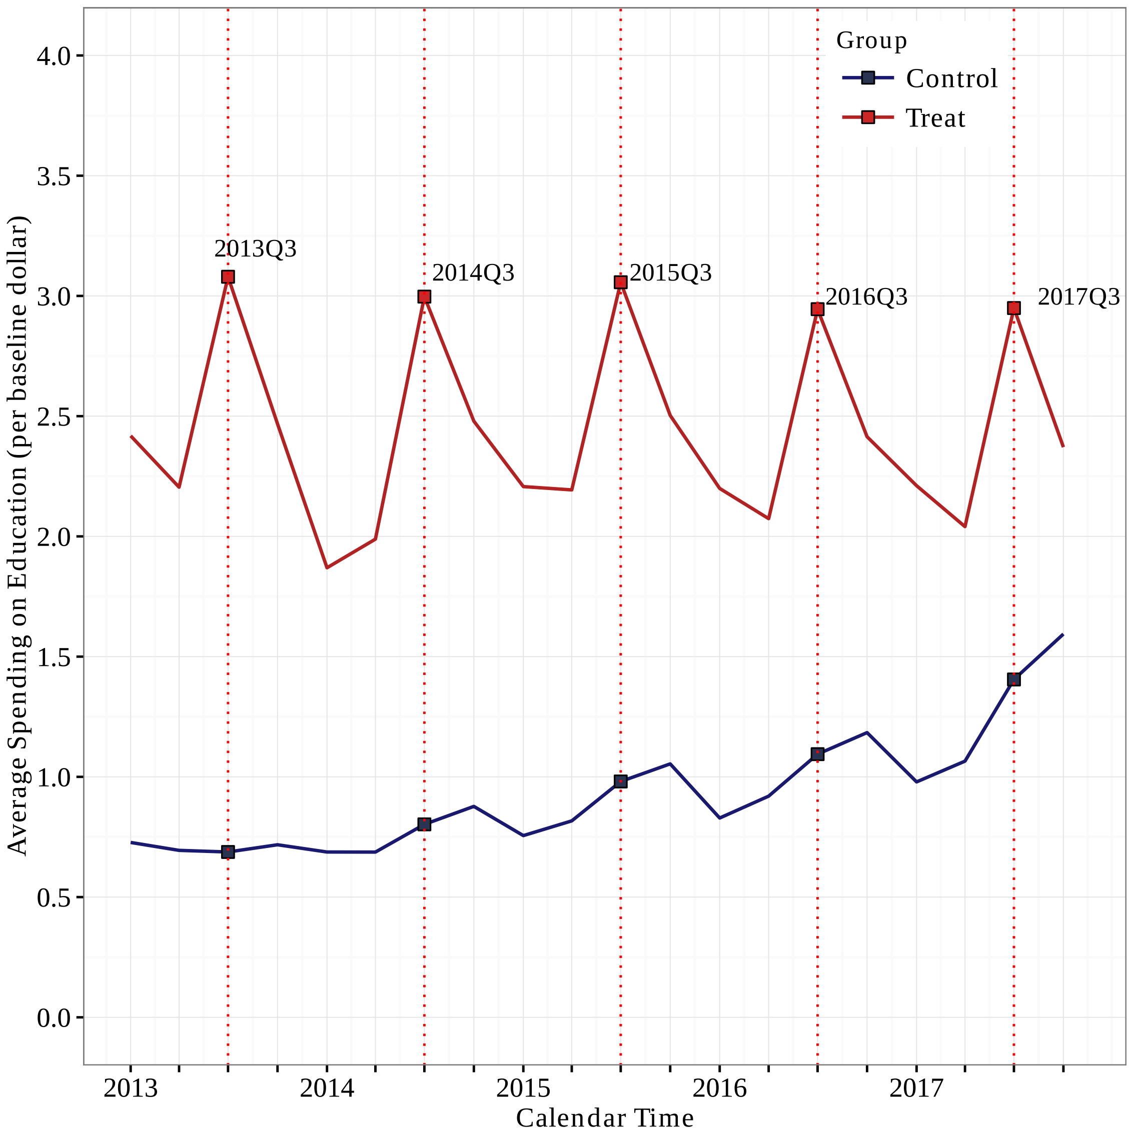

Figure 2-5: The Effect of Education Spending on Business Performance

Notes: This figure plots the first stage and reduced form estimates β_t from equation 2.2, which capture the effect of a dependent becoming college-entering age on education spending and business outcomes for each quarter-to-quarter transition. The sample is restricted to self-employed households with dependents aged between 15 and 19, where households with younger than 18 year olds are in the control group and those with 18 or 19 year olds are in the treated. The IV estimates – ρ – from equation 2.3, which capture the business effects induced by increased education spending, are annotated in panel B. All regressions control for the age of the business and its owner, the number of dependents in a household, business industry, state of residence, and employer status of a business. The estimates are also reported in table 2.3. Whiskers show 95% confidence intervals and standard errors are clustered at the household-level.

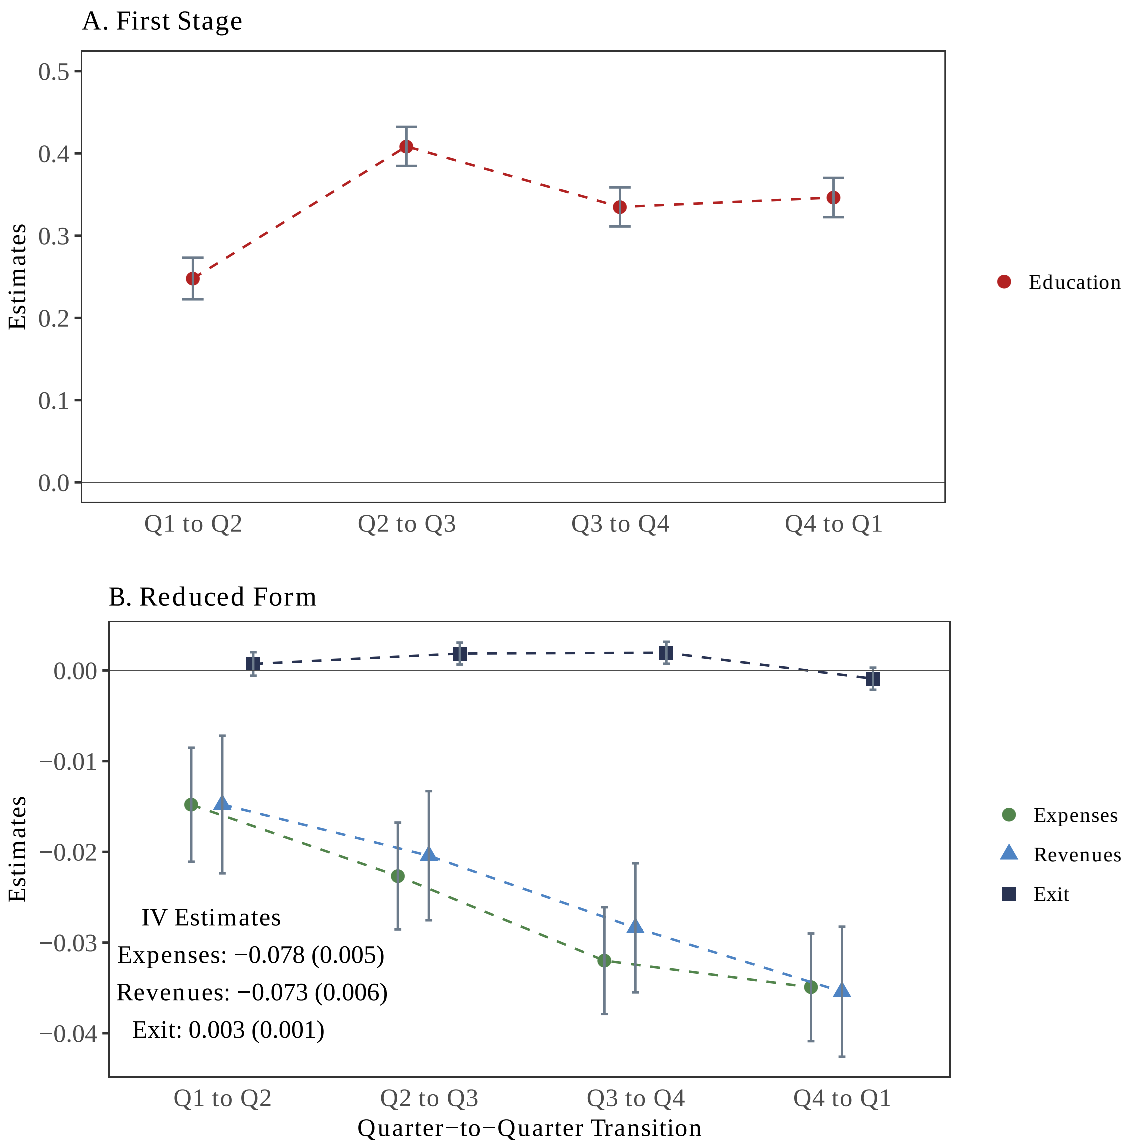

Figure 2-6: The Effect of Education Spending on Business Expenditures

Notes: This figure plots the reduced form and IV estimates of the effect of education spending on various business expenditures. The sample is restricted to self-employed households with dependents aged between 15 and 19, where households with younger than 18 year olds are in the control group and those with 18 or 19 year olds are in the treated. Panel A plots the reduced form estimates β_t, which capture the effect of a dependent becoming college-entering age on various business expenditures for each quarter-to-quarter transition. Panel B plots the IV estimates – ρ – from equation 2.3, which capture the business spending effects induced by increased education spending. Auto maintenance includes spending at auto repair shops. Home improvement includes spending on office furnitures, restoration services, or upholstery. Machinery includes spending on industrial equipments, durable appliances, or expenditures at auto dealerships. Utilities include cable, electric, gas, telecommunications, or water. All regressions control for the age of the business and its owner, the number of dependents in a household, business industry, state of residence, and employer status of a business. The estimates are also reported in table 2.3. Whiskers show 95% confidence intervals and standard errors are clustered at the household-level.

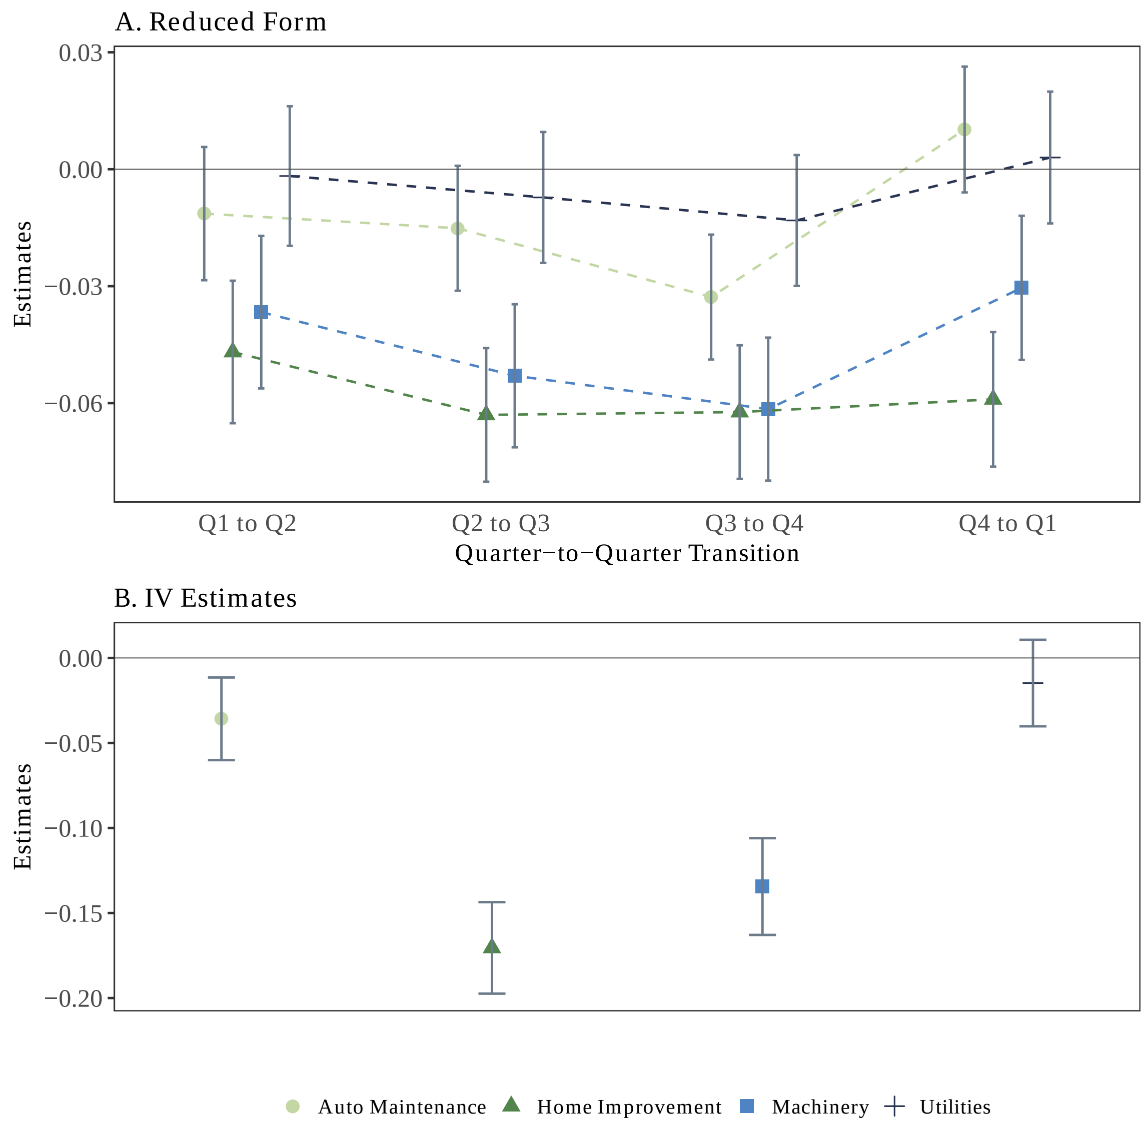

Figure 2-7: The Effect of Education Spending on Household Consumption

Notes: This figure plots the reduced form and IV estimates of the effect of education spending on household consumption *net* of education spending. The sample is restricted to self-employed households with dependents aged between 15 and 19, where households with younger than 18 year olds are in the control group and those with 18 or 19 year olds are in the treated. Panel A plots the reduced form estimates β_t, which capture the effect of a dependent becoming college-entering age on household consumption for each quarter-to-quarter transition. Panel B plots the IV estimates – ρ – from equation 2.3, which capture the consumption effects induced by increased education spending. Detailed information on consumption categories are available in table B.6. All regressions control for the age of the business and its owner, the number of dependents in a household, business industry, state of residence, and employer status of a business. The estimates are also reported in table 2.4. Whiskers show 95% confidence intervals and standard errors are clustered at the household-level.

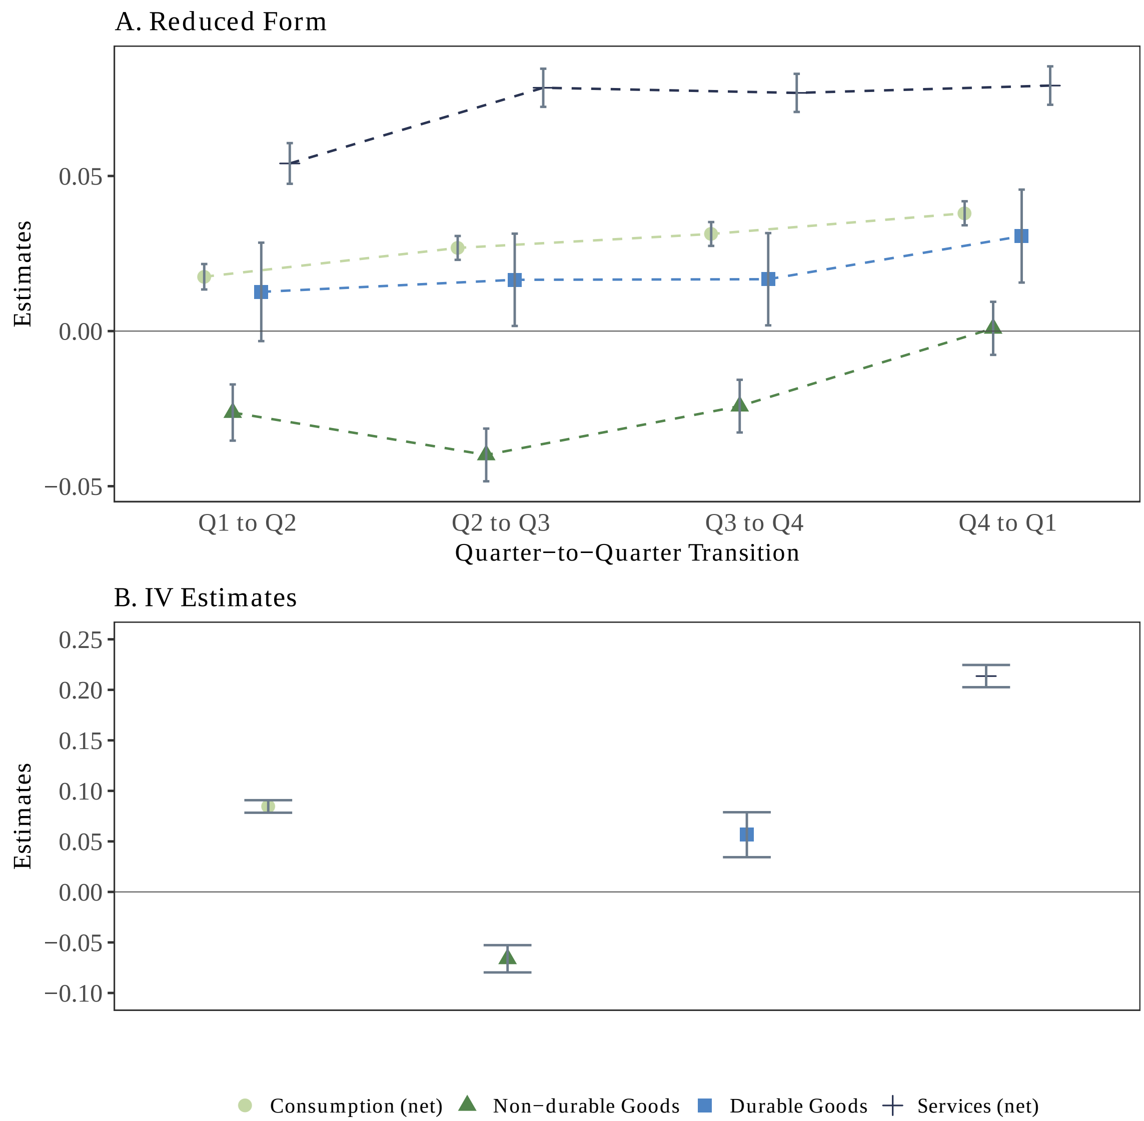

Figure 2-8: The Effect of Education Spending on Consumption Composition

Notes: This figure plots the reduced form and IV estimates of the effect of education spending on household consumption *net* of education spending. The sample is restricted to self-employed households with dependents aged between 15 and 19, where households with younger than 18 year olds are in the control group and those with 18 or 19 year olds are in the treated. Panel A plots the reduced form estimates β_t, which capture the effect of a dependent becoming college-entering age on household consumption for each quarter-to-quarter transition. Panel B plots the IV estimates – ρ – from equation 2.3, which capture the consumption effects induced by increased education spending. Groceries include spending at grocery stores or supermarkets. Restaurants include spending at restaurants, fast food chains, coffee shops, or bakeries. Medical expenses include dentist, doctor, or health practitioner visits. Utilities include cable, electric, gas, telecommunications, or water. Mortgage indicates mortgage payments. All regressions control for the age of the business and its owner, the number of dependents in a household, business industry, state of residence, and employer status of a business. The estimates are also reported in table 2.4. Whiskers show 95% confidence intervals and standard errors are clustered at the household-level.

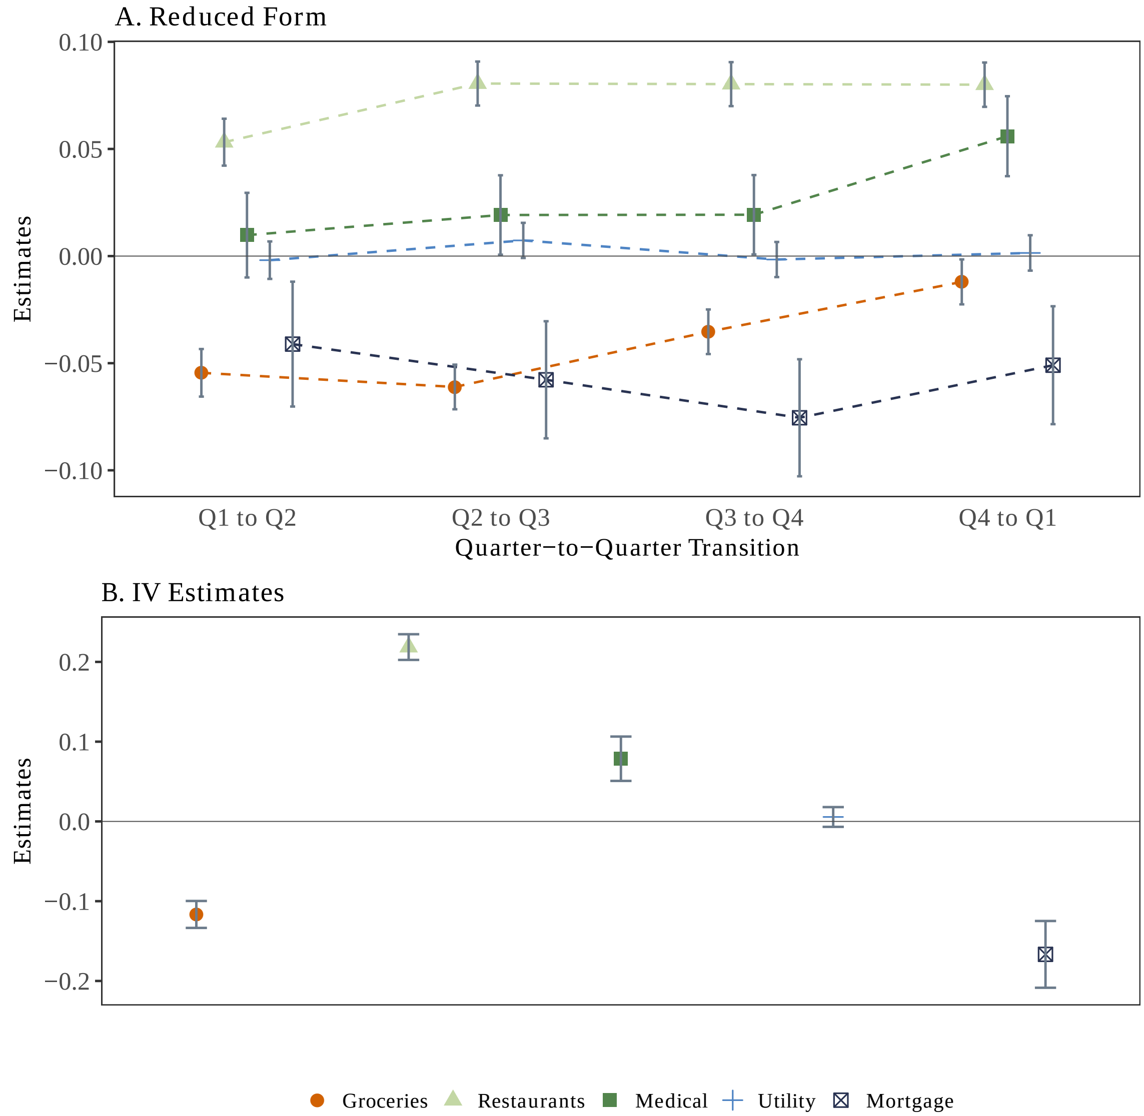

Figure 2-9: Business Effects by Average Growth Rates

Notes: This figure plots β_a of equation 2.4, which captures the effect of self-employed households having an a-year old dependent on business expenses and revenues by businesses with different business growth propensities. Business growth rates are calculated as the average year-over-year revenue growth *before* a child turns 18 years old. Self-employed households are grouped into quartile bins by average growth rates. Bin 1 includes firms with the lowest pre-18 average growth rates and bin 4 includes those with the highest pre-18 growth rates. The sample is restricted to self-employed households that ever had 18 or 19 year olds during the sample period I analyze. All regressions control for the age of the business and its owner, the number of dependents in a household, business industry, state of residence, and employer status of a business. Whiskers show 95% confidence intervals and standard errors are clustered at the household-level.

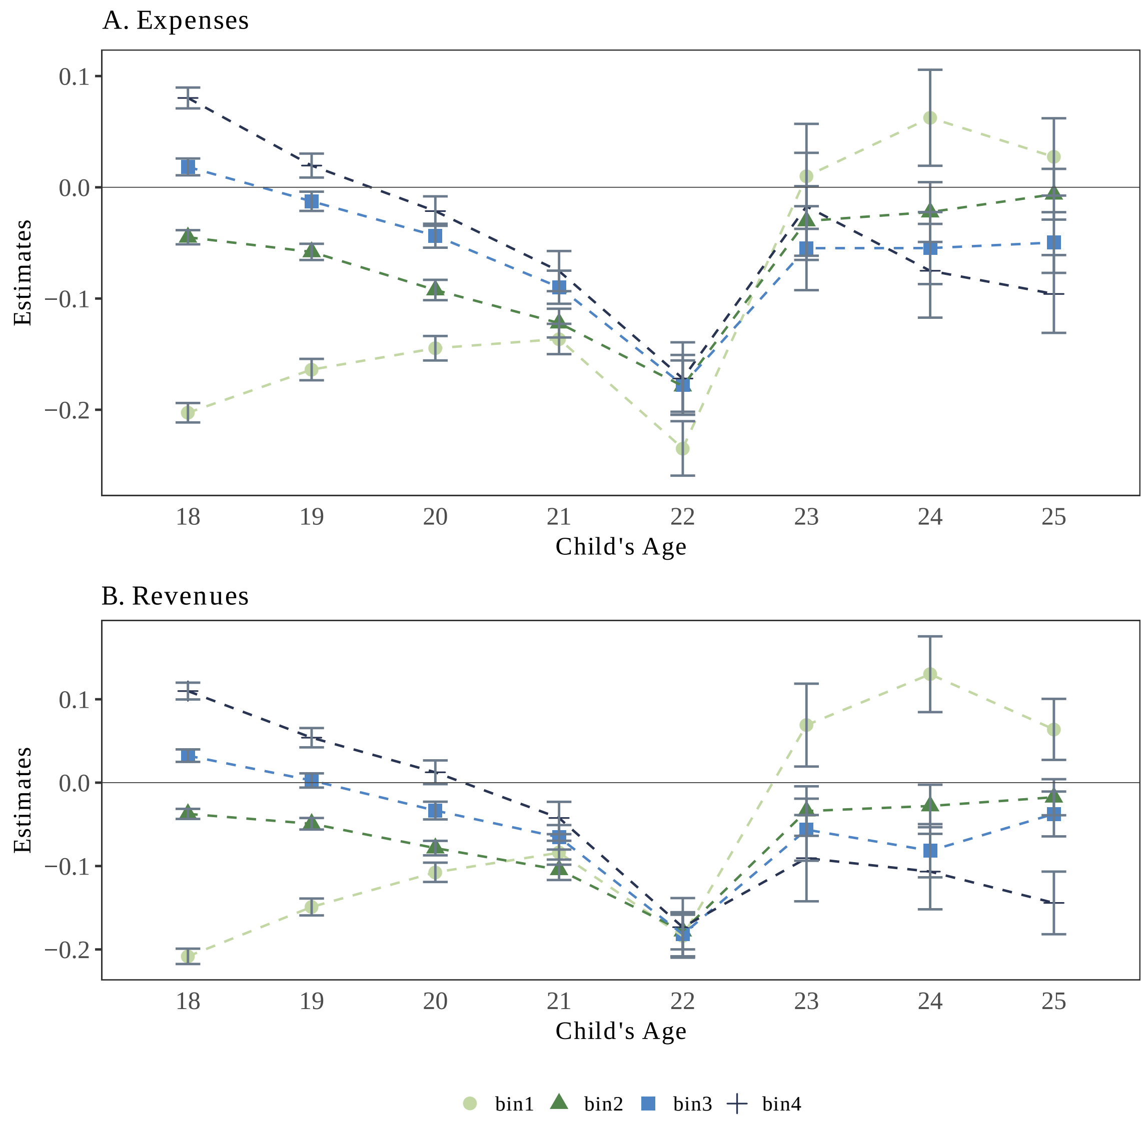

Figure 2-10: Household Consumption Effects by Average Growth Rates

Notes: This figure plots β_a of equation 2.4, which captures the effect of self-employed households having an a-year old dependent on household consumption *net* of education expenditure by businesses with different business growth propensities. Business growth rates are calculated as the average year-over-year revenue growth *before* a child turns 18 years old. Self-employed households are grouped into quartile bins by average growth rates. Bin 1 includes firms that has the lowest pre-18 average growth rates and bin 4 includes those with the highest pre-18 growth rates. The sample is restricted to self-employed households that ever had 18 or 19 year olds during the sample period. All regressions control for the age of the business and its owner, the number of dependents in a household, business industry, state of residence, and employer status of a business. Whiskers show 95% confidence intervals and standard errors are clustered at the household-level.

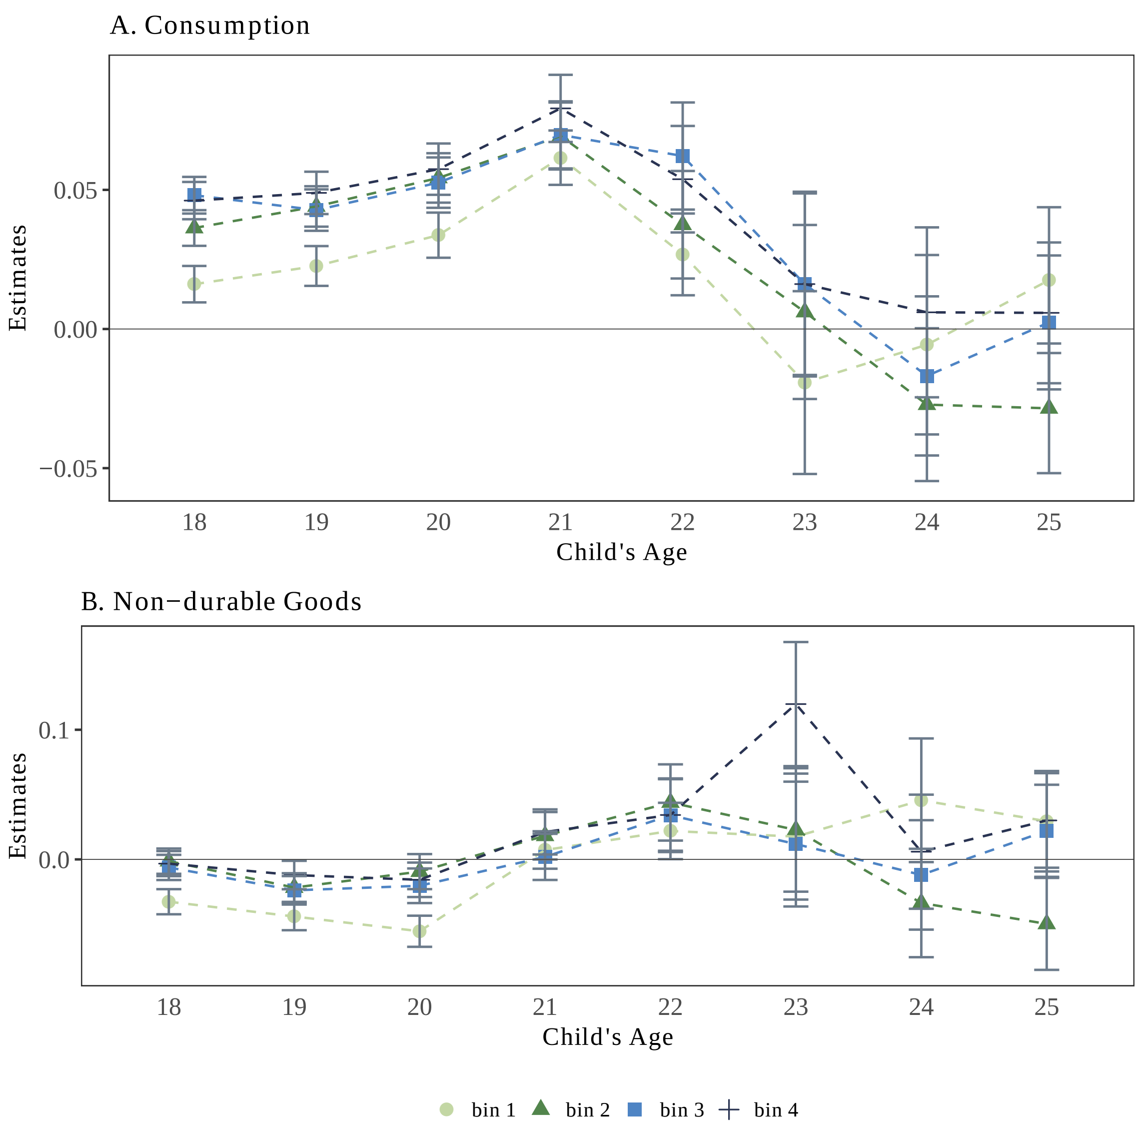

Figure 2-11: Heterogeneity in Post-Exit Response

Notes: This figure plots β_s from equation 2.5 by subgroups of households that exit before a child is 18 years old and those that exit after. The exit timing may capture the degree of financial constraint that a households experiences– households that exit before a child reaches college-entering age may be forward-looking households whereas those that exit after may have been induced to exit due to high spending burden. Consumption captures total household consumption net of education spending. Nondurable goods capture household spending on nondurable goods. Labor income refers to any direct deposits and payroll income from employer or payroll processor companies. Gig income refers to any income derived from participating in the online platform economy (labor, capital non-transport, and leasing platforms). Solid fitted lines are estimated from local regressions, and 95% confidence bands of the fitted lines are shown in grey. All regressions control for the age of the business owner, the number of dependents in a household, business industry and location, and employer status of a business.

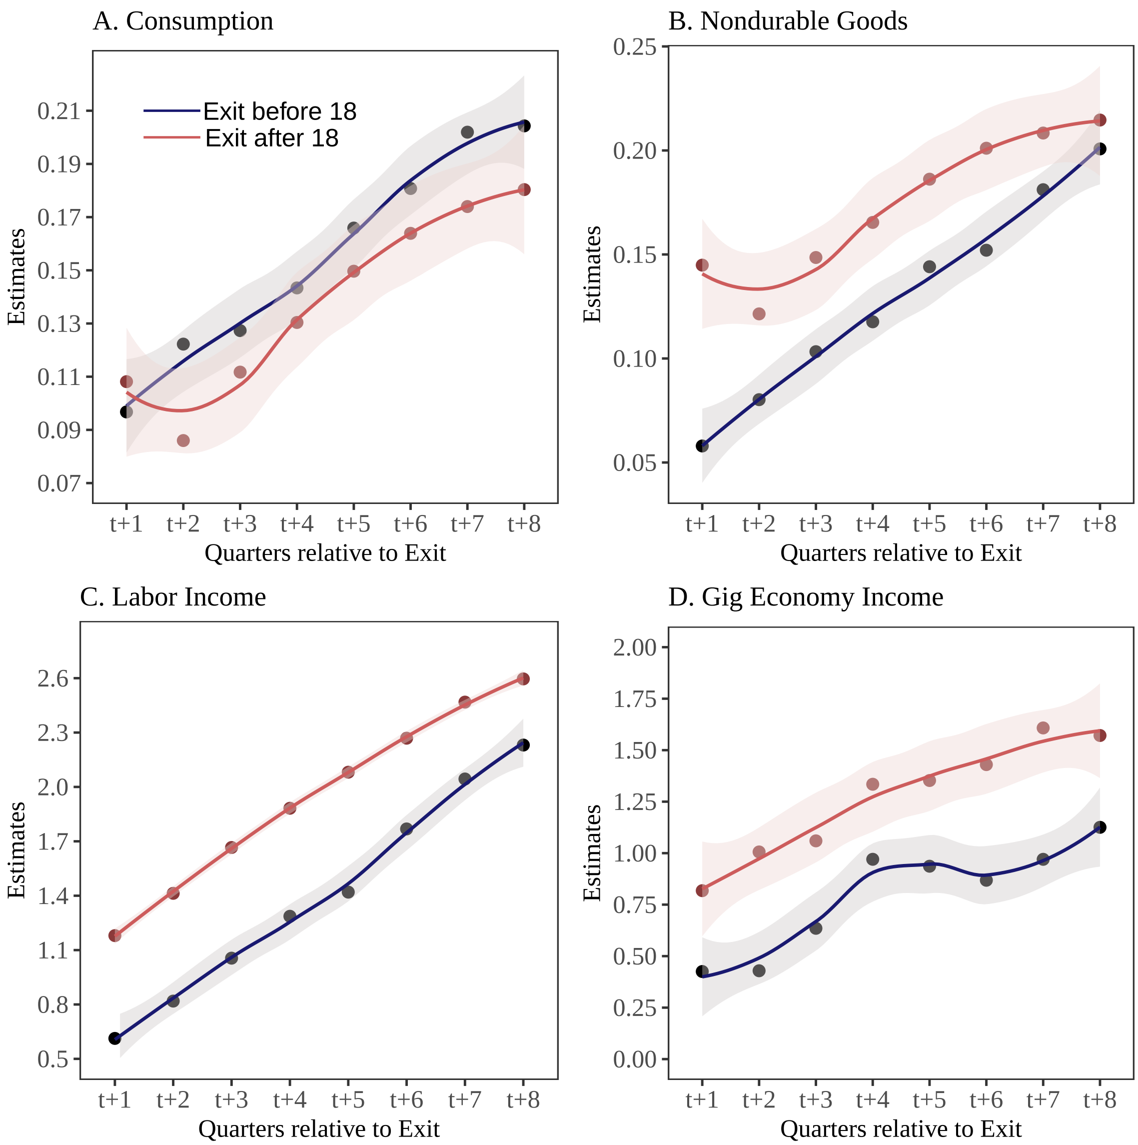

Table 2.1: Descriptive Statistics

Notes: This table reports descriptive statistics for household and business characteristics. Business outcomes are reported at the quarterly frequency in USD. To preserve anonymity, statistics are presented as means of ten observations in the p^{th} percentiles. The top panel reports descriptive statistics of all self-employed households with children of age between 14 and 25. The bottom panel restricts the sample size to the baseline regression sample with children of age between 15 and 19. Parent's age corresponds to the age of the oldest member in a household, and the child's age corresponds to the age of the oldest dependent in a household. Sample ranges from 2012 Q4 to 2018 Q2.

	Mean	SD	p25	p50	p75
	(1)	(2)	(3)	(4)	(5)
All Sample					
A. Household Characteristics					
Number of family members	3	1	3	3	4
Number of dependents	1	0.6	1	1	2
Parent's age	52	11	47	52	58
Child's age	21	3	18	21	23
B. Business Characterisristics					
Business years in operation	6	6	2	5	8
Business revenue	73,165	115,810	5,303	23,344	78,581
Business expense	62,472	103,755	4,029	16,967	63,700
C. Share of firms in top 3 industry and location					
Professional Services		0.16	California		0.18
Other Services (exc. Public Services)		0.12	New York		0.17
Construction		0.10	Texas		0.15
Number of households	148,275	148,275	148,275	148,275	148,275
Regression Sample					
A. Household Characteristics					
Number of family members	3	1	3	3	4
Number of dependents	1	0.5	1	1	1
Parent's age	51	10	45	49	55
Child's age	17	1.5	16	17	19
B. Business Characteristics					
Business years in operation	6	6	2	4	8
Business revenue	76,726	118,482	5,894	25,324	84,395
Business expense	65,278	106,097	4,383	18,260	67,986
C. Share of firms in top 3 industry and location					
Professional Services		0.17	California		0.17
Other Services (exc. Public Services)		0.11	Texas		0.16
Construction		0.11	New York		0.15
Number of households	61,684	61,684	61,684	61,684	61,684

Table 2.2: Quarterly Spending in Education

Notes: This table reports mean and standard deviation of quarterly education spending for self-employed households with children. Panels A and B compare statistics for households with 15-17 year olds ("Near College-Entering Sample") and those with 18-19 year olds ("College-Entering Sample"). Panels C and D compare statistics for households with 18-22 year olds ("College-Going Sample") and those with 23-25 year olds ("College-Graduating Sample"). Each panel reports unconditional and conditional statistics where unconditional statistics pool households with dependents of referenced age group, and conditional statistics pool those with positive education spending. Total education spending includes any payments to post-secondary institutions, testing service agencies, student loan servicing companies, and savings to 529 accounts. 529 drawdowns and savings are defined by transfers between 529 accounts and personal or business accounts.

Outcomes ($)	Mean (1)	SD (2)	Mean (3)	SD (4)	Mean (5)	SD (6)	Mean (7)	SD (8)
	A. Near College-Entering Sample				B.College-Entering Sample			
	Unconditional		Conditional		Unconditional		Conditional	
Total Education Spending	321	1,690	803	2,597	595	3,114	1,401	4,650
529 Drawdowns	5	509	11,037	20,075	26	847	11,103	13,523
529 Savings	26	572	2,199	4,763	26	676	2,590	6,207
Student Loan Payments	125	3,567	1,216	11,056	203	15,485	1,762	45,648
	C. College-Going Sample				D.College Graduating Sample			
	Unconditional		Conditional		Unconditional		Conditional	
Total Education Spending	617	2,951	1,588	4,563	356	3,052	1,311	5,741
529 Drawdowns	26	803	10,485	12,385	12	467	8,537	9,205
529 Savings	21	525	2,536	5,247	9	285	2,149	3,822
Student Loan Payments	253	22,825	1,971	63,699	255	2,003	1,480	4,634

Table 2.3: The Effect of Education Spending on Business Outcomes

Notes: This table reports first stage, reduced form, and 2SLS estimates of business response of sending kids to college using the college-entering sample that compares business outcomes of self-employed households with near college-entry age children (15-17 year-olds, or "control") to that of households with college-going age (18-19 year-olds, or "treatment") children. Column 1 reports control mean in levels. Columns 2 - 5 report β_t from respective first stage or reduced form regressions. The first row reports the estimates from the first stage equation 2.2, and the bottom rows report the reduced form estimates. Column 6 reports the TSLS-IV estimate ρ from equation 2.3. All regressions control for the age of the business and its owner, the number of dependents in a household, business industry, state of residence of a household, and employer status of a business. Standard errors are clustered at the household-level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.10

Outcome Variables	Control Mean	Quarter Transitions $\times$ $\mathbb{1}$(College-entering age)				TSLS Estimates
		I to II	II to III	III to IV	IV to I	
	(1)	(2)	(3)	(4)	(5)	(6)
			First Stage			
Education Spending	$354	0.248 ***	0.409 ***	0.335 ***	0.346 ***	
		(.013)	(.012)	(.012)	(.012)	
			Reduced Form			
Business Expenses	$67,822	-0.015 ***	-0.023 ***	-0.032 ***	-0.035 ***	-0.078 ***
		(.003)	(.003)	(.003)	(.003)	(.005)
Auto maintenance	$155	-0.011	-0.015 *	-0.033 ***	0.010	-0.036 ***
		(.009)	(.008)	(.008)	(.008)	(.012)
Office supplies/tools	$735	-0.047 ***	-0.053 ***	-0.062 ***	-0.059 ***	-0.171 ***
		(.009)	(.009)	(.009)	(.009)	(.016)
Machinery	$375	-0.037 ***	-0.053 ***	-0.062 ***	-0.030 ***	-0.134 ***
		(.01)	(.009)	(.009)	(.009)	(.015)
Utilities	$854	-0.002	-0.007	-0.013	0.003	-0.015
		(.009)	(.009)	(.009)	(.009)	(.013)
Business Revenues	$80,557	-0.015 ***	-0.020 ***	-0.028 ***	-0.035 ***	-0.073 ***
		(.004)	(.004)	(.004)	(.004)	(.006)
Exit	0.009	0.001	0.002 ***	0.002 ***	-0.001	0.003 ***
		(.001)	(.001)	(.001)	(.001)	(.001)
Number of Observations	382,063	382,063	382,063	382,063	382,063	382,063

Table 2.4: The Effect of Education Spending on Household Consumption

Notes: This table reports first stage, reduced form, and 2SLS estimates of household spending using the college-entering sample that compares business outcomes of self-employed households with near college-entry age children (15-17 year olds, or "control") to that of households with college-going age (18-19 year olds, or "treatment") children. Household consumption Standard errors are clustered at the household-level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.10

Outcome Variables	Control Mean	Quarter Transitions $\times$ $\mathbb{1}$(College-going age)				TSLS Estimates
		I to II	II to III	III to IV	IV to I	
	(1)	(2)	(3)	(4)	(5)	(6)
		First Stage				
Education Spending	$354	0.248***	0.409***	0.335***	0.346	
		(.013)	(.012)	(.012)	(.012)	
		Reduced Form				
Consumption	$17,462	0.018***	0.027***	0.031***	0.038***	0.085***
		(.002)	(.002)	(.002)	(.002)	(.003)
Nondurable goods	$2,435	-0.026***	-0.040***	-0.024***	0.001	-0.066***
		(.005)	(.004)	(.004)	(.004)	(.007)
Groceries	$967	-0.054***	-0.061***	-0.035***	-0.012**	-0.117***
		(.006)	(.005)	(.005)	(.005)	(.009)
Durable goods	$615	0.013	0.017**	0.017**	0.031***	0.057***
		(.008)	(.008)	(.008)	(.008)	(.011)
Services	$2,886	0.054***	0.078***	0.077***	0.079***	0.214***
		(.003)	(.003)	(.003)	(.003)	(.006)
Restaurant	$656	0.053***	0.081**	0.080***	0.080***	0.219***
		(.006)	(.005)	(.005)	(.005)	(.01)
Medical	$183	0.010	0.019**	0.019**	0.056***	0.079***
		(.01)	(.009)	(.009)	(.009)	(.014)
Utility	$1,183	-0.002	0.007*	-0.002	0.001	0.006
		(.004)	(.004)	(.004)	(.004)	(.006)
Mortgage	$2,270	-0.041***	-0.058***	-0.075***	-0.051***	-0.167***
		(.015)	(.014)	(.014)	(.014)	(.021)
Number of Observations	382,063	382,063	382,063	382,063	382,063	382,063

Table 2.5: Business Effects by Exit Decisions

Notes: This table reports the business spending response over a dependent's age profile by subsample of firms that do and don't exit from self-employement. The left panel ("Stay") sample contains households that remain self-employed, and the right panel ("Exit") sample contains households that exit from self-employment at some point during my sample period. Thus, the outcomes for the "exit" sample reflect the business response before the owners exit. Standard errors are clustered at the household level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.10

	Stay				Exit			
Child's Age	Education Spending	Business Expenses	Business Revenues	Machinery	Education Spending	Business Expenses	Business Revenues	Machinery
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
18	0.196 ***	-0.010 ***	-0.006 ***	0.017 ***	0.131 ***	-0.050 ***	-0.062 ***	-0.047 ***
	(.007)	(.002)	(.002)	(.004)	(.020)	(.007)	(.007)	(.011)
19	0.155 ***	-0.013 ***	-0.008 ***	0.007 **	0.073 ***	-0.058 ***	-0.068 ***	-0.056 ***
	(.006)	(.002)	(.002)	(.003)	(.019)	(.007)	(.007)	(.011)
20	0.106 ***	-0.015 ***	-0.008 ***	-0.001	-0.005	-0.062 ***	-0.071 ***	-0.058 ***
	(.006)	(.002)	(.002)	(.003)	(.018)	(.007)	(.007)	(.011)
21	0.090 ***	-0.018 ***	-0.011 ***	0.002	0.005	-0.056 ***	-0.068 ***	-0.061 ***
	(.006)	(.002)	(.002)	(.003)	(.018)	(.007)	(.007)	(.011)
22	0.039 ***	-0.026 ***	-0.019 ***	-0.007 **	-0.070 ***	-0.063 ***	-0.075 ***	-0.061 ***
	(.006)	(.002)	(.002)	(.003)	(.018)	(.007)	(.007)	(.011)
23	-0.024 ***	-0.023 ***	-0.017 ***	-0.010 ***	-0.097 ***	-0.074 ***	-0.091 ***	-0.064 ***
	(.006)	(.002)	(.002)	(.003)	(.017)	(.007)	(.007)	(.010)
24	-0.087 ***	-0.026 ***	-0.014 ***	-0.015 ***	-0.133 ***	-0.082 ***	-0.095 ***	-0.069 ***
	(.006)	(.002)	(.002)	(.003)	(.017)	(.007)	(.007)	(.010)
25	-0.139 ***	-0.029 ***	-0.018 ***	-0.015 ***	-0.178 ***	-0.102 ***	-0.106 ***	-0.097 ***
	(.005)	(.002)	(.002)	(.003)	(.015)	(.006)	(.006)	(.010)
Number of Observations	1,557,267	1,481,397	1,481,397	1,557,267	167,762	167,733	167,733	167,762

Table 2.6: Consumption Effects by Exit Decisions

Notes: This table reports the consumption response over a dependent's age profile by subsample of firms that do and don't exit from self-employement. The left panel ("Stay") sample contains households that remain self-employed, and the right panel ("Exit") sample contains households that exit from self-employment at some point during my sample period. Thus, the outcomes for the "exit" sample reflect the business response before the owners exit. Standard errors are clustered at the household level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.10

	Stay				Exit			
Child's Age	Consumption	Groceries	Restaurants	Medical	Consumption	Groceries	Restaurants	Medical
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
18	0.021***	0.004*	0.038***	0.010***	0.014***	-0.013*	0.015**	-0.002
	(.001)	(.002)	(.002)	(.004)	(.004)	(.007)	(.007)	(.012)
19	0.018***	0.001	0.038***	0.003	0.000	-0.027***	0.002	-0.008
	(.001)	(.002)	(.002)	(.004)	(.004)	(.007)	(.007)	(.012)
20	0.016***	0.000	0.037***	-0.008**	-0.002	-0.030***	-0.004	-0.056***
	(.001)	(.002)	(.002)	(.004)	(.004)	(.007)	(.007)	(.011)
21	0.021***	0.006***	0.046***	-0.009**	-0.005	-0.024***	0.007	-0.045***
	(.001)	(.002)	(.002)	(.004)	(.004)	(.007)	(.007)	(.011)
22	0.020***	0.001	0.048***	-0.019***	-0.011***	-0.035***	-0.005	-0.035***
	(.001)	(.002)	(.002)	(.004)	(.004)	(.007)	(.007)	(.011)
23	0.014***	-0.009***	0.041***	-0.026***	-0.004	-0.034***	0.001	-0.064***
	(.001)	(.002)	(.002)	(.004)	(.004)	(.007)	(.007)	(.011)
24	0.013***	-0.017***	0.036***	-0.038***	-0.005	-0.045***	-0.006	-0.061***
	(.001)	(.002)	(.002)	(.004)	(.004)	(.007)	(.006)	(.011)
25	0.020***	-0.018***	0.040***	-0.035***	0.001	-0.045***	-0.004	-0.065***
	(.001)	(.002)	(.002)	(.003)	(.004)	(.006)	(.006)	(.010)
Number of Observations	1,557,267	1,557,267	1,557,267	1,557,244	167,762	167,762	167,762	167,762

Table 2.7: Post-Exit Consumption and Income Paths

Notes: This table reports consumption and income paths of households that exit from self-employment using the reduced form equation 2.5. The table reports β_s, which captures the average effect of post-exit response s periods after exiting. Consumption measure captures total household consumption net of education expenditure. Labor income refers to any direct deposits and payroll income from employer or payroll processor companies. Gig income refers to any income derived from participating in the online platform economy (labor, capital non-transport, and leasing platforms). UI receipt can be identified from UI inflows. Standard errors are clustered at the household level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.10

	Outcomes Relative to Pre-Exit Mean							
Quarters from Exit	Consumption		Labor Income		Gig Income		UI Receipt	
	(1)		(2)		(3)		(4)	
$t+1$	0.110	***	1.122	***	0.776	***	0.000	
	(.004)		(.029)		(.049)		(.002)	
$t+2$	0.093	***	1.345	***	0.939	***	0.000	
	(.005)		(.033)		(.055)		(.002)	
$t+3$	0.117	***	1.585	***	1.003	***	-0.003	
	(.005)		(.037)		(.059)		(.002)	
$t+4$	0.135	***	1.793	***	1.279	***	-0.004	*
	(.005)		(.041)		(.070)		(.003)	
$t+5$	0.156	***	1.973	***	1.284	***	0.000	
	(.006)		(.045)		(.074)		(.003)	
$t+6$	0.171	***	2.178	***	1.329		0.000	
	(.006)		(.050)		(.080)		(.003)	
$t+7$	0.183	***	2.383	***	1.485	***	-0.001	*
	(.007)		(.054)		(.088)		(.003)	
$t+8$	0.189	***	2.515	***	1.477	***	-0.005	*
	(.007)		(.054)		(.093)		(.003)	
Number of Observations	298,408		298,408		298,408		298,408	

Chapter 3

Loan Guarantees and Credit Supply

3.1 Introduction

Indirect government loan guarantees reimburse unrecovered dollars to private lenders and are an increasingly common type of credit subsidy. In 2019 alone, \$1.4 out of the \$1.5 trillion dollars in projected federal credit assistance came in the form of loan guarantees, with a projected subsidy value of \$37.9 billion.[1] This paper studies how private lenders respond to federal loan guarantees. In markets affected by asymmetric information and credit rationing, government loan guarantees can increase aggregate welfare if they restore lending to an efficient level (see, e.g., Gale, 1991; Stiglitz and Weiss, 1981; Smith, 1983; Mankiw, 1986). Whether this occurs is ultimately an empirical question and depends in part on the responsiveness of lenders to the guarantee. Whether federal guarantee programs have any effects on increasing access to credit, or simply act as a subsidy to lenders, depends on the elasticity of credit provision to the loan guarantee. If credit supply is inelastic, guarantees will not increase the level of borrowing and will simply reimburse lenders on their losses. In this case, government loan guarantees can also crowd out more efficient private borrowing and encourage excessive risk-taking. Despite the large and growing volume of federally guaranteed debt, there remains relatively little work exploring the effects of federal guarantees on lending.

[1] See the report "Fair-value estimates of the cost of federal credit programs in 2019" by the Congressional Budget Office Research (https://www.cbo.gov/publication/55278).

In this paper, we focus on how guarantees affect the supply of credit to small businesses. Credit constraints are well-known barriers to growth for small firms, and these problems are especially severe given imperfect information and a lack of collateral (see, e.g., Fazzari, Hubbard, Petersen, Blinder and Poterba, 1988; Petersen and Rajan, 1994, 1995; Kaplan and Zingales, 1997; Barrot, 2016). Prior work has shown that these programs can alleviate barriers to entrepreneurship (Clair Lelarge, David Sraer and David Thesmar, 2010). We employ data from the Small Business Administration (SBA), the government agency tasked with providing assistance to small businesses. Specifically, we utilize data on loans originated under the 7(a) loan program. Under the SBA 7(a) loan program, a portion of loans from commercial lenders are insured against losses from defaults. Loans of up to $150,000 carry a higher maximum guarantee rate than loans larger than $150,000. This feature of the federal guarantee program leads to sharply different levels of risks for lenders originating loans above and below the threshold.

We employ a bunching estimator to measure the excess mass at the threshold and use this to estimate the elasticity of loan supply to the guarantee rate. We use a simple model to translate the observed excess borrowing at the mass into an elasticity of credit supply. The degree of bunching identifies the elasticity of lending supply to the guarantee; if lending supply is inelastic and lenders do not adjust loan size in response to the guarantee, we will not observe bunching. On the other hand, if lending supply is highly elastic, we will observe bunching, as a significant number of loans will be moved to the side of the threshold with higher guarantees.

We find significant bunching directly below the threshold, which translates to a highly elastic lending supply response to loan guarantees. Interpreted in dollar magnitudes, this means that a one percentage point change in the guarantee net subsidy rate (expressed as a percentage of loan principal) generates $19,054 in additional lending. Guarantee thresholds change over time, and we find that the observed bunching is stronger in years when guarantee amounts across the threshold are higher. We find that the elasticity varies slightly from year to year, and consistent with optimization frictions, we find smaller elasticities in years immediately after guarantee notches have changed. Moreover, the guarantee notch was eliminated during a two-year period from 2009 to 2010, as part of the American Recovery and Reconstruction Act (ARRA). During this period, we find no excess mass across the threshold, which serves as a placebo test to rule out the possibility that alternative factors might be changing across the threshold and driving our results.

The validity of the bunching estimate relies on two key assumptions: first, the counterfactual distribution is smooth in the absence of a notch, and second, there exists a well defined marginal buncher. Consistent with our identifying assumptions, we find no excess mass in years when the guarantee notch is eliminated, making it unlikely that other factors are changing at the threshold. Additionally, we find no differences in loan terms around the threshold: interest rates, maturities, revolving loan percentages, and charge-off percentages appear similar at or near the notch. We rule out several alternative explanations and threats to identification. According to SBA rules, lenders are only able to issue one loan to borrowers that have exhausted other borrowing options. We confirm in the data that lenders are not issuing multiple loans to the same borrower to take advantage of guarantees. We also find no difference in interest rates at or around the threshold, which is likely due to a particular institutional detail–the majority of loans in this program have binding interest rate caps, and thus there is very little room to vary the interest rate. This supports a channel from distributional responses driven by supply, rather than demand, side forces.

Our analysis sheds light on an ongoing policy debate regarding the efficiency of government loan guarantees. Proponents of lending guarantee programs argue that guarantees provide credit to borrowers that would otherwise be unable to access funds. However, opponents of loan guarantee programs contend that these programs simply serve as a subsidy to lenders. Major pushes to shut down the SBA were undertaken by the executive branch and Congress in 1984 and 1996, with pressure continuing into the 2010s. For example, a 2012 Wall Street Journal op-ed noted that "Congress created the Small Business Administration in 1953 to fix a specific problem: Lenders allegedly were turning away large numbers of small businesses that, if given a loan, would generate untapped economic growth. It is questionable whether this problem ever existed... The SBA loan program is best understood as a subsidy to banks. Borrowers apply to an SBA-certified bank. The SBA guarantees 75% to 85% of the value of loans made in the flagship program. The banks then boost their earnings by selling the risk-free portion of the loans on a secondary market." See the Congressional Budget Office (CBO) for a discussion of proposals to eliminate the SBA. As well as being important to policy, the effect of loan guarantees on the supply of funds is a key parameter in many models of the effects of guarantees. For example, Smith (1983) notes that "To be effective, it must be demonstrated that there is some impact of these policies on supply elasticities of credit." Gale (1991) states that "Perhaps the single most important and controversial parameter is the elasticity of supply of funds." Finally, Lucas (2016) notes that "The elasticity of credit supply affects the extent to which additional borrowing in government credit programs is offset by reductions in private borrowing."

We inform this debate by focusing on the guarantee program that serves as a major source of small business financing in the United States. The SBA is an important source of small business financing, with $25.4 billion in SBA-guaranteed loans made in 2018, mainly through the 7(a) program. This funding is typically provided to young firms at a critical point in the firm's life cycle when they are unable to access other sources of capital. A number of well-known major companies secured SBA loans in early stages. These include Apple, FedEx, Nike, Intel, Under Armour, Whole Foods, and Chipotle.

This paper contributes to a body of work on federal lending subsidies and guarantees by estimating a key parameter from classic theory models. Despite the growing volume of federal lending in recent years, the area remains underexplored relative to other credit markets. Notable exceptions include Gale (1990), Gale (1991), Smith (1983), and Lucas (2016). To our knowledge, this is the first empirical paper to estimate how lending supply responds to federal loan guarantees. This literature largely focuses on calibrated models, and different papers use a wide range of estimates of the elasticity of credit supply to guarantee rates for calibrations.

Other work has focused on different aspects of government credit guarantees. Rafael La Porta, Florencio Lopez de Silanes and Andrei Shleifer (2002) examine the effect of government ownership of banks, and find a positive correlation between government intervention and slower subsequent financial development that is consistent with government crowding out efficient private borrowing. Marianne Bertrand, Antoinette Schoar and David Thesmar (2007) examine the effect of the French Banking Act of 1985, which eliminated government subsidies to banks intended to help small- and medium-sized firms. Clair Lelarge, David Sraer and David Thesmar (2010) study the effects of a French guarantee program on entrepreneurship. Atkeson, d'Avernas, Eisfeldt and Weill (2018) emphasize the role of government guarantees in bank valuation by arguing that the decline in banks' market-to-book ratio since the 2008 crisis is due to changes in the value of government guarantees. Bryan Kelly, Hanno Lustig and Stijn Van Nieuwerburgh (2016) show that government guarantees lower financial sector index prices.

Prior theory work has shown that under information asymmetries, government interventions in credit markets such as loan guarantees and loan subsidies can increase welfare (see, e.g., Stiglitz and Weiss, 1981; Mankiw, 1986; Greenwald and Stiglitz, 1986). More recent work by Hanson, Scharfstein and Sunderam (2018) has focused on tradeoffs between private and social costs, and Fieldhouse (2018) shows that housing policies subsidizing an expansion in residential mortgage lending crowd out commercial mortgages and loans. While in theory loan guarantees can increase welfare, whether this is true in practice is ultimately an empirical question. We show that private lending is indeed responsive to federal loans guarantees, suggesting that these programs have real effects beyond simply subsidizing lenders.

This paper also links to a literature on credit access for entrepreneurs and small firms. Financing constraints are well known to be a significant barrier to growth for small firms (see, e.g., David S. Evans and Boyan Jovanovic, 1989; Toni Whited and Guojun Wu, 2006; Joshua Rauh, 2006; William R. Kerr and Ramana Nanda, 2010; Jean-Noel Barrot, 2016; Manuel Adelino, Song Ma and David Robinson, 2017). A large body of work studies small enterprises' financial frictions and various policy responses. Petersen and Rajan (1994), Petersen and Rajan (1995), and Darmouni (2017) show that, for small firms, close ties with institutional lenders increases the availability of credit. Darmouni and Sutherland (2018) show that lenders to small firms are highly responsive to competitors' offers.

More recent work has focused on how federal programs can affect the supply of credit and entrepreneurship. Brown and Earle (2017) and Joao Granja, Christian Leuz and Raghuram Rajan (2018) study the SBA program and, respectively, find that access to credit has large effects on employment and the average physical distance of borrowers from banks' branch matters for ex-post loan performance. Jean-Noel Barrot, Thorsten Martin, Julien Sauvagnat and Boris Vallee

(2019), Mullins and Toro (2017), and Gonzalez-Uribe and Wang (2019) study similar programs to stimulate small business lending in France, Chile, and the UK. Howell (2017) demonstrates that federal grants have large effects on future fundraising, patenting and revenue. This paper shows that the volume of small business lending is highly responsive to loan guarantees and loan guarantees can be a relatively low cost way to increase lending to small enterprises.

Beyond the use of these estimates directly for the growing literature on loan guarantees, our estimates and their implications for the supply of credit to small businesses are relevant for structural models of entrepreneurship and firm dynamics. For example, David S. Evans and Boyan Jovanovic (1989) assume that the lending rate equals the borrowing rate, which implies that the supply curve of capital is not upward sloping over a wide range. Neus Herranz, Stefan Krasa and Anne Villamil (2015) additionally assume that debt is provided by a risk-neutral competitive lender with an elastic supply of funds. Our estimates are also of use in terms of estimating the marginal value of public funds (e.g Hendren (2014, 2016)), specifically in terms of the SBA program for welfare analysis. The marginal value of public funds maps the causal estimates of a policy change into welfare analysis by comparing the ratio of the beneficiaries' willingness to pay for a program with the net cost to government, in other words, cost-benefit analysis.

The remainder of this paper is organized as follows. Section 3.2 discusses institutional details on SBA loans and federal guarantees and describes the SBA data used in our analysis. Section 3.3 presents an illustrative model and discusses how our identification strategy is linked to this model. Section 3.4 introduces the bunching estimator and discusses the empirical approach. Section 3.5 presents the main results and demonstrates significant lending response to government guarantees. Section 3.6 discusses alternative explanations and presents placebo results. Section 3.7 concludes the paper and discusses avenues for further research.

3.2 Institutional background and data

Federal loan guarantee programs operate in a fashion similar to insurance contracts. Lenders pay a fee to the government, and in return, the government reimburses a portion of dollars that are charged off when a loan goes into default. Loan guarantees exist or have existed in several loan markets, such as in student, mortgage, and small business lending markets. In this paper, we focus on loan guarantees in the small business lending market. This section discusses the institutional details surrounding the SBA 7(a) program studied in our empirical analysis.

3.2.1 SBA 7(a) loans

The SBA is an independent federal government agency created in 1953 with the mission of providing assistance to small businesses. We focus on the lending program, designed to improve access to capital for young small businesses that may not be eligible to obtain credit through traditional lending channels. The SBA lending programs are guarantee programs where the SBA guarantees a portion of loans originated by commercial lending institutions against losses from defaults rather than lending directly to qualifying borrowers. We focus on the SBA's flagship loan guarantee program, the 7(a) loan program.

SBA 7(a) guarantees consist of two components, a reimbursement rate and a fee. The reimbursement rate is the fraction of each dollar charged off that the bank receives back from the SBA, and the fee is the amount that the bank must pay to participate in the 7(a) program. There are several features of guarantee components that are relevant to this study. Most importantly, the maximum guarantee rate is based on a nonlinear size cutoff rule: loans up to $150,000 carry a maximum guarantee rate of 85%, which drops sharply to 75% for loans larger than $150,000. The guarantee fees also increase at the same threshold, making the overall guarantee less generous for loans larger than $150,000. We exploit this guarantee notch around $150,000 to identify our parameters of interest. Features of the SBA 7(a) program have remained relatively stable over the last decade, except during 2009-2010, when the SBA temporarily raised the guarantee rate on either side of the $150,000 threshold to 90% and waived fees with the signing of the ARRA of 2009. This time period provides a helpful placebo test for our analysis since no lending response should occur in a year when there is no discrete change in the guarantee rate.

To qualify for a 7(a) loan, a borrower must meet several requirements. First, it must be a for-profit business that meets SBA size standards. Size standards vary by industry and are based on the number of employees or the amount of annual receipts ("total income" plus "the costs of goods sold"). In addition to the size requirement, a business must be independently owned and operated and not be nationally dominant in its field. The business must also be physically located and operate in the US or its territories. Last, small businesses must demonstrate the need for a loan by providing loan application history, business financial statements, and evidence of personal equity investment in the loan proposal.

To qualify, borrowers must exhaust other funding sources, including personal sources, before seeking financial assistance and be willing to pledge collateral for the loan.[2] SBA 7(a) loans are intended as a last resort, and to ascertain that borrowers cannot access credit elsewhere, lenders are required to conduct "credit elsewhere" tests. The SBA provides further information regarding credit elsewhere tests. In addition, Appendix Table C.3 shows the fraction of firms accessing multiple sources of credit in the 2003 Federal Reserve Survey of Small Business Finances (SSBF) that have loans from a government agency, including the SBA. The table indicates that very few firms that have SBA loans are accessing credit from multiple sources. Lenders are required to demonstrate that borrowers cannot obtain the loan on reasonable terms without the SBA guarantee and that the funds are not unavailable from the resources of the applicant. The personal resources of any applicant who owns more than 20% of the small business are reviewed. The SBA monitors lenders' compliance with the credit elsewhere test through targeted reviews. Failure to comply with credit elsewhere tests can lead to the denial of a guarantee, exclusion from the lending program and other enforcement actions from the Office of Credit Risk Management.

The 7(a) loans are disbursed through private lending institutions. This loan submission and disbursement procedure depends largely on the lender's level of authority (i.e., delegated or non-delegated) provided by the SBA. The SBA conducts its own analysis of the application and approves the originating lender's decision to lend, which can be expedited depending on a lender's experience. In practice, SBA lenders have meaningful bargaining power over credit supply. In a typical case, a borrower requests a loan to a lender, and the lender decides whether the SBA loan would be suitable for a given borrower upon reviewing the borrower's background. Given that

[2]The following reports contain detailed information on the eligibility requirements: "Small Business Administration 7(a) loan guaranty program" by the Congressional Research Service; "Bankers' guide to the SBA 7(a) loan guaranty program" by the Office of Comptroller of the Currency Community Affairs Department; "Lender and development company loan programs. SOP 50 10 5(H)" by the SBA.

lenders cannot provide more than one loan to a single borrower such that the SBA-guaranteed loan is secured with a junior lien position, lenders have incentives to retain this bargaining power and to be selective in choosing borrowers.

Note that the reimbursement rate and fees are typically determined by an Office of Management and Budget (OMB) model, vary from year to year, and have been changed through legislation such as the ARRA. The CBO notes that "One of the SBA's goals is to achieve a zero subsidy rate for its loan guaranty programs," which entails generating revenue from fees and recoveries to offset program costs. In practice, the SBA is sometimes successful and sometimes not in terms of achieving a zero subsidy rate. Between 2007 and 2009, and between 2014 to the present day, the program operated at zero subsidy. The CBO report on the SBA 7(a) loan guaranty program provides further detail regarding the goals and subsidy rates of the program.

3.2.2 Data

We obtain the 7(a) loan data from the SBA. The SBA requires all participating lenders in the 7(a) program to submit loan applications (Forms 1919 and 1920) to the 7(a) Loan guaranty Processing Center ("LGPC") when they request a new loan. Delegated lenders must complete the form, sign and date, and retain in their loan file before processing a loan for faster processing. The information included in these forms are then compiled into a data set and are provided publicly pursuant to the Freedom of Information Act (FOIA). This loan origination data set includes basic information about the participants (i.e., the identity of the borrower and the lender and their addresses, city, zip code, and industry), nonpricing terms (i.e., loan volume, guarantee amount, or approval date), pricing term (i.e., loan spread plus base rate), ex-post loan performance (such as the total loan balance that has been charged off), and other administrative details such as the delegation status of the lender and the SBA district office that processed the loans.

For our analysis, we only consider loans originated over the last decade—2008 to 2017—under the SBA 7(a) program. We exclude SBA 7(a) Express loans and drop 22 loans that appear to contain data errors (i.e., loans for which the guaranteed share is greater than a hundred percent of the amount originated). Under these restrictions, the sample covers 199,013 loans originated by

3,066 lenders to 177,049 borrowers. Table 3.1 presents summary statistics for the main analysis variables.

The median SBA loan size is \$460,000, and the guaranteed amount is \$356,400. The median loan maturity and interest rate at the time of origination are ten years and 6%, respectively. Since the median prime rate is 3.25% in our sample, the maturity and interest rates are consistent with the SBA's maximum interest rate rule. Loans with a maturity of over seven years and amount greater than \$50,000 can carry a maximum rate of 2.75% over the prime rate. The median charge-off amount is zero, while the mean is \$11,706, indicating that the share of loans that are eventually charged off is small. Panel B of Table 3.1 reports the same statistics for subsample of loans used for notch estimation, where we restrict the loan size to be between \$75,000 and \$225,000. We restrict to the left of the threshold to loans above \$75,000 to avoid the excluded region from a second interest rate notch. Loans below \$50,000 carry a higher interest rate cap, which can additionally change lender incentives and lead to bunching. We take an equal range to the right of the \$150,000 threshold to arrive at the upper bound, \$225,000. Once we apply this restriction, we include 41,460 loans in the main analysis sample.

While the distribution is relatively similar to that in other papers using SBA data, such as Brown and Earle (2017), we only include 7(a) loans between 2008 and 2017. The difference in means relative to Brown and Earle (2017) comes from the fact that they include 504 loans that are up to \$5.5 million, whereas we only examine loans below \$350,000 in our main analysis sample. For certain heterogeneity analysis, we also link our main data to the Federal Deposit Insurance Corporation (FDIC) Statistics on Depository Institutions Data. This data set and sample construction is discussed in Appendix Section C.3. Additional robustness checks vary the main analysis sample to include some loans from the sample shown in panel A.

We use these data to estimate private lenders' responsiveness to federal loan guarantees. It is important to note that lenders cannot manipulate the lending structure by issuing multiple guaranteed loans to the same borrower. As discussed in the institutional details section, the SBA prohibits lenders from originating loans with a "piggyback" structure where multiple loans are issued to the same borrower at the same time, and the guaranteed loan is secured with a junior lien position.

While this policy does not prevent lenders from having a shared lien position with the SBA loans (i.e., Pari Passu), we confirm in our data that more than 99% of the borrowers receive only one loan from the same lender at the same time. As reported in Table 3.1, the average number of loans a given borrower receives from the same lender and year is one. The data also suggest that lenders are not "evergreening" loans–only (0.03%) of loans are categorized as “revolving” debt–and we remove these loans from the estimation sample. The institutional features of the SBA 7(a) program allows us to conduct a notch estimation for studying the impact of federal loan guarantees on credit supply.

3.3 Model and identification strategy

We model entrepreneurs as borrowing D at interest rate R from banks to fund their projects. Their projects are characterized by a productivity type that determines output and therefore the probability of success of the project. An entrepreneur's type is drawn from a distribution $F(r, n)$, characterized by the average type r and variance n. While r and n are known to both borrower and bank, the realized type $\underline{r}$ is revealed only after the loan is made and the project is attempted.

Once a project's output is realized, borrowers decide whether to repay the loan to the bank. The borrower pays nothing in default and pays $D(1 + R)$ otherwise. We assume that the borrower pays back as long as the realized output $\underline{r}$ is greater than the amount owed to the bank. Thus, the lender's payoffs are

$$\Pi = \begin{cases} -D & \text{if} \qquad \underline{r} < D(1+R) \\ D\bar{R} & \text{if} \qquad \underline{r} > D(1+R). \end{cases} \tag{3.1}$$

The lender loses the capital lent D if the borrower defaults and gains $D\bar{R}$ if the borrower repays.

Lenders have market power but are restricted to charge a regulated interest rate of $\bar{R}$, which is consistent with interest rate caps in SBA programs. They decide how much capital D to lend to a borrower by maximizing the expected profits:

$$E[\Pi] = \int_{D(1+\bar{R})} D \cdot \bar{R} \cdot f(r)dr - \int^{D(1+\bar{R})} D \cdot f(r)dr \tag{3.2}$$
$$= D \cdot \bar{R} \cdot Pr(r > D(1+\bar{R})) - D \cdot Pr(r < D(1+\bar{R})).$$

The first term is positive and represents revenue made from a repaid loan. While the mechanical revenue, $D \cdot \bar{R}$, is increasing in loan size, the probability of repayment, $Pr(r > D(1+\bar{R}))$, is decreasing. This term is concave in D so that it is equal to zero when D is zero or infinite and is otherwise positive. The second term represents the expected costs to the lender from borrower default. The probability of default is given by $Pr(r < D(1+\bar{R}))$, which is increasing in loan size. Thus $-D \cdot Pr(r < D(1+\bar{R}))$ is negative, convex, and increasing in D.

We remain agnostic about the exact distribution of r, and write the probability of default as $\pi(D, \bar{R})$, an increasing function of D. Lender profits are:

$$E[\Pi] = D \cdot \bar{R} \cdot (1 - \pi(D, \bar{R})) - D \cdot \pi(D, \bar{R}). \tag{3.3}$$

Given the tradeoff between increased revenue and a higher default probability, lenders choose the loan size that maximizes their expected profits. Optimal loan size is implicitly a function of $\pi(\cdot)$ and satisfies the first-order equation:

$$D^* = \frac{\bar{R}}{\pi'(D^*, \bar{R}) \cdot (1+\bar{R})} - \frac{\pi(D^*, \bar{R})}{\pi'(D^*, \bar{R})}. \tag{3.4}$$

We focus only on the loans for which a positive D^* exists given the distribution of realized output and the set interest rate, $\bar{R}$. Note that the optimal loan size will depend on the interest rate $\bar{R}$ as well as the mean and variance of realized productivity, which determine the shape of $\pi(D, \bar{R})$. All else equal, a borrower with a higher mean probability of default or higher variance will have a lower optimal loan size.

3.3.1 Lender's problem with a loan guarantee

We now analyze what happens to loan size when the lender receives an indirect loan guarantee. There are two key components of the federal loan guarantee program: a reimbursement rate and a fee. If a bank makes a loan that is ultimately charged off, the government will reimburse $\gamma\%$ of the losses. In return, the bank pays a certain fee equal to σ percent of the loan principal to the government. Given a charge-off probability, $\pi(D, \bar{R})$, the total expected subsidy S provided by the government on loan amount D is given by

$$S = \gamma \cdot \pi(D, \bar{R}) \cdot D - \sigma \cdot D = D \cdot \Gamma. \tag{3.5}$$

where the net generosity of the guarantee per unit of lending is given by $\Gamma = \gamma \cdot \pi(D, \bar{R}) - \sigma$. We assume that banks are risk-neutral so that a change in the reimbursement rate is isomorphic to a change in the fee.

The guarantee does not change the borrower's behavior since it is targeted toward and given only to the lender. Indeed, the guarantee is a contract between the lender and the government and hence should not directly affect borrowers other than through lender behavior. The lender's payoffs are now

$$\Pi = \begin{cases} \gamma D - (1+\sigma)D & \text{if} \quad \underline{r} < D(1+R) \\ \bar{R}D - \sigma D & \text{if} \quad \underline{r} > D(1+R). \end{cases} \tag{3.6}$$

and the expected profits are

$$E[\Pi] = D \cdot (\bar{R} - \sigma) \cdot (1 - \pi(D, \bar{R})) - (1 + \sigma - \gamma) \cdot D \cdot \pi(D, \bar{R}). \tag{3.7}$$

The guarantee decreases marginal revenue since it requires paying a fixed percentage fee, σ. However, in the case of a subsidy, the reimbursement component also decreases the marginal cost of lending from $\pi'(D, \bar{R})$ to $(1 + \sigma - \gamma) \cdot \pi'(D, \bar{R})$. We analyze what happens when the guarantee is made more generous using this profit function–specifically, what happens to profit and loan size when γ increases holding all else constant?

We focus on positive subsidy guarantees such that $\sigma = 0$, $\gamma > 0$, and $\Gamma = \gamma \cdot \pi(D)$. Taking the derivative of expected profit with respect to loan size gives us a new formula for D^* that relies on the guarantee generosity:

$$D^* \quad = \quad \frac{\bar{R}}{(1+\bar{R})\cdot\pi'(D^*,\bar{R}) - \Gamma\frac{\pi'(D^*,\bar{R})}{\pi(D^*,\bar{R})}} - \underbrace{\frac{\pi(D^*,\bar{R})}{\pi'(D^*,\bar{R})}}_{\text{Default effect}}. \tag{3.8}$$

This expression shows that the elasticity of loan size to the guarantee depends not only on the generosity of Γ but also on the size and shape of the default function. While a more generous guarantee decreases the costs of default borne by the lender–inducing lenders to increase the loan supply, D^*–a larger loan carries a higher probability of default. The magnitude of the elasticity of loan size to the guarantee is therefore inversely related to the local slope of the default function.

The local slope of the default function is determined by the productivity type distribution, $F(r,n)$. Specifically, as the variance of the productivity type increases, an equal sized change in D will cause a smaller change in the default probability. Thus, an increase in n flattens the slope of the default function and leads to higher lending supply elasticity with respect to Γ.

The top panel of Fig. C-1 simulates how loan size responds to a varying type of n. The change in the loan size is positively related to the variance of productivity type distribution, illustrating that the increase in the variance of expected returns leads to higher lending response with respect to Γ through a flattening of the slope of the default function. The bottom panel simulates changes in D^* as Γ increases for a high and low variance distribution of expected returns. Again, the guarantee has a larger loan size effect for the high variance distribution, which becomes amplified as the subsidy increases in generosity.

3.3.2 Impact of guarantee subsidy on lender profits versus additional lending

An increase in guarantee generosity is costly for the government. To what extent does this spending simply subsidize lenders, and how does the subsidy versus loan creation effect depend on the loan size elasticity? Given that D^* is implicitly a function of Γ, we rewrite lender profit as

$E[\Pi] = D^*(\Gamma) \cdot (\bar{R} \cdot (1 - \pi(D, \bar{R})) - \pi(D, \bar{R}) + \Gamma)$ and take the derivative with respect to Γ:

$$\frac{\partial E[\Pi]}{\partial \Gamma} = \underbrace{D'(\Gamma) \cdot (\bar{R} \cdot (1 - \pi(D, \bar{R})) - \pi(D, \bar{R}) + \Gamma) + D}_{\text{Increased revenue from larger loan}} - \underbrace{D(\Gamma) \cdot (1 + \bar{R}) \cdot \pi'(D, \bar{R}) \cdot D'(\Gamma)}_{\text{Decreased prob. of repayment}} \tag{3.9}$$

While complex, this derivative shows that profits change due to both loan size adjustment and the mechanical decrease in expected costs. Therefore, the extent to which the guarantee acts as a lender subsidy relies on the responsiveness of loan size to the guarantee rate, $D'(\Gamma)$.

Recall that the expected total cost of the guarantee subsidy is $S = D \cdot \Gamma$. If loan size is completely inelastic, the change in profits will be exactly equal to the change in costs, and the guarantee will act as a pure subsidy to lenders. As loan size becomes more responsive to the guarantee, the expected costs of the guarantee and net-of-guarantee losses for the lender increase. While the loan size increases more dramatically, less of the guarantee subsidy is retained by the lender. Fig. 3-1 illustrates this logic. The left panel shows that the fraction of the guarantee subsidy that goes to the lender declines as $D'(\Gamma)$ increases, while the right panel shows that the lending supply expands with $D'(\Gamma)$.

3.3.3 Identification strategy

Our identification strategy and the interpretation of our estimated elasticity relate closely to the curvature of the lenders' profit function modeled in Section B.1. Specifically, since lenders' profit functions are concave, we assume that there is a global optimum of the amount of capital D_i^* that the bank should lend to each borrower with mean productivity type, r_i. This optimum is shown in panel a of Fig. 3-2; the red dot indicates the point where a lender maximizes profit for a given borrower type.

We assume that there is a distribution of mean productivity types in the population, and thus the optimal amount of capital varies by the expected type. Therefore, even in the absence of the guarantee notch, this leads to wide variation in the amount of capital lent to the borrowers. Fig. C-2 illustrates this point–the observed loan distribution is wide, even in the placebo years when the guarantee notches were eliminated, indicating that the heterogeneity in loan size is driven by the underlying productivity types.

In our setting, we observe a loan size-specific guarantee subsidy that creates a discontinuity in the profit function with respect to D. A more generous guarantee applies to all loans below a specific loan size threshold, D^T. All else constant, this shifts the bank's profit function upwards in this region. As shown in the panel b of Fig. 3-2, the notch creates a new optimum for a certain subset of productivity types. In particular, for some borrowers that normally would be optimally located to the right of the notch, the notch will distort the distribution of observed loans, as it will now be profit maximizing for the bank to offer $D_i = D^T$. It is important to note that this will only impact the placement of loans that were previously located to the right of the notch. If the optimum was previously to the left of the notch, the guarantee will change the level of profit received by the bank but not the location of D_i^*.

Whether a given productivity type is affected by the notch is determined by how profit changes between D_i^* and D^T. Fig. 3-2 illustrates this point. For the borrower in panel b, the notch creates a new optimum. However, panel d shows that for the borrowers with an original optimum D_i^* further to the right away from D^T will be less likely to be relocated to the notch. This is because there is a smaller difference in the profit at D_i^* and D^T as D_i^* increases. Finally, panel c shows the borrowers that we refer to as the "marginal buncher," or those that the bank is indifferent between giving a loan at either D_i^* or D^T.

Our estimation strategy, which recovers the local slope of the profit function, relies on identifying the marginal buncher. We do this by comparing the observed distorted and counterfactual undistorted loan distributions. We identify the point to the right of the notch where the observed loan distribution is no longer distorted or impacted by the notch. This corresponds to the location of the marginal buncher. We define the distance between the undistorted optimal location of the marginal buncher and D^T as ΔD, and it is the key empirical determinant of our reduced-form elasticity. The method assumes homogeneity in profit function across types r_i. In our current setup, implying that the riskiness of the realized draw n does not vary with mean expected returns, r.

The location of the marginal buncher, and hence the measured elasticity, depends on the curvature of the profit function. Panel c of Fig. 3-2 plots both a very steep profit function (in blue) and flat profit function (in black) that both face a guarantee notch with the same size and location. It denotes the location of the marginal buncher in each case. The reduced-form elasticity that we estimate maps approximately to the inverse of the slope near and to the left of the optimum. As discussed above, this underlying curvature is determined by the underlying distribution of realized types $F(r, n)$ and the interest rate $\bar{R}$.

3.4 Empirical approach

As explained in Section 3.3.3, we identify and estimate the elasticity of lending to a change in the guarantee rate using the discrete change in the level of the guarantee rate in the SBA 7(a) lending program. The notch point created by the change in the guarantee rates creates incentives for lenders to shift loans below the guarantee notch point. If lending is elastic to the guarantee rate, lenders will be more likely to shift loans to a point below the notch where a loan carries a higher guarantee rate, whereas if lending is inelastic, lenders will not alter their behavior. Specifically, an elastic response will lead to "bunching" at the notch point, with excess mass below the notch point where guarantee rates are higher and missing mass above the notch point where guarantee rates are lower.

A bunching approach uses the excess mass at the threshold to estimate an implied lending response to the change in the guarantee rate and provides nonparametric estimates of the elasticity of credit supply. Recent papers employing bunching estimators include Kleven (2016), Best and Kleven (2018), DeFusco and Paciorek (2017), Saez (2010), Kleven and Waseem (2013). The method is related to, but distinct from a regression discontinuity approach. Regression discontinuity design exploits notched incentives when there is no manipulation of an assignment variable. In a bunching design, the manipulation of the assignment variable is used to identify the parameter of interest (see Kleven (2016) for a general overview of bunching estimators). In the subsequent analysis, we closely follow the methodology outlined in Kleven and Waseem (2013).

To implement the approach, we first recall that a bank i decides how much to lend, D_{ij}, to entrepreneur j using the objective function that maximizes returns in D_{ij}:

$$\max_{D_{ij}} D_{ij} \cdot (\bar{R} \cdot (1 - \pi(D_{ij}, \bar{R})) - \pi(D_{ij}, \bar{R}) + \Gamma_{ij}). \tag{3.10}$$

We calculate Γ_{ij} as the observed ex-post return on a loan, net of realized charge-offs, guarantee fee payments, and guarantee reimbursements. We use our loan-level data to first model an indicator for loan default as a function of loan size. We multiply the predicted default probabilities (π) by the guaranteed reimbursement rate (γ) to find the expected reimbursement rate on a given loan–this implicitly assumes a hundred percent charge-off rate on defaulted loans. We then subtract the loan fees (σ) paid to the SBA, which are expressed as a percentage of loan principal. This provides the net subsidy rate provided to banks by the SBA, the empirical analogue to $\Gamma = (\gamma \cdot \pi - \sigma)$ in Section B.1. A full description of the methodology we use to estimate Γ can be found in Appendix C.2.

Empirically, the default probability varies little across the threshold, whereas γ and σ vary significantly. We make the assumption that banks are risk neutral, which means that lenders treat a change in the reimbursement rate equivalently to a change in the fee. This generates a discrete drop in the return the bank makes on lending right above the threshold. Specifically,

$$\Gamma(D_{ij}) = \begin{cases} \Gamma, & \text{if } D_{ij} \leq D^T \\ \Gamma - \Delta\Gamma, & \text{otherwise.} \end{cases}$$

In the absence of a notch, we assume there would have been a smooth distribution of loans made that would satisfy the banks' first-order condition, conditional on and mapping directly to a smooth underlying distribution of loan demand, n_j. The notch, however, creates a region directly above the threshold for a subset of loans where marginal revenue is strictly lower than the marginal cost. The marginal bunching loan is made at the point $D^T + \Delta D$ where the bank is indifferent between making a smaller loan under the more generous guarantee and making a larger loan under the less generous guarantee:

$$\begin{aligned} &D^T \cdot (\bar{R} \cdot (1 - \pi(D^T, \bar{R})) - \pi(D^T, \bar{R}) + \Gamma) = \\ &\qquad (D^T + \Delta D) \cdot (\bar{R} \cdot (1 - \pi(D^T + \Delta D, \bar{R})) - \pi(D^T + \Delta D, \bar{R}) + (\Gamma - \Delta\Gamma)). \end{aligned} \tag{3.11}$$

Therefore, ΔD captures the reduction in dollars lent in response to the change in the guarantee rate for this marginal buncher, and it is the key empirical parameter needed to calculate the elasticity of lending. The substantial excess mass we observe in the data at the point $D_{ij} = D^T$ comes from this region of strictly dominated lending for the bank $(D^T, D^T + \Delta D)$ directly above the notch point. This approach allows us to map the amount of excess mass to the loan response ΔD using the bunching methodology we discuss below in Section 3.4.1.

Within the dominated region, the bank can always increase its return by making smaller loans under the higher guarantee rate, Γ. As discussed in Section 3.3.3, the size of the dominated region (and therefore the reduced-form elasticity of lending with respect to the guarantee rate) relates to the slope of the default function $\pi(D)$–if a small change in D generates a sharp increase in costs, there will be a small dominated region and a small elasticity of lending. If a change in D has little impact on costs, then there will be a larger dominated region, more bunching at the threshold, and a larger elasticity of lending with respect to the guarantee rate. The analysis treats the guarantee parameter as exogenously set by the SBA.

3.4.1 Bunching methodology

This section describes the estimation methodology in detail. Our objective is to estimate the reduced-form lending elasticity with respect to the guarantee generosity, or the percentage change in dollars lent that results from a corresponding percentage change in the guarantee generosity:

$$\varepsilon_{D,\Gamma} \equiv \frac{\Delta D}{D^T} \times \frac{(1 + \Gamma^T)}{\Delta \Gamma}. \tag{3.12}$$

Here $\Delta\Gamma$ is the change in the marginal guaranteed return faced by the bank. We estimate the elasticity by noting that a notch in the marginal guarantee rate allows us to approximate the implicit marginal guarantee rate , Γ^T, created by the notch $\Gamma^T \approx \Gamma + \Delta\Gamma \frac{D^T}{\Delta D}$. We can then write the reduced form elasticity as

$$\varepsilon_{D,\Gamma} \approx \left(\frac{\Delta D}{D^T}\right)^2 \times \frac{(1 + \Gamma)}{\Delta \Gamma}. \tag{3.13}$$

The validity of the bunching estimate relies on three keys assumptions: first, the counterfactual distribution would be smooth in the absence of a notch; second, bunchers come from a continuous set such that there exists a well defined marginal buncher; third, optimization frictions are locally constant. Interpreting this as purely the effect of a change in the guarantee rate requires a fourth assumption that contract terms do not change at the notch point due to the presence of a guarantee.

The first assumption rules out that other factors are changing at the threshold, which might bias our estimates. The assumption effectively means that there are no other policies at the threshold that would induce borrowers to move to the notch point. To our knowledge, there are no other relevant contract parameters, and we confirm this empirically for observable contract parameters in the data. This assumption also captures extensive margin responses and implies that locally borrowers move to the notch rather than choosing not to embank on projects.

While the second assumption is technical and fairly weak, the third assumption is stronger. The assumption that optimization frictions are locally constant allows the use of the dominated region to the right of the notch to identify behavioral responses and parameters of interest. This assumption requires that the mass of set of movers equals the total area under a counterfactual on the other side of the notch point.

The assumption that contract terms do not change at the notch point due to the presence of a guarantee is likely to hold in our context due to the particular institutional details. The main parameter that lenders might vary in response to the guarantee is price. Empirically, we observe that interest rates trend smoothly across the notch. This is likely due to the presence of interest rate caps–the vast majority of lenders price right at the cap.

It is important to note that there are a wide variety of current and historical government guarantee programs, ranging from the mortgage and student loan markets to the small business loans that we study. Our estimates are for small business loans between \$75,000 and \$225,000. It is possible that outside the range, lending supply could be more or less responsive, and it is also possible that effects are different in other loan markets.

We obtain the parameters for elasticity estimation from the SBA data. The threshold D^T is \$150,000 for the years in our sample. We calculate $(1 + \Gamma)$ as the observed ex-post return on a loan, net of realized charge-offs, guarantee fee payments, and guarantee reimbursements. As noted earlier, interest rates and ex-post charge-off rates trend smoothly through the threshold. Therefore, all systematic variation in returns come from changes in the generosity of the guarantee contract at the threshold. Over our time period, loans less than or equal to $150,000 had lower guarantee fees and higher guarantee reimbursement rates than loans to the right of the threshold. Given that the generosity varies over time, we estimate the excess mass and elasticity separately by year. Note that while the effective generosity of the guarantees vary over time, the charge-offs are low and stable over time. Fig. C-3 shows that the three-year cohort default rates are relatively stable except during the Great Recession. Therefore our elasticity measure is identified through the variation in the guarantee rates rather than time-varying default rates.

To calculate ΔD empirically, we must locate the counterfactual loan amount provided to the marginal buncher. This occurs at the point where the excess mass at the threshold is equal to the missing mass to the right of the threshold. To measure the excess and missing mass, we estimate the counterfactual loan distribution that would have occurred in the absence of a notch by fitting a polynomial of degree six with a vector of round number dummies for multiples of 1, 5, 10, 25, and 50 thousand and excluding a region at and to the right of the threshold:

$$N_j = \sum_{k=0}^{6} \beta_k (d_j)^k + \sum_{i=d_l}^{d_u} \delta_{ij} \mathbb{1}(d_j = i) + \sum_{n \in \{1k, 5k, 10k, 25k, 50k\}} \delta_n \mathbb{1}(d_j = n) + \eta_j. \tag{3.14}$$

where N_j is the number of loans in bin j, d_j is the loan amount midpoint of interval j, $\{d_l, d_u\}$ is the excluded region, δ_{ij} are dummies for bins for the excluded region, and δ_n are dummies for multiples of prominent round numbers. For estimation, we cut the data into $500 dollar bins and restrict the loan size to be between $75,000 and $225,000 to limit the estimation range. For robustness, we repeat the estimation with bin sizes of $200, $1,000, and $2,000, polynomials of degree four, five, and seven, and for various ranges of loan samples; these results are shown in the appendix. While the results are very robust to the different bin and polynomial choices, they are sensitive to the inclusion of $50,000 within the range. Another interest-rate-related threshold exists at the $50,000 mark, which causes additional bunching, and therefore we exclude it from our estimation. The counterfactual distribution, $\hat{N}_j$, is estimated as the predicted values from Eq. 3.14 using the β_k and the δ_n terms

$$\hat{N}_j = \sum_{k=0}^{6} \hat{\beta}_k (d_j)^k + \sum_{n \in \{1k,5k,10k,25k,50k\}} \hat{\delta}_n \mathbb{1}(d_j = n). \tag{3.15}$$

Excess mass is defined as the difference between the observed and counterfactual bin counts between the lower limit of the excluded region (d_l) and the threshold, $\hat{B} = \sum_{j=d_l}^{D^T}(N_j - \hat{N}_j)$, whereas the missing mass, $\hat{M} = \sum_{j=D^T}^{d_u}(N_j - \hat{N}_j)$, is defined as the same bin counts but is in the range between the threshold and the upper limit of the excluded region (d_u).

While the lower limit d_l is easily observable visually as the notch point, we do not observe a sharp valley to the right of the cutoff. This is common in bunching estimators (Kleven and Waseem, 2013). Thus, to identify the upper limit d_u, we follow an iterative procedure. We identify the upper limit (i.e., $d_u = D^T + \Delta D$) by requiring that the excess mass $\hat{B}$ be equal to the missing mass, $\hat{M}$.

The estimation procedure proceeds in four steps. First, the estimation begins with a starting value of d_u right above D^T. Second, we calculate $(\hat{B} - \hat{M})$. The next step is to increase d_u by a step size of \$500 if $(\hat{B} - \hat{M} \neq 0)$. Finally, we repeat these steps until the result converges. We pool together all banks in our main estimation. However, to test whether the elasticity and bunching is driven by a specific bank, we have also repeated the estimation on a conditional distribution that controls for bank fixed effects. The bunching and elasticities are very similar.

We calculate standard errors for Eq. (3.13) using a non parametric bootstrap procedure in which we draw a large number of loan distributions following Chetty, Friedman, Olsen and Pistaferri (2011). We create new bins of loans by drawing randomly with replacement from the estimated vector of $\eta_{\mathbb{j}}$ and by adding those to the estimated distribution implied by the coefficients $\{\hat{\beta}_k, \hat{d}_j, \hat{d}_u\}$ from Eq. (3.14). Finally, we apply the bunching estimator technique described above again to calculate a new estimate $\hat{\varepsilon}^b_{D,\Gamma}$. We repeat this procedure 1,000 times and define the standard error as the standard deviation of the distribution of $\hat{\varepsilon}^b_{D,\Gamma}$s created. As we observe the universe of SBA 7(a) loans over the years considered, the standard error represents error due to misspecification of the polynomial and the number of dummies included in the exclusion zone used in rather than representing sampling error.

Fig. 3-3 visually illustrates the variation that we use to identify the elasticity of credit supply to the loan guarantee. The figure shows the raw data in 2013, where the guarantee notch is relatively small, and again in 2015 when the guarantee notch is larger. Fig. 3-3 illustrates the striking contrast in bunching in 2013 when there was a small notch, and in 2015 when there was a large notch. The left panel shows the number of loans, in discrete \$2,000 bins, while the right panel shows the total expected guarantee benefits. In 2015, where the total expected guarantee is comparatively higher, we see more bunching relative to 2013.

The bunching technique captures intensive margin responses. If banks reject applications simply because they are above the threshold, this would lead us to underestimate the credit supply response to the guarantee further away from the notch and to make our estimates more sensitive to the choice of polynomial. Since banks have considerable power when deciding how much to lend and could increase returns by reducing D_{ij} rather than not lending at all, these extensive margin responses are unlikely in our setting. However, we still test the sensitivity of our estimates to the choice of parameters. Kleven and Waseem (2013) show that these extensive margin responses should only occur in a region far to the right of the notch, with the intensive margin response concentrated in the area directly next to the notch. They note that extensive margin bias will mainly enter via functional form misallocation, and therefore sensitivity analysis should be conducted with respect to the polynomial. We show in Table 3.3 that our results are robust to using a range of polynomials, which suggests that extensive margin responses do not play a large role in our setting.

3.5 Main results

3.5.1 Visual evidence

We begin by showing the change in guarantees at the \$150,000 threshold. The top panel of Fig. 3-4 shows average guarantees and fees by loan amounts as a percentage of the loan principal amount in \$2,000 bins across the threshold between 2008 and 2017. Consistent with the policy rule, the guarantee benefit jumps sharply across the threshold–loans below \$150,000 receive a guarantee rate nearly twice as generous as loans above the threshold. Appendix Fig. C-4 breaks down the guarantee benefit by the average expected guarantee fees and reimbursement rate separately.

To determine whether the guarantee benefit notch affects lending volumes, we analyze the density of borrowing. The bottom panel of figure 3-4 shows bunching directly below the threshold. The figure shows the number of loans in $2,000 bins across the threshold between 2008 and 2017. Visual evidence indicates that there are significantly more loans at the threshold relative to other points nearby, consistent with banks lending fewer dollars in response to a lower guarantee rate (i.e. moving borrowers to loan volumes below the notch).

Fig. 3-5 shows the observed and counterfactual density of loans. The solid line shows the observed number of loans, while the dashed line shows the counterfactual number of loans. The counterfactual is determined according to the method discussed in Section 3.4 and is estimated as specified in Eq. 3.15. Several patterns are immediately clear from Fig. 3-5. First, there are significantly more loans disbursed just at the threshold, which is consistent with guarantees affecting credit supply. Second, there is also missing mass to the right of the guarantee notch. In other words, the counterfactual distribution is higher than the observed distribution. Third, the observed number of loans is lower to the right of the threshold. Finally, there is significant round number bunching, which is captured by our modeling procedure.

The presence of two placebo years in 2009 and 2010 in the middle of our sample period, when no notch existed, provides a direct test of the first assumption that the counterfactual distribution would be smooth in the absence of a notch. Fig. C-5 shows that the bunching disappears completely in these years and suggests that there are no other unobserved factors generating bunching at the threshold. These years also allow us to test the fit of our estimated counterfactual not only at the notch but also across the rest of the loan distribution. The observed and estimated distributions in Fig. C-5 are almost identical, which indicates that our counterfactual specification accounts well for the round number bunching in the distribution.

3.5.2 Elasticity estimates

Table 3.3 formalizes and scales the bunching noted above relative to the change in the size of the guarantee and presents estimates of $\varepsilon_{D,\Gamma}$ as described in Section 3.4. The first column shows the degree of the polynomial used to estimate the counterfactual distribution- we vary this to test sensitivity to the parameter choices and to gauge whether extensive margin responses are playing a large role. The second column shows the estimated excess mass, $\hat{B}$, in terms of the number of loans. The third column shows estimates of ΔD, the distance of the marginal buncher in dollar terms from the threshold. The fourth column presents $\Delta\Gamma$, the change in the generosity of the guarantee rate at the notch. Over the years in our sample, $\Delta\Gamma$ varied between 0 and 0.078. For this estimate, we take a weighted average of $\Delta\Gamma$ in non zero years to pool across years; in the appendix we also list estimates by year. The final column shows estimates of $\varepsilon_{D,\Gamma}$, the elasticity of dollars of loans made with respect to the guarantee rate.

The first panel show estimates from placebo years when the notch was eliminated as part of the ARRA stimulus of 2009. Reassuringly, we see very little excess mass when loan guarantees are identical across the notch. This assuages potential concerns that other factors could be changing across the threshold and is discussed further in the next section. Note that we cannot compute elasticity estimates in 2009 and 2010, as there is no variation in the notch.

The second panel shows estimates from years when the guarantee notch was binding. The estimates of the elasticity are between 4.5 to 5.2, depending on the polynomial used. Interpreted in dollar magnitudes, this means that a 3.8 percentage point change in the guarantee subsidy rate (Γ) generates an approximate \$70,500 in additional lending.

It is important to note that we estimate a reduced-form elasticity, which could be affected by optimization frictions. Optimization frictions are factors that prevent agents from locating at notch or kink points. For example, in labor supply estimates workers might be unable to alter their hours worked due to contractual arrangements, and in our context, projects might need a certain amount of capital. Notches are particularly useful in bunching estimators, relative to kinks, because in the absence of optimization frictions, theory predicts an excluded region to the right of the notch. Kleven and Waseem (2013) show methods to identify upper and lower bounds for the true structural elasticity. Under the admittedly strong assumption that the structural elasticity is homogenous, the reduced-form estimate is a structural elasticity. Otherwise, if there is heterogeneity in elasticities, the upper bound is represented by the response of the most sensitive individual.

We can place an additional restriction to identify the lower bound for the true structural elasticity, ϵ. This approach requires identifying α, or the share of individuals with sufficiently high adjustment costs that they are unresponsive to the notch. In this case, the term $\varepsilon_{D,\Gamma} = (1-\alpha)\epsilon$ is a lower bound for the true structural elasticity. We can use the share of individuals who do not optimize in a given year (i.e., the number of loans that are larger than \$150,000 as a share of all loans in the dominated region) to estimate that α. In years with a very high guarantee notch, we see approximately 40% of borrowers locating in the dominated region, suggesting that $\alpha \approx .4$. Thus we obtain a lower bound for the structural elasticity of approximately $\epsilon \approx 8.3$ (i.e., $5 = (1-0.4) \times \epsilon$). Intuitively, the reduced-form elasticity $\varepsilon_{D,\Gamma}$ is the observed elasticity attenuated by optimization frictions, α. Therefore, the structural elasticity is greater than the reduced-form elasticity in the presence of frictions.

3.5.3 Bunching over time and placebo estimates

The observed amount of bunching varies over time with the size of the guarantee notch. Fig. 3-6 shows bunching at the guarantee notch for each year between 2008 and 2017. The figure groups years into three broad groups, years during which there is a high, low, or no guarantee notch. Between 2014 and 2017, the size of the notch was between 0.04 and 0.08 of the average expected guarantee benefit as a percentage of the loan principal. In 2008, and between 2011 and 2013, the notch was between 0.02 and 0.03 of the average expected guarantee benefit as a percentage of the loan principal. In 2009 and 2010 the notch was eliminated as part of the ARRA.

We see a very close relation between the guarantee change and observed bunching at the threshold, defined as the difference in the share of loans between the observed and counterfactual density. In years with a large change in the guarantee, we see greater excess mass relative to years with a lower guarantee change at the notch. However, in years when the notch was eliminated (i.e., "placebo years"), there is no excess mass at the threshold.

Fig. 3-7 provides additional reduced-form evidence that this bunching is indeed driven by guarantees. The generosity of the guarantee across the notch has varied significantly over time, which allows us to explore dynamic aspects of the lending response. Consistent with the bunching being driven by loan guarantees, and not by any other factors changing across the threshold, we find higher excess mass in years when the difference in the guarantee across the threshold is greater.

Fig. 3-7 shows the relation between share of excess mass at the threshold and the guarantee rate in each year. For this figure, we again use our reduced-form measure of excess mass: we observe some bunching at round number points, as is shown in Fig. C-6.

To account for this bunching, we calculate excess mass at the threshold relative to intervals of \$50,000 between \$50,000 and \$300,000. The figure shows the amount of bunching occurring at the \$150,000 threshold against the size of the guarantee change at the threshold between 2008 and 2017, in ten bins absorbing bank fixed effects. The left panel plots the share of excess mass and the change in the guarantee at the threshold. There is a striking linear relation between the share of excess mass and the guarantee rate. The right panel shows the relation between the share of excess mass and guarantee rates over time. The figure shows that the observed excess mass comoves with the guarantee rate, indicating a strong relation between the incentives to bunch and the amount of bunching.

Table 3.4 repeats the main analysis, showing estimates year by year. While estimates are relatively stable between 2008 and 2013, and similar in 2017, the estimates of $\varepsilon_{D,\Gamma}$ are about one-third the size of estimates in other years in 2014 and 2015. We see little excess mass in years when the notch was eliminated, and excess mass starts to grow sharply in 2014 when the guarantee notch becomes larger.

This growth in excess mass suggests the existence of adjustment frictions, where banks could take some time to increase credit supply. This can be seen in the bottom panel of Fig. 3-7. While there is a sharp jump in the guarantee notch between 2013 and 2014, approximately doubling from 0.033 to 0.077, the increase in excess mass is more gradual and increases year by year. The pattern translates to an initially lower elasticity, which increases to between 4.5 and 6 in 2017. Similarly, we observe some loans being made in the dominated region directly to the right of the threshold, suggesting that banks face optimization frictions when trying to adjust loan sizes. Therefore we estimate a reduced-form elasticity that is inclusive of adjustment costs rather than a structural elasticity.

The 2009 ARRA stimulus provides a placebo check. As part of the stimulus, the SBA temporarily raised the guarantee rate to 90% and waived fees in 2009. This effectively eliminated the guarantee notch at $150,000. It is immediately evident graphically that the lending response drops when guarantee notches are eliminated. The bottom rows of Fig. 3-6 shows the excess mass during years in which the notch was eliminated. Between 2009 and 2010, when guarantee rates were identical across the threshold, we do not observe any excess mass beyond round number bunching. The fact that excess bunching is only present in years when the guarantee rate is discontinuous assuages a potential concern that other factors might change discontinuously across the threshold.

3.5.4 Magnitudes

This section discusses the implied magnitudes of our estimates. The average guarantee subsidy rate over all years and for loans below $350,000 is 5.1% of loan principal. Thus a lender making a loan through the guarantee program receives a subsidy worth 5.1% of the loan size. The subsidy rate includes the expected reimbursement the lender will receive on any losses minus the guarantee fees ($\Gamma = \pi \cdot \gamma - \sigma$). Empirically, the guarantee subsidy generosity varies over years and loan size from -4%–when the guarantee fees outweigh the expected reimbursement–to 11.6%.

Our elasticity estimate suggests that an increase in one percentage point of the guarantee subsidy rate (Γ) for a given loan would generate an intensive margin response of $19,054 in additional lending. To increase the overall guarantee subsidy rate, the SBA could either increase the reimbursement portion (γ) or decrease the guarantee fees (σ). Increasing the reimbursement rate on a loan from 80% to 90% would increase the overall subsidy rate by $10\% \times \pi = 0.37\%$. The average charge-off rate over all years in our data is 3.7% and generates $8,002 in additional lending. This rate is based on the three-year cohort default rate. Decreasing the loan fee (σ) from 2.89% of loan principal (the average rate in 2008) to 0% (the rate in 2009) would increase the overall subsidy rate by 2.89% and generate $55,066 in intensive margin additional lending. Analyzed from the perspective of our model in Section B.1, the elasticity suggests that additional lending has little impact on marginal default probabilities. Thus, lenders capture a relatively small portion of the subsidy.

These elasticities are on the higher end of estimates used for calibrations in Gale (1991) and consistent with elasticities used for model parameters in Lucas (2016). Lucas (2016) notes that supply elasticity is high in times of high levels of bank reserves and loose monetary policy. Overall, we argue that loan guarantees do indeed impact lending to small businesses by increasing loan volume.

3.5.5 Risk-shifting

A natural question is whether guarantees lead lenders to issue riskier loans. A higher portion of charged off dollars could induce lenders to be more lax in screening borrowers or to take fewer steps in monitoring borrowers and preventing defaults. One possibility is that the generosity of the guarantee rate pushes banks to lend to riskier borrowers (adverse selection) or deteriorates incentives to prevent charge-offs of loan applicants (moral hazard). Moral hazard and adverse selection on the part of the entrepreneur are unlikely in our context. Lenders, not borrowers, interface with the SBA programs. Borrowers rarely know that they are borrowing through the SBA program, and all changes in fees and reimbursement rates impact the bank directly, not the borrower. On the other hand, lenders might screen borrowers less thoroughly due to the guarantee.

We explore this question by exploiting temporal variation in the guarantee notch, testing whether banks shift loans more likely to be charged off to the notch when the guarantee benefit is higher. Table 3.5 shows estimates of variants of the following specification:

$$\pi_{i,t} = \alpha_i + \alpha_t + \alpha_m + \alpha_l + \delta \mathbf{1}(\mathrm{D} > D^T) + \zeta \mathbf{1}(\mathrm{D} > D^T) \times \Gamma + \xi \mathbf{1}(\mathrm{D} = D^T) \times \Gamma + \nu_{i,t}, \quad (3.16)$$

where $(\mathrm{D} = D^T)$ is an indicator of whether a loan is at the notch, $(\mathrm{D} > D^T)$ is an indicator of whether the loan is above the notch, and Γ is the guarantee generosity. The outcome of interest is π_{it}, which is various measures of loan charge-off. Specifications include year or year-month fixed effects α_t, lender fixed effects α_i, maturity fixed effects α_m, and loan size bin fixed effects α_l. The main coefficient of interest is ξ, which captures the difference in charge-offs at the notch.

The results in Table 3.5 suggest that lenders indeed do shift riskier loans to the notch, where the guarantee rate is higher. The odd columns include only year fixed effects α_t, while the even columns include year-month fixed effects as well as lender fixed effects α_i, maturity fixed effects α_m, and loan size bin fixed effects α_l. In the first pair of estimates the dependent variable is an indicator of whether a loan is charged off, in the second pair of estimates the dependent variable is the percentage of the principal charged-off, while in the third pair of estimates the dependent variable is the log of the charged off amount.

Table 3.5 indicates that higher guarantee amounts are associated with higher charge-offs at the notch. All three dependent variables indicate higher levels of loan charge-off when the guarantee generosity is higher. A 10 percentage point increase in the generosity of the notch is associated with a 1.8 to 3.2 percentage point increase in charge-offs, a 1.7% to 2.5% increase in the amount of principal charged off, and a 0.23% to 0.37% increase in the amount charged off at the notch relative to the rest of the distribution. The estimates on the interaction term are statistically significant at the 0.01 level in all specifications. The table thus provides strong evidence that lenders are shifting risky loans to the notch when the guarantee generosity is higher.

Fig. 3-8 presents similar results graphically. Specifically, the top panel figure plots point estimates ξ_t and a 95% confidence interval from the following specification:

$$\pi_i = \alpha_b + \alpha_y + \alpha_m + \alpha_l + \delta \mathbf{1}(\mathrm{D} > D^T) + \zeta \mathbf{1}(\mathrm{D} > D^T) \times \Gamma + \sum_{2008}^{2016} \xi_t \mathbf{1}(\mathrm{D} = D^T) \times \Gamma + \nu_i. \quad (3.17)$$

Fig. 3-8 indicates that the patterns in the difference in charge-offs largely track the generosity of the notch in each year, which is shown in the bottom panel. The coefficients ξ_{2014}, ξ_{2015}, and ξ_{2016} are particularly high when the size of the notch is greatest. We see very small and insignificant estimates of ξ_{2009} and ξ_{2010} when the notch was eliminated.

3.6 Alternative channels, robustness, and placebo estimates

3.6.1 Demand and supply elasticities

One concern is that our estimates do not identify lenders' elasticity of supply to the guarantee rate but rather identify borrowers' elasticity of demand. It is in theory possible that guarantees are passed through to borrowers in the form of lower interest rates. Specifically, borrowers might be more likely to apply for a smaller $150,000 loan if the guarantee is passed through via a lower interest rate or lower risk standards. However, there are several institutional details that make a demand channel unlikely. As noted earlier, lenders are unable to issue multiple loans to the same borrower under the SBA program, making manipulation of the notch unlikely. Furthermore, borrowers must have exhausted all other financing options to qualify for an SBA loan, which rules out the possibility that banks or borrowers are topping up their SBA loans with additional private funding. Indeed, the eligibility criteria listed on the SBA website specifically states that to qualify for a 7(a) loan, "the business cannot get funds from any other financial lender." The observed data are also inconsistent with this demand side hypothesis. We find that a negligible portion (0.03%) of loans are categorized as revolving debt, i.e., a line of credit that can be drawn down by the borrower, and could also lead to demand-driven manipulation of the notch.

Despite the fact that institutional details make this demand channel unlikely, we verify whether the notch induces borrowers to bunch at the threshold by observing whether interest rates or ex-post charge-off rates (a measure of borrower risk) change discretely at the threshold. Fig. 3-9 shows average interest rates and the guarantee notch. Interest rates evolve smoothly despite the sharp guarantee notch. Fig. C-7 provides some insight as to why this could be the case–the majority of loans are priced at the cap on each side of the threshold.

Fig. C-8 shows that other factors trend smoothly across the threshold. Interest rates, revolving loan status, charge-offs, and loans terms all evolve smoothly, which suggests that the generosity in the guarantee is not passed on to the borrower through either an intensive margin interest rate effect or an extensive margin rationing effect. The figure implies that borrowers have no incentives to bunch at the threshold because requesting smaller loans to bunch at the notch only gives them less capital with no added benefits. Given this lack of incentives to bunch from the perspective of the borrowers, it is unlikely that the bunching is demand driven.

It is also possible in theory that borrowers request smaller loans than they otherwise would have if they believed that bunching at the notch improves their odds of getting the loan approved.

If this is the case, this is still interpretable as a supply elasticity since it is operating through a supply side mechanism: the approval rate. If the supply side was not reducing credit supply to the right of the notch, borrowers would not modify their loan requests.

3.6.2 Competition and loan substitution

Loan substitution

One potential concern is that we are not measuring a supply elasticity but rather a substitution elasticity (i.e., the loan guarantee is not increasing total credit supply, but rather incentivizing banks to shift loans from their SBA small business portfolio into the non-SBA portfolio or vice versa). Such within-bank substitution would generate a discontinuity in the number of loans originated at the \$150,000 size cutoff. While this channel would not generate excess mass at the \$150,000 notch, it could generate spurious missing mass to the right of the notch if banks place low-guarantee loans in their non-SBA portfolio.

To assuage a concern that spurious missing mass can confound our elasticity estimates, we estimate and compare elasticities on loans originated by a subsample of lenders that do and do not specialize in making SBA loans. A number of lenders, such as Live Oak Bank, specialize in making SBA-guaranteed loans and offer few, if any, other products. Thus, if the elasticity estimates between specialized and nonspecialized lenders are similar, it implies that it is unlikely that lenders shift loans between SBA and non-SBA products.

To identify lenders who specialize in SBA lending, we link SBA lenders to Call Report data and compute the total share of SBA loans originated by each lender. Next, we merge the SBA data set with quarterly Statistics on Depository Institutions (SDI) data from the FDIC to capture non-SBA loans. We match the majority of banks in our data (including federal credit unions) at an overall rate of 83% and a rate of 96% conditional that Call Report data exist (prior to Q1 2010, SDI reports were only provided yearly in Q2). The SDI data record the total number and amount of small business loans outstanding at a quarterly level per institution and further split small business lending into categories of loan size and purpose. We specifically look at small business commercial and industrial loans under \$1 million since these are most comparable to those provided through the 7(a) program. We also aggregate the SDI statistics to the yearly level. Appendix C.3 provides further information on the FDIC SDI data and a description of how we compute the SBA loan share by lender.

The top two panels of Table 3.6 show sample splits by lenders that do and do not specialize in SBA lending. Plots of the estimated counterfactual density for both splits are in Fig. C-13. The first panel splits lenders by whether the share of SBA loans is above or below 60% of their entire loan portfolio, while the second panel splits lender by above and below 80% share. The elasticity estimates are slightly higher at SBA specialized lenders, but overall the estimates are very similar. Thus, we do not find evidence that our results are biased by lenders substituting loans between SBA and non-SBA products.

In addition to comparing elasticities, we directly test whether banks that specialize in SBA lending are more likely to substitute non-SBA for SBA loans when the guarantee generosity is higher. If higher guarantees incentivize banks to shift their small business portfolio to SBA loans, we would expect this effect to be concentrated among banks with higher propensity to issue SBA loans relative to other small business loans. We explore this in the appendix by comparing differential responses between high- and low-SBA share lenders when guarantee rates were increased during 2009 and 2010. Table C.4 shows estimates of the following specification:

$$D_{i,t} = \alpha_i + \alpha_t + \delta \mathbf{1}(\text{Treat}) + \zeta \mathbf{1}(\text{Treat}) \times \theta_i + \varepsilon_{i,t}, \tag{3.18}$$

where the outcomes include the log of total, non-SBA, or SBA loan amounts; $\mathbf{1}(\text{Treat})$ is an indicator that equals one for years 2009 and 2010; θ_i is a bank-specific share of the amount of SBA lending relative to its overall small business lending portfolio in 2008; and α_i and α_t are bank and year fixed effects. Since the reimbursement rate increased to 90% on both sides of the $\$150,000$ threshold in 2009 and 2010 due to the ARRA stimulus, the estimated coefficients δ and ζ, respectively capture the effect of increased guarantees on the composition of small business lending portfolio and the differential response for banks with higher pre-ARRA propensity to issue SBA loans.

Table C.4 shows that while banks with higher pre-ARRA share of SBA lending increased SBA lending more in response to higher guarantees, the effects on non-SBA and total loan supply are not statistically significant. Fig. C-9 illustrates this finding graphically by plotting ζ_t from the following equation for total and non-SBA and SBA small business loans:

$$D_{i,t} = \alpha_i + \alpha_t + \delta \mathbf{1}(\text{Treat}) + \sum_{2008}^{2016} \zeta_t \mathbf{1}(\text{Year} = t) \times \theta_i + \varepsilon_{i,t}. \tag{3.19}$$

Consistent with the results in Table C.4, while banks with higher propensity to issue SBA loans differentially increase SBA loan supply in 2009 and 2010, with ζ_{2009} and ζ_{2010} being positive and significant for SBA loans, the effect on non-SBA loans are not statistically different from zero. These results confirm that within-bank substitution between SBA and non-SBA loans are unlikely.

Competition

A notch can incentivize borrowers to smooth the lenders' bunching behavior through borrowing from other sources. To the extent that a notch leads borrowers to seek funds from other sources, this can mitigate the credit supply effect. The institutional detail suggests that this is unlikely. SBA 7(a) loans typically carry higher interest rates than most other loan products, making it unlikely that borrowers would seek SBA loans if other financing options are available. Moreover, the SBA requires that lenders document and verify that a borrower passed the credit elsewhere requirement, which demonstrates that a borrower has "exhausted" all options for getting funds and cannot obtain funds without undue hardship.

While the SBA loans are intended to serve borrowers that cannot obtain loans elsewhere, it is still possible that this test is ineffective or poorly enforced. To explore this channel, we conduct sample splits by the number of banks operating in a borrower's county. In geographic areas with fewer operating banks, it can be more difficult for firms to access other forms of credit because the market is concentrated. Thus, if estimates are similar across areas with varying bank competition, we can infer that credit availability in a local market plays a small role in how lenders respond to changes in the guarantee rate.

The bottom two panels of Table 3.6 report the results. The first panel splits the sample by loans originated in areas where the number of banks is above or below three, while the second panel splits the sample by areas where the number of banks is above or below seven. While the estimates in counties with fewer banks are slightly lower relative to counties with more bank competition, we still observe significant excess mass and large elasticities between three and five in counties with fewer banks. This suggests that we see similar bunching effects in less competitive markets.

The top two panels of Table 3.6 report that we see similar elasticities for specialized lenders that are very likely to be compliant with the credit elsewhere test. Lenders can be excluded from the guarantee program if they repeatedly fail to verify credit elsewhere tests. Since exclusion from the program is extremely costly for lenders that specialize in making SBA loans, they are very likely to be compliant with the credit elsewhere test. Thus, this result supports the idea that there is a significant lending supply response from lenders who are compliant with the credit elsewhere test.

3.6.3 Alternative ranges

Table C.5 varies the range used in the estimation. We vary the loan sample range and bin size. The first column denotes alternative loan size ranges, while the top row denotes alternative bin sizes. The elasticities remain large and significant–between three and seven–when using alternative ranges and bin sizes, similar to those reported in Table 3.4 when we vary the bin size. There is some variation in the elasticity estimates stemming from varying the estimation range. There are two factors that make the estimates somewhat sensitive to the choice of estimation range.

First, there are two additional thresholds to the left and the right of the guarantee notch, which constrain our ability to extend the estimation region further and cause estimates to change in the vicinity of these notches. First, there is an interest rate cap notch at \$50,000. For loans below \$50,000, lenders can charge a one percentage point higher interest rate. Second, there is a collateral notch at \$350,000. Lenders are required to collateralize the loan to the maximum extent possible up to the loan amount when they borrow over \$350,000. If business assets do not fully secure the loan, the lender must take equity in the principal and/or personal real estate as collateral. This leads to significant bunching at \$350,000. We thus avoid ranges near the other policy notches.

Second, much of the variation in the calculated elasticities stemming from varying the upper bound of the range is driven by the estimation procedure itself. We calculate the elasticity with the upper bound when the terminal condition of the estimation is reached. In other words, the mass of $\hat{B}$ and $\hat{M}$ are in a given tolerance range. Given the iteration in the estimation procedure, this sometimes leads to the result that $\hat{M} > \hat{B}$. This is because we do not assign the upper bound as the value that best equalizes $\hat{M}$ and $\hat{B}$ but assign the d_U at which the estimation is terminated because the tolerance threshold is reached. This terminal d_U can in some cases be greater than the d_U that equalizes $\hat{M}$ and $\hat{B}$. For most of the specifications–including our main preferred spec–this does not happen (i.e., the terminal $\hat{M}$ = $\hat{B}$ is within a 0%-5% margin). However in the rare cases when it does happen, this can cause the elasticities to jump.

We further explore alternative estimation ranges graphically. Fig. 3-10 shows point estimates of the elasticity using various ranges, along with a 95% confidence interval. The estimates are generally centered around 5, with a range between 2.5 and 11. The estimates generally increase when we expand the range. Appendix Fig. C-10 shows that the estimates are more sensitive to varying the upper limit rather than the lower limit.

3.6.4 Additional robustness

We show in the appendix that the elasticity estimates are robust to a number of alternative specifications. Our main estimates use a polynomial of degree six to estimate the counterfactual loan distribution, the sample of loan size between \$75,000 and \$225,000, and a step size of \$500 when iterating through the estimation routine to find the upper limit of the excluded zone. Tables C.6 tests the sensitivity of our elasticity estimates by varying key parameters. In Table C.6, we vary the polynomial (top panel) to degree five and seven while keeping the step size constant and vary the step size while keeping the polynomial constant. The elasticity estimates are robust to the choice of polynomial and do not exhibit a specific direction of bias (smaller or larger) when we increase the polynomial degree.

We explore heterogeneity in elasticity estimates by borrower's project location and sector characteristics. Appendix Table C.7 shows that the elasticity is larger in project areas with high local bank competition relative to areas with low bank competition, where bank competition is measured using the Herfendahl-Hirschman Index (HHI). This result is consistent with the results found in Table 3.6 that shows heterogeneity by SBA market concentration. Heterogeneity by sector characteristics show that the elasticity is larger for the borrowers in high exit, services-producing, and tradable sectors relative to those in low exit, goods-producing, and nontradable sectors. One interpretation of this result is that lenders may want to bunch more when lending to borrowers that operate in risky sectors with higher chances of default or in sectors that traditionally require high initial capital. Appendix Table C.8 reports elasticity estimates by industry for the top five industries in terms of the share of SBA loans provided to a given industry relative to the total SBA loans.

Our assumption of risk neutrality does not qualitatively affect the bunching estimation procedure, but it does affect the underlying structural elasticities estimated. Under risk neutrality, a change in the reimbursement rate γ is equivalent to a change in the fee σ. However, if the banks are risk averse, a decrease in reimbursement rate γ does not linearly correspond to an increase of participation fee σ, and banks' expected subsidy rate can no longer be represented as a netted quantity $\Gamma(D_{ij})$, which drops discretely at threshold D^T. Even if the optimal D_{ij} remains unchanged, the slope of the indifference curve could potentially change–it will be steeper as banks become more risk averse. This will reduce the size of the marginal bunching loan relative to the modeled scenario with risk-neutral banks. As a consequence, we would expect more bunching for more risk averse banks. Intuitively, staying above the threshold and getting a lower guarantee rate is a greater potential loss to more risk -averse banks.

3.7 Conclusion

The efficiency of federal credit guarantees depends crucially on how responsive the lending supply is to subsidies. Specifically, the marginal change in costs per dollar of lending generated affects the elasticity of loan supply to the guarantee. This paper uses notches in SBA lending rules to provide the first empirical estimates of credit supply response to guarantees. We find that supply is responsive to loan guarantees-significantly more loans are disbursed below a threshold where guarantees are higher and we find that this bunching is stronger in years when the discontinuity in guarantee rates is greater and disappears when the discontinuity is temporarily eliminated. We translate these estimates into an elasticity of credit supply with respect to a loan guarantee of approximately five, which is on the higher end of those used in previous studies. We thus conclude that lending is highly sensitive to loan guarantees, and these programs have significant potential to increase lending levels when borrowing is inefficiently low.

While we have shown that lending supply is responsive to guarantee rates–a key parameter when considering the welfare effects of federal credit programs–important questions remain unanswered. Perhaps most importantly, the efficiency of loan guarantees ultimately rests on the efficiency of the rate of return on investments made by marginal loans and whether this is greater than the risk-free rate. Moreover, federal credit programs can have allocative effects, transferring credit from one rationed group to another. Future work should attempt to study both the allocative effects of federal credit programs and the return of loans being made under these programs.

Bibliography

Adelino, Manuel, Song Ma, and David Robinson. 2017. "Firm Age, Investment Opportunities, and Job Creation." *Journal of Finance*, 130(2): 237–264.

Atkeson, Andrew, Adrien d'Avernas, Andrea Eisfeldt, and Pierre-Olivier Weill. 2018. "Government Guarantees and the Valuation of American Banks." *NBER Macroeconomics Annual*, 33: 81–145.

Barrot, Jean-Noel. 2016. "Trade Credit and Industry Dynamics: evidence from Trucking Firms." *Journal of Finance*, 71(5): 1975–2016.

Barrot, Jean-Noel, Thorsten Martin, Julien Sauvagnat, and Boris Vallee. 2019. "Employment Effects of Alleviating Financing Frictions: worker-level Evidence from a Loan Guarantee Program." *Unpublished working Paper. HEC Paris School of Management, Bocconi University, and Harvard Business School*.

Bertrand, Marianne, Antoinette Schoar, and David Thesmar. 2007. "Banking Deregulation and Industry Structure: evidence From the French Banking Reforms of 1985." *Journal of Finance*, 62(2): 597–628.

Best, Michael, and Henrik Kleven. 2018. "Housing Market Responses to Transaction Taxes: evidence from Notches and Stimulus in the UK." *Review of Economic Studies*, 85: 157–193.

Brown, J. David, and John S. Earle. 2017. "Finance and Growth at the Firm Level: evidence from SBA Loans." *The Journal of Finance*, 72(3): 1039–1080.

Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri. 2011. "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: evidence from Danish Tax Records." *The Quarterly Journal of Economics*, 126(2): 749–804.

Darmouni, Olivier. 2017. "Estimating Informational Frictions in Sticky Relationships." *Unpublished working paper. Columbia University*.

Darmouni, Olivier, and Andrew Sutherland. 2018. "Learning about Competitors: evidence from SME Lending." *Unpublished working paper. Columbia University and MIT*.

DeFusco, Anthony, and Andrew Paciorek. 2017. “The Interest Rate Elasticity of Mortgage Demand: evidence from Bunching at the Conforming Loan Limit.” *American Economic Journal: Economic Policy*, 1(9): 210–240.

Evans, David S., and Boyan Jovanovic. 1989. “An Estimated Model of Entrepreneurial Choice under Liquidity Constraints.” *Journal of Political Economy*, 97(4): 124–149.

Fazzari, Steven, Glenn Hubbard, Bruce Petersen, Alan Blinder, and James Poterba. 1988. “Financing Constraints and Corporate Investment.” *Brookings Papers on Economic Activity*, 1988: 141–206.

Fieldhouse, Andrew J. 2018. “Crowd-out Effects of U.S. Housing Credit Policy.” *Unpublished working Paper. Cornell University*.

Gale, William. 1990. “Federal Lending and the Market for Credit.” *Journal of Public Economics*, 42: 177–93.

Gale, William. 1991. “Economic Effects of Federal Credit Programs.” *American Economic Review*, 81(1): 133–52.

Gonzalez-Uribe, Juanita, and Su Wang. 2019. “Dissecting the Effect of Financial Constraints on Small Firms.” *Unpublished working paper. London School of Economics and University of Amsterdam*.

Granja, Joao, Christian Leuz, and Raghuram Rajan. 2018. “Going the Extra Mile: distant Lending and Credit Cycles.” *Unpublished working paper. National Bureau of Economics Research*, , (25196).

Greenwald, Bruce, and Joseph Stiglitz. 1986. “Externalities in Economies with Imperfect Information and Incomplete Markets.” *The Quarterly Journal of Economics*, 101(2): 229–264.

Hanson, Samuel G., David S. Scharfstein, and Adi Sunderam. 2018. “Social Risk, Fiscal Risk, and the Portfolio of Government Programs.” *Review of Financial Studies*, 32: 2341–2382.

Hendren, Nathaniel. 2014. “Efficient Welfare Weights.” *Unpublished working paper. National Bureau of Economic Research*.

Hendren, Nathaniel. 2016. “The Policy Elasticity.” *Tax Policy and the Economy*, 30(1): 51–89.

Herranz, Neus, Stefan Krasa, and Anne Villamil. 2015. “Entrepreneurs, Risk Aversion and Dynamic Firms.” *Journal of Political Economy*, 123(5): 1133–1176.

Howell, Sabrina. 2017. “Financing Innovation: evidence from R&D Grants.” *American Economic Review*, 4(107): 1136–64.

Kaplan, Steven, and Luigi Zingales. 1997. “Do Investment-Cash Flow Sensitivities Provide Useful Measures of Financing Constraints?” *Quarterly Journal of Economics*, 112(1): 169–215.

Kelly, Bryan, Hanno Lustig, and Stijn Van Nieuwerburgh. 2016. "Too-Systemic-to-Fail: what Option Markets Imply about Sector-Wide Government Guarantees." *American Economic Review*, 106(6): 1278–1319.

Kerr, William R., and Ramana Nanda. 2010. "Democratizing Entry: banking Deregulations, Financing Constraints, and Entrepreneurship." *Journal of Financial Economics*, 94(1): 124–149.

Kleven, Henrik. 2016. "Bunching." *Annual Review of Economics*, 8: 435–464.

Kleven, Henrik J., and Mazhar Waseem. 2013. "Using Notches to Uncover Optimization Frictions and Structural Elasticities: theory and Evidence from Pakistan." *Quarterly Journal of Economics*, 128(2): 669–723.

La Porta, Rafael, Florencio Lopez de Silanes, and Andrei Shleifer. 2002. "Government Ownership of Banks." *Journal of Finance*, 57(1): 265–301.

Lelarge, Clair, David Sraer, and David Thesmar. 2010. "Entrepreneurship and credit constraints: evidence from a French loan guarantee program. Lerner, J., Schoar, A. (Eds.), International Differences in Entrepreneurship." , ed. J. Lerner and A. Schoar, 243–273. University of Chicago Press, Chicago.

Lucas, Deborah. 2016. "Credit Policy as Fiscal Policy." *Brookings Papers on Economic Activity*, 47(Spring): 1–57.

Mankiw, N Gregory. 1986. "The Allocation of Credit and Financial Collapse." *The Quarterly Journal of Economics*, , (101): 455–70.

Mullins, William, and Patricio Toro. 2017. "Credit Guarantees and New Bank Relationships." *Unpublished working paper. UC San Diego and Banco Central de Chile*, 8241.

Petersen, Mitchell A., and Raghuram G. Rajan. 1994. "The Benefits of Lending Relationships: evidence from Small Business Data." *The Journal of Finance*, 49(1): 3–37.

Petersen, Mitchell A., and Raghuram G. Rajan. 1995. "The Effect of Credit Market Competition on Lending Relationships." *The Quarterly Journal of Economics*, 110(2): 3–37.

Rauh, Joshua. 2006. "Investment and Financing Constraints: evidence from the Funding of Corporate Pension Plans." *Journal of Finance*, 61(1): 33–71.

Saez, Emmanuel. 2010. "Do Taxpayers Bunch at Kink Points?" *American Economic Journal: Economic Policy*, 2(3): 180–212.

Smith, Bruce. 1983. "Limited Information, Credit Rationing, and Optimal Government Lending Policy." *American Economic Review*, 73(3): 305–18.

Stiglitz, Joseph, and Andrew Weiss. 1981. "Credit Rationing in Markets with Imperfect Information." *American Economic Review*, 71: 393–410.

Whited, Toni, and Guojun Wu. 2006. "Financial Constraints Risk." *Review of Financial Studies*, 19(2): 531–559.

Figures and Tables

Figure 3-1: Changes in bank profit and lending by elasticity

Notes: Here we simulate what happens to bank profits and additional dollar lending when the reimbursement rate changes from $\gamma = 0.5$ to 0.8. The top figure plots the change in bank profit retained by the lender as a fraction of the change in the guarantee subsidy cost, S. The figure illustrates that if the elasticity of loan size to guarantee rate is inelastic, the share of guarantee subsidy retained by the lender is high and this share declines as the elasticity increases. The bottom figure plots the change in additional lending as a fraction of the change in the guarantee subsidy cost over lending elasticities. As the lending supply becomes more responsive to the guarantee, the expected costs of the guarantee and the net-of-guarantee losses for the lender increase. Therefore, loan size increases as the elasticity increases, while the guarantee subsidy retained by the lender decreases.

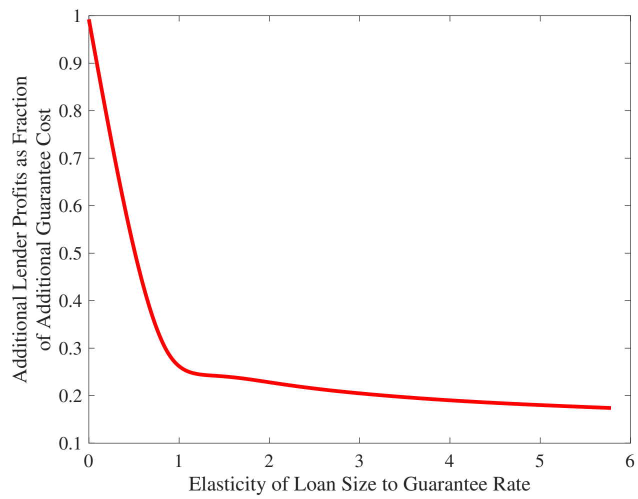

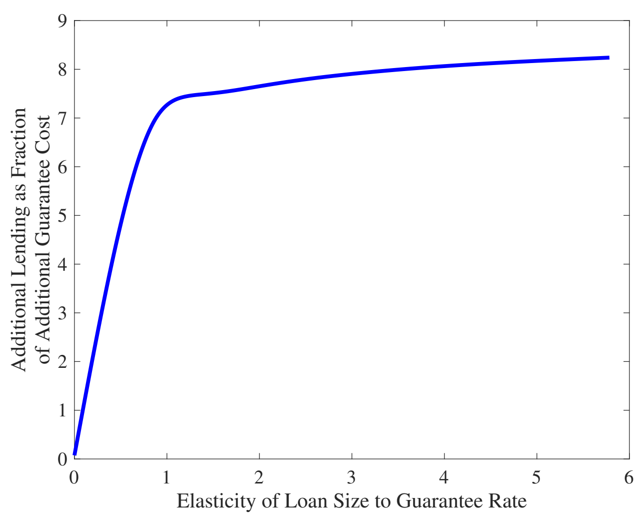

Figure 3-2: Bank profit as a function of D for a given observable type r

Notes: This figure illustrates the shape and the slope of a bank's profit as a function of loan size D for a given productivity type r. Panel a illustrates a concave profit function in the absence of a notch where the location of the global optimum D^* is indicated by a red dot. Panel b shows a new profit function when a notch point is introduced at D^T. Panel c shows a profit function for a marginal buncher as well as an additional profit function with a steeper slope (in blue). The shape of the profit function is determined by the distribution of realized productivity types, $F(r, n)$ and the interest rate $\bar{R}$. Panel d shows a profit function of an unimpacted borrower.

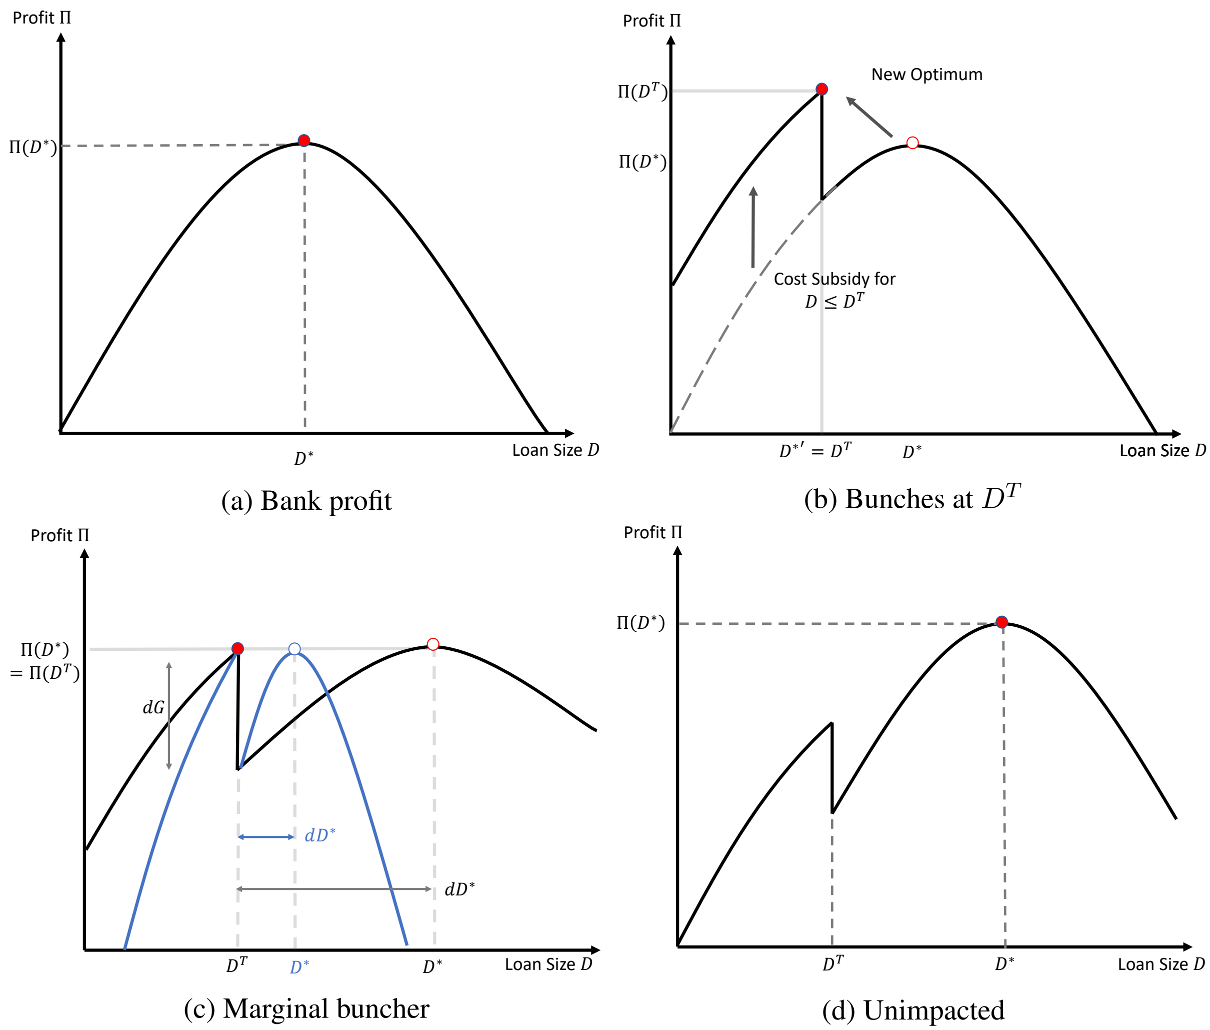

(a) Bank profit

(b) Bunches at D^T

(c) Marginal buncher

(d) Unimpacted

Figure 3-3: Bunching at the guarantee notch in 2013 and 2015

Notes: The left panel shows the number of loans made in discrete $2,000 bins made in 2013 (red) and 2015 (black). The right panel shows the change in the guarantee rate at the threshold in these two years. In 2015, when the change in the guarantee at the threshold was much larger than in 2013, there was substantially more excess mass at the threshold. Source: SBA.

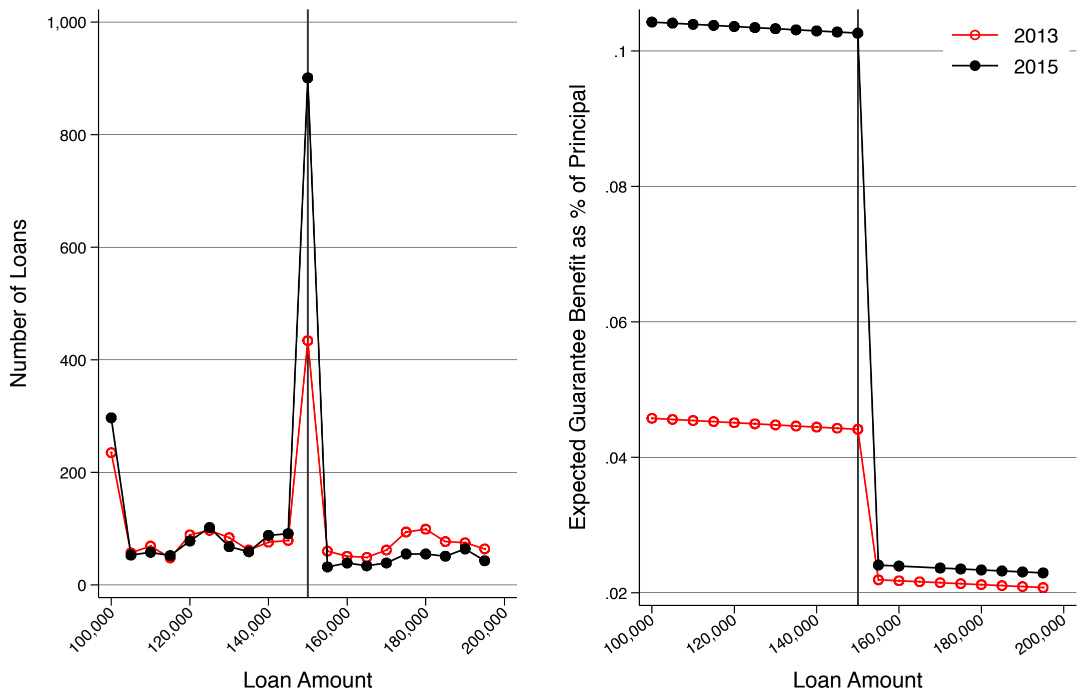

Figure 3-4: Guarantees and fees by loan amount

Notes: The top panel shows the average expected guarantee benefit as a percentage of the loan principal amount for discrete $2,000 bins across the threshold. This net benefit is calculated as the guaranteed reimbursement on expected losses minus guarantee fees. The bottom panel shows the number of loans made in discrete $2,000 bins across the threshold. The figures pool over all years 2008-2017. Source: SBA.

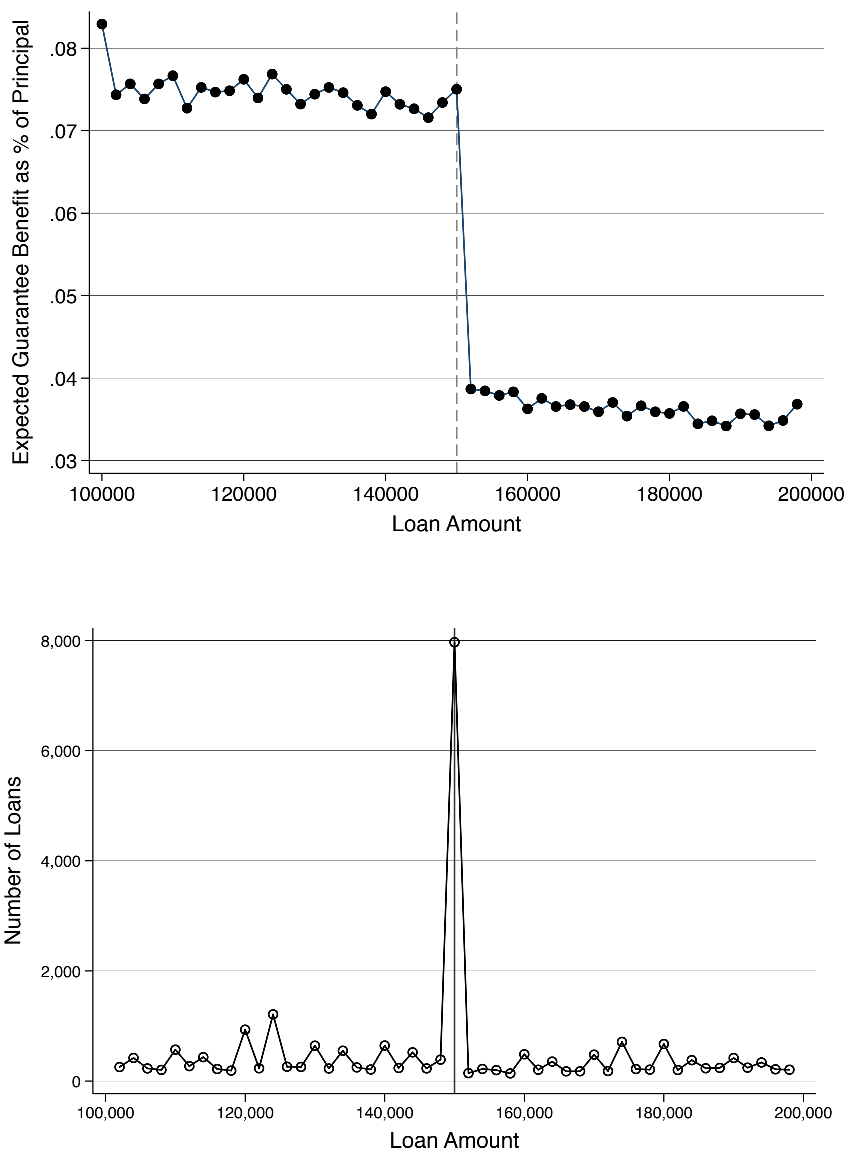

Figure 3-5: Observed and counterfactual distributions

Notes: This figure shows the observed and counterfactual density of loans. The solid line shows the observed number of loans, in \$5,000 bins. The dashed line shows the counterfactual number of loans. The counterfactual is estimated for each notch separately by fitting a sixth-order polynomial with round number fixed effects to the empirical distribution using step size of 5,000 and excluding data around the notch, as specified in Eq. 3.15. The dotted vertical lines mark the estimated excluded range $[d_L, d_U]$, where d_U is the location of the marginal buncher. We calculate ΔD as $d_U - d_L$. Source: SBA.

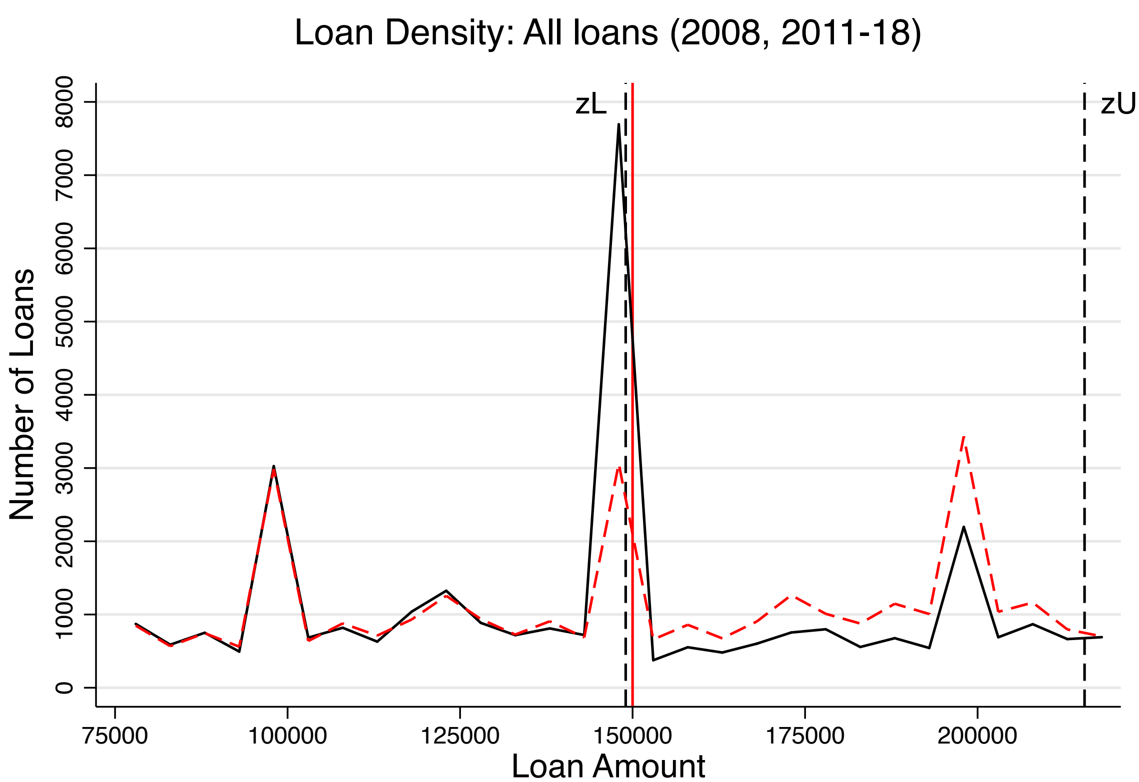

Figure 3-6: Bunching at the guarantee notch by year

Notes: This figure shows the fraction of loans made in discrete $2,000 bins across the threshold by year. We divide the loans by years when the notch was either positive and above (high) or below (low) the median, or are nonexistent. Source: SBA.

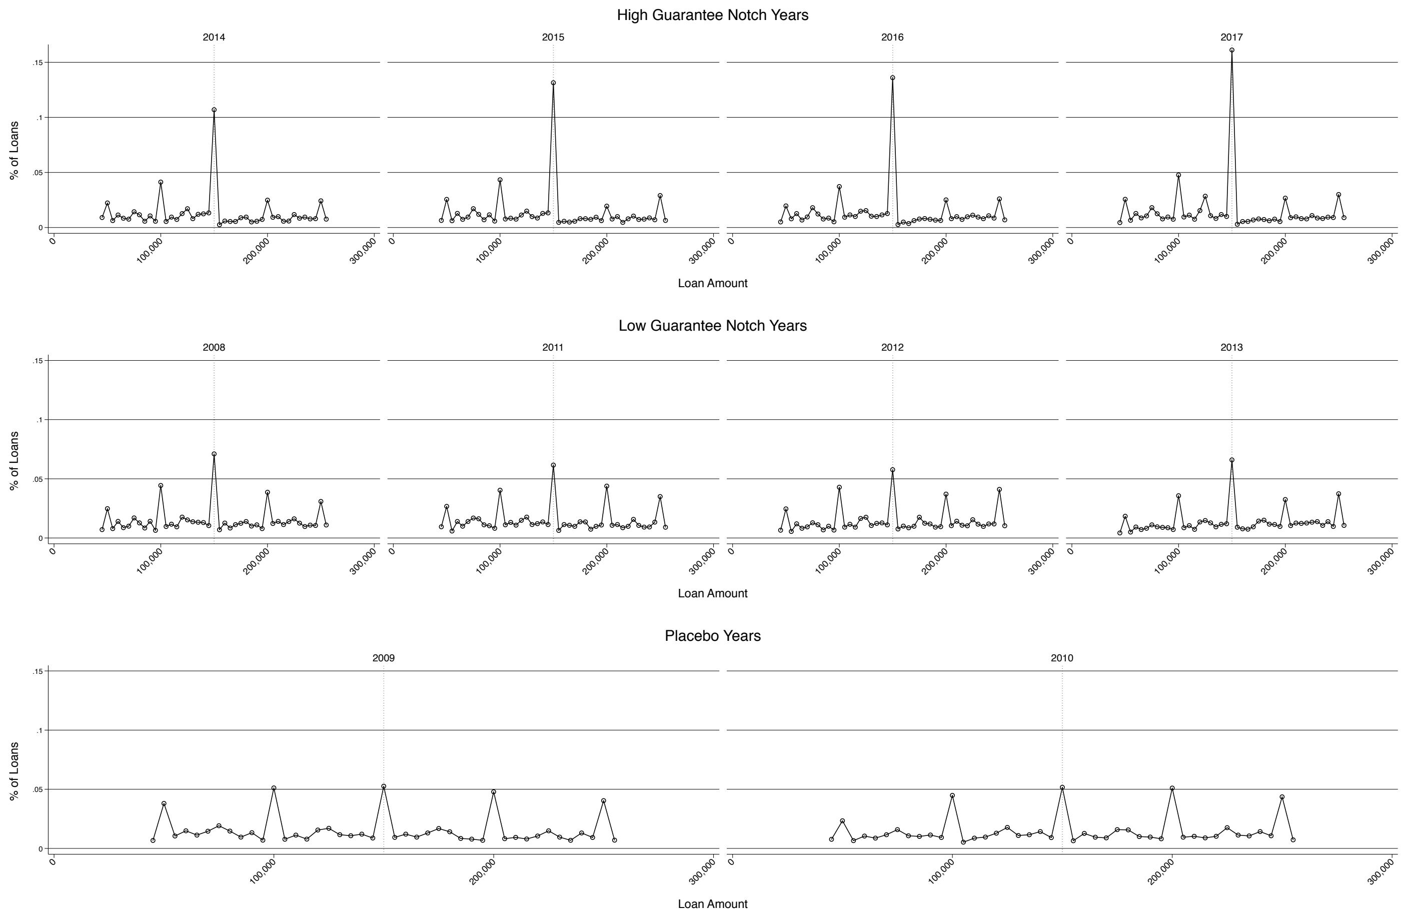

Figure 3-7: Relation between size of notch and excess mass

Notes: The top figure plots the share of excess mass against the size of the guarantee rate change at the $150,000 threshold. The excess mass at the $150,000 threshold is measured as the difference in the percentage of loans at the threshold relative to other round numbers. The share of excess mass is the estimated excess mass as a share of the total number of loans in the estimation range. The change in the guarantee rate is the change in the average expected guarantee benefit as a percentage of the loan principal. The bottom figure plots the share of excess mass and the size of the guarantee rate change at the threshold over time to show the tight correlation between the two measures. Both figures show that there is a positive correlation between the incentive to bunch (the size of the guarantee rate change) and the amount of bunching. Both graphs pool over all years 2008-2017 and control for bank fixed effects. Source: SBA.

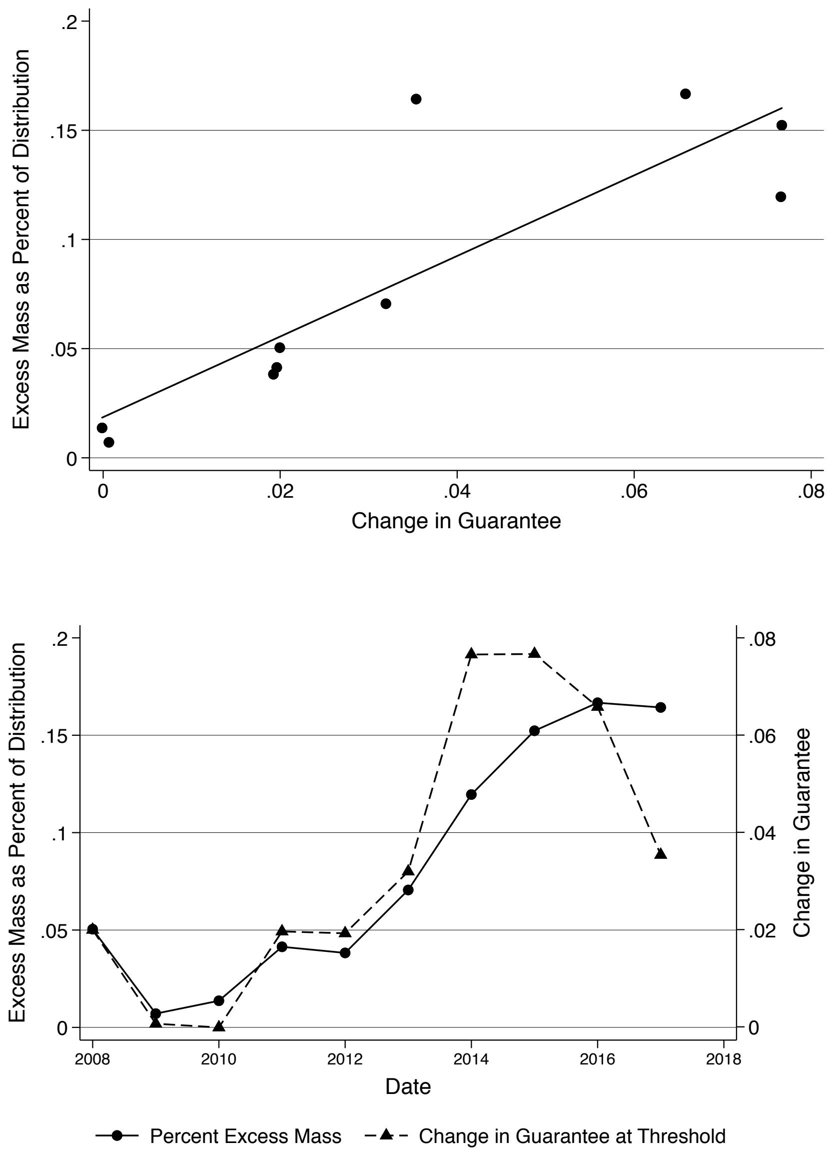

Figure 3-8: Risk-shifting: charge-offs at notch

Notes: The top panel plots the coefficient ξ_t on the interaction terms between whether a loan is at the guarantee notch and year indicators in Eq. 3.17, along with a 95% confidence interval. The dependent variable is an indicator of whether a loan is charged off. The baseline is 2017. The bottom panel shows the magnitude of the expected benefit at the guarantee notch. Source: SBA.

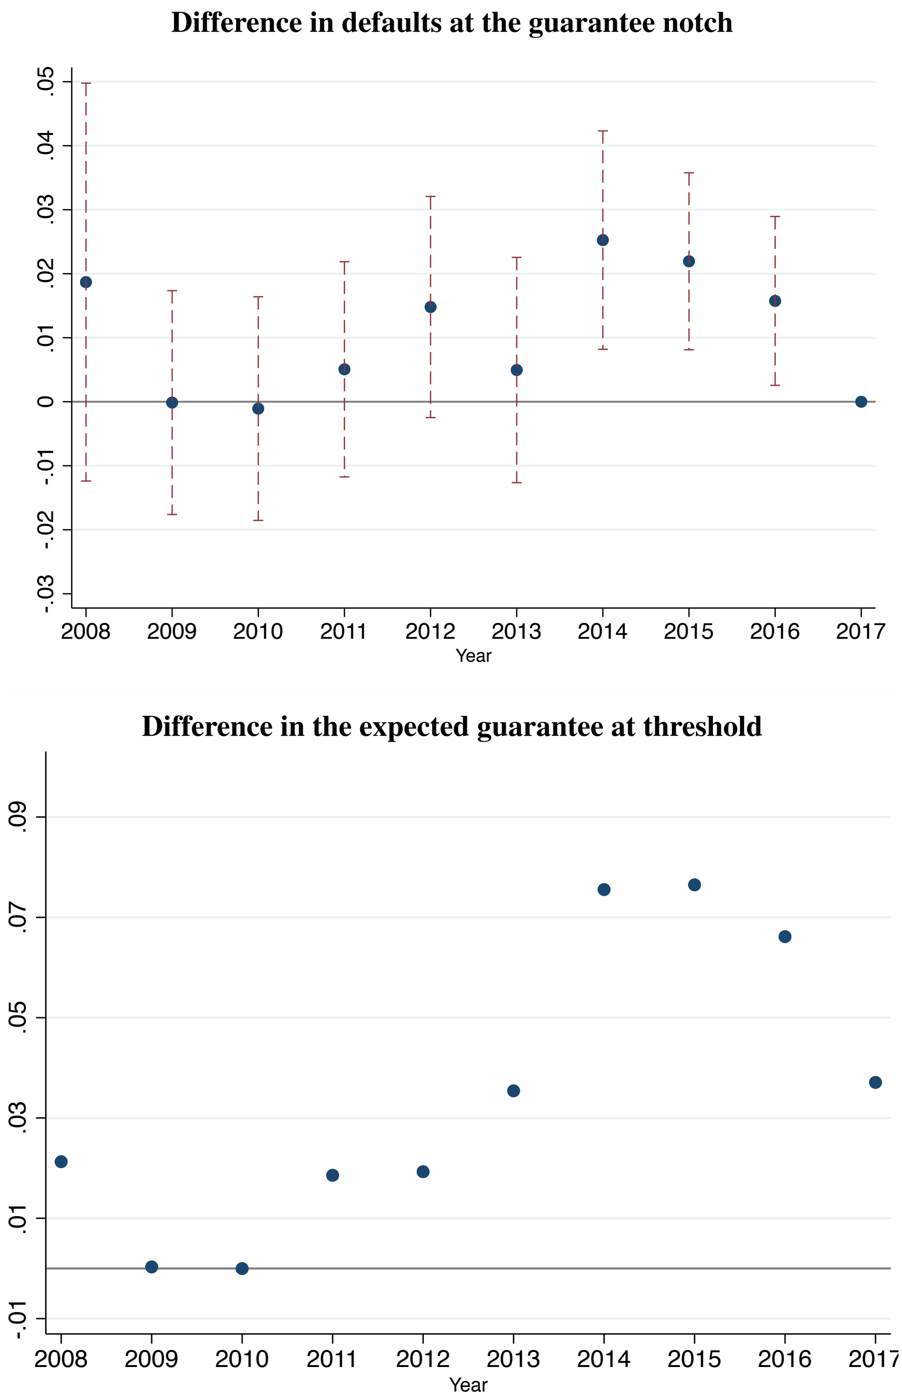

Figure 3-9: Average interest rate and guarantee rate across the threshold

Notes: This figure shows interest rates and guarantee rates in discrete $2,000 bins across the threshold. While the guarantee rate drops dramatically at the threshold, the interest rate remains flat. The guarantee rate is the average expected guarantee benefit as a percentage of the loan principal. The graph pools over all years 2008-2017. Source: SBA.

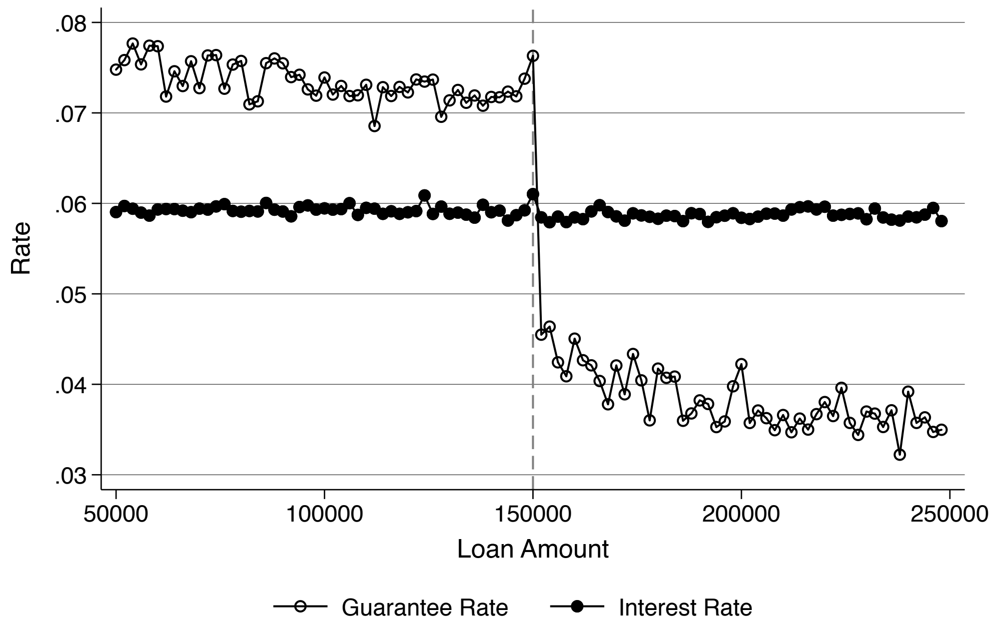

Figure 3-10: Estimates varying estimation range

Notes: This figure shows elasticity estimates varying the starting and ending range around the $150,000 threshold. A sixth-order polynomial is used. The bars show a 95% confidence interval. Standard errors are computed using a bootstrap. Source: SBA.

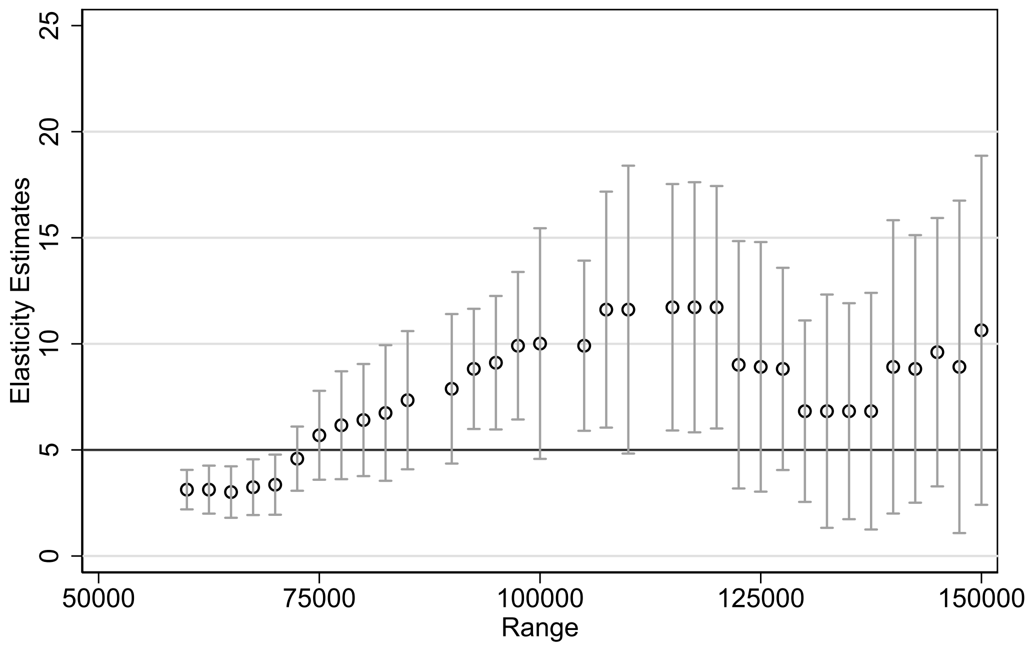

Table 3.1: Summary statistics

Notes: This table shows summary statistics for the main analysis variables. The first two columns report the mean and the standard deviation, and the third to fifth columns report the 25th, median, and the 75th percentile, respectively. Panel A reports summary statistics for full sample, and panel B reports statistics for the sample of loans used in the notch estimation (loan size between $75,000 and $225,000). Loan amount is the loan size. Reimbursement rate refers to the SBA determined reimbursement rate pooling across all years in the sample (2008-2017). Reimbursed amount is the guaranteed portion of the loan balance. Interest rate is the total interest rate (base plus spread) at the time of loan origination. Maturity is the length of loan terms, and charge-off amount is the total loan balance charged off, including the guaranteed and non-guaranteed portion of loan. Expected guarantee subsidy (Γ) is the predicted guarantee amount as a share of loan principal net of one-time and yearly fees. A full description of the methodology we use to estimate Γ can be found in Appendix C.2. Loans per firm-lender pair reports the number of loans that a given firm borrows from the same lender in the same year. The excess mass reports an estimate of the amount of excess mass ($\hat{B}$) at the 150k notch, which we measure as the difference between the observed and counterfactual bin counts in the excluded region at and to the left of the notch. The estimate is reported as the share of bunching relative to the total number of loans in the estimation range. Excess mass is only reported in panel B, as it is estimated using the notch sample only. Source: SBA.

Outcome	Mean	Std. dev.	25^{th} Pctile.	Median	75^{th} Pctile.
			A. Full sample		
Loan amount ($)	746,107	826,485	215,000	460,000	950,000
Reimbursement rate	80	6	75	75	85
Reimbursed amount ($)	574,195	626,519	168,750	356,400	735,000
Interest rate (%)	5.73	0.74	5.25	5.96	6.00
Maturity (in years)	15	8	10	10	25
Charge-off amount ($)	11,706	85,383	0	0	0
Expected guarantee subsidy (Γ)	0.02	0.04	0	0.02	0.05
Loans per firm-lender pair	1.05	0.27	1.00	1.00	1.00
Observations	199,013	199,013	199,013	199,013	199,013
			B. Sample for notch estimation		
Loan amount ($)	147,359	41,330	112,000	150,000	180,000
Reimbursement rate	84	5	85	85	90
Reimbursed amount ($)	120,575	31,354	93,750	127,500	141,110
Interest rate (%)	6	1	6	6	6
Maturity (in years)	10	5	7	10	10
Charge-off amount ($)	6,221	26,704	0	0	0
Expected guarantee subsidy (Γ)	0.06	0.03	0.03	0.06	0.07
Share of excess mass	0.08	0.06	0.04	0.05	0.16
Loans per firm-lender pair	1.03	0.27	1.00	1.00	1.00
Observations	41,460	41,460	41,460	41,460	41,460

Table 3.2: Guarantees and fees by loan amount

Notes: This table includes fees and guarantee rates for loans with maturities longer than 12 months. Fees are calculated as a percentage of the loan principal. The reimbursement rate is expressed as a percentage of charged-off principal. The net benefit combines the fees and reimbursement rate to measure the average expected generosity of the guarantee and is expressed as a percentage of the loan principal amount. This net benefit is calculated as the guaranteed reimbursement on expected losses minus guarantee fees. Loan amounts smaller than $150,000 refers to loans between $0-150,000. Loan amounts larger than $150,000 refers to loans between $150,000-700,000. Source: SBA.

	Loan amount smaller than $150,000				Loan amount larger than $150,000			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fiscal year	Yearly fee	One time dee	Reimbursement rate	Net benefit	Yearly fee	One time fee	Reimbursement rate	Net benefit
2008	0.55	2	85	4.6	0.55	3.42	75	2.6
2009	0.55	0	90	7.4	0.55	0	90	7.4
2010	0.55	0	90	7.4	0.55	0	90	7.4
2011	0.55	2	85	4.9	0.55	3.42	75	2.9
2012	0.55	2	85	4.6	0.55	3.42	75	2.7
2013	0.55	2	85	5.8	0.55	3.42	75	2.6
2014	0	0	85	10.5	0.52	3.42	75	2.9
2015	0	0	85	10.5	0.52	3.42	75	2.9
2016	0	0	85	9.6	0.47	3.42	75	2.9
2017	0.55	0	85	6.3	0.55	3.42	75	2.7

Table 3.3: Excess mass and elasticity estimates

This table reports estimates of excess mass and the main elasticity estimates. The top panel shows placebo years (2009 and 2010) where there was no change in the reimbursement rate at the 150,000 threshold. The bottom panel shows years where a notch existed (2008, 2011-2017). Elasticity estimates are reported in the latter sample. For estimation, we restrict the loan sample with size between \$75,000 to \$225,000, use the step size of 500, and include round number dummies for multiples of 1,5, 10, 25, and 50 thousand. The polynomial used is denoted in the second column. The change in the guarantee rate ($\Delta\Gamma$) at the threshold for years in which a notch existed is computed as the weighted average of the average expected guarantee benefit as a percentage of the loan principal, where the weights correspond to the number of loans across years 2008, 2011-2017. Standard errors are reported in italics and are obtained by empirical bootstrap with 1,000 repetitions of resampling the distribution of loans made. The bunching estimation routine is run at every bootstrap iteration until convergence. Source: SBA.

Year	Polynomial	Excess mass	ΔD	$\Delta\Gamma$	Elasticity
		A. Placebo years - no notch			
2009-2010	5	67	21,000	-	-
		(21.57)	*(14,796)*	-	-
	6	66	21,000	-	-
		(41.36)	*(14,285)*	-	-
	7	0	9,500	-	-
		(16.76)	*(13,131)*	-	-
		B. Pooled years - with notch			
2008, 2011-2017	5	4,744	66,500	0.038	4.519
		(98.9)	*(1,326)*	-	*(0.186)*
	6	4,747	66,000	0.038	4.589
		(44.38)	*(2,806)*	-	*(0.395)*
	7	4,745	70,500	0.038	5.235
		(102.6)	*(1,240)*	-	*(0.181)*

Table 3.4: Excess mass and elasticity estimates, by Year

This table shows elasticities for years in which a notch existed and estimates of the excess mass for the two years (2009 and 2010) in which there was no change in the guarantee rate at the 150,000 threshold. For this estimation, the stepsize = 500, the range was limited to 75,000-225,000, we included round number dummies for multiples of 1,5, 10, 25, and 50 thousand, and we used a polynomial of degree 6. The change in the guarantee rate ($\Delta\Gamma$) at the threshold for years in which a notch existed is computed as the weighted average of the average expected guarantee benefit as a percentage of the loan principal, where the weights correspond to the number of loans across years 2008, 2011-2017. Source: SBA.

Year	Excess mass	ΔD	$\Delta\Gamma$	Elasticity
Placebo years - no notch				
2009	19.12	2,500	0	NA
2010	35.02	6,000	0	NA
Years with notch				
2008	248.39	52,000	0.02	5.32
2011	151.81	40,500	0.02	3.36
2012	132.64	60,500	0.02	7.62
2013	199.91	71,500	0.03	6.41
2014	233.02	62,000	0.08	2.01
2015	457.83	55,500	0.08	1.61
2016	564.04	60,500	0.07	2.24
2017	1,386.12	69,500	0.04	5.47

Table 3.5: Effect of notch size on loan charge-off

This table reports estimates of the notch size on charge-off probabilities. The column titles report the dependent variables. (D $> D^T$) equals one if a loan size is greater than $150,000. Γ is the notch size. (D $= D^T$) indicator equals one if the size of the loan is $150,000. The sample is restricted to loan size between $50,000 and $225,000. Standard errors are reported in parentheses. The inclusion of fixed effects and controls is denoted beneath each specification. Source: SBA.

	Probability of charge-off		Percent charged-off		Log of charged-off amount	
	(1)	(2)	(3)	(4)	(5)	(6)
(D $> D^T$)	-.008	-.002	-.005	-.002	-.085	-.037
	(.004)	(.012)	(.003)	(.010)	(.047)	(.142)
(D $> D^T$)$\times$ Γ	.034	.100	-.015	.016	.904	1.54
	(.102)	(.083)	(.066)	(.063)	(1.15)	(.938)
(D = D^T)$\times$ Γ	.184	.317	.172	.251	2.39	3.70
	(.036)	(.092)	(.026)	(.070)	(.406)	(1.04)
Year FEs	X		X		X	
Size bin controls		X		X		X
Bank FEs		X		X		X
Year-month FEs		X		X		X
Maturity FEs		X		X		X
Number of observations	40,751	40,751	40,751	40,751	40,751	40,751

Table 3.6: Estimates split by number of banks and SBA share

This table reports estimates of excess mass and the main elasticity estimates by subsamples of loans that are and are not originated by lenders that specialize in SBA lending (top two panels) and by subsamples of loans originated in counties with more or less operating lenders (bottom two panels). The share of SBA lending in the top two panels is computed from calculating the share of SBA loans relative to a lender's overall small business lending. The estimation restricts the sample to loans to be of size between \$75,000 to \$225,000, uses the step size of 500, and includes round number dummies for multiples of 1,5, 10, 25, and 50 thousand. The degree of the polynomial used in the estimation is denoted in the second column. The change in the guarantee rate ($\Delta\Gamma$) at the threshold for years in which a notch existed is computed as the weighted average of the average expected guarantee benefit as a percentage of the loan principal, where the weights correspond to the number of loans across years 2008, 2011-2017. Standard errors are shown in parentheses. Source: SBA and FDIC SDI.

Year	Polynomial	Excess mass	ΔD	$\Delta\Gamma$	Elasticity	Excess mass	ΔD	$\Delta\Gamma$	Elasticity
		SBA share $> 60\%$				SBA share $\leq 60\%$			
2008, 2011-2017	6	2,336	69,500	0.038	5.087	2,413	55,500	0.038	3.244
		(20.63)	(3,286)	–	(0.466)	(40.18)	(7,254)	–	(0.942)
	No. obs.	8,931				24,333			
		SBA share $> 80\%$				SBA share $\leq 80\%$			
2008, 2011-2017	6	2,231	69,500	0.038	5.087	2,518	55,500	0.038	3.244
		(20.41)	(3,299)	–	(0.465)	(41.04)	(7,489)	–	(0.975)
	No. obs.	7,958				25,306			
		Unique banks > 3				Unique banks ≤ 3			
2008, 2011-2017	6	4,363	66,000	0.038	4.589	383	53,500	0.038	3.015
		(46.23)	(6,149)	–	(0.852)	(9.18)	(6,889)	–	(0.887)
	No. obs.	28,851				4,413			
		Unique banks > 7				Unique banks ≤ 7			
2008, 2011-2017	6	3,931	70,000	0.038	5.161	818	60,500	0.038	3.855
		(37.86)	(5,819)	–	(0.812)	(17.00)	(7,858)	–	(1.031)
	No. obs.	24,509				8,755			

Appendix A

Appendix for Chapter 1

A.1 Public vs. Private Goods

To illustrate how empowering secondary earner with more borrowing capacity affects private vs. public consumption, consider a simple Cobb-Douglas individual preferences. Then for each spouse $i \in (P, S)$:

$$u^i(x^i, Q) = \sum_k \alpha_k^i log q_k^i + \sum_j \delta_j^i log Q_j \tag{A.1}$$

where the coefficients are positive and normalized by $\sum_k \alpha_k^s + \sum_j \delta_j^s = 1$. Let the relative decision power of two spouses $\mu = \frac{\theta^S}{\theta_P}$ denote S' Pareto weight. Prices are normalized to 1, so that the BC is simply:

$$\sum_k (q_k^a + q_k^b) + \sum_j Q_j = x \tag{A.2}$$

where the LHS is consumption of the two members: $c^P + c^S = \sum_k (q_k^P + q_k^S) + \sum_j Q_j$, and the RHS is $x = Y - a$. Household problem:

$$\mathcal{L} = \Big(\sum_k \alpha_k^P log q_k^P + \sum_j \delta_j^P log Q_j \Big) + \mu \Big(\sum_k \alpha_k^S log q_k^S + \sum_j \delta_j^S log Q_j \Big) + \tag{A.3}$$

$$\lambda \Big(x - \sum_j Q_j - \sum_k (q_k^P + q_k^S) \Big) \tag{A.4}$$

FOCs:

$$\lambda q_k^P = \alpha_k^P \tag{A.5}$$

$$\lambda q_k^S = \mu \alpha_k^S \tag{A.6}$$

$$\lambda Q_j = \delta_j^P + \mu \delta_j^S \tag{A.7}$$

Sum all of these terms:

$$\lambda \Big(\sum_k (q_k^P + q_k^S) + \sum_j Q_j \Big) = (\sum_k \alpha_k^P + \sum_j \delta_j^P) + \mu (\sum_k \alpha_k^S + \sum_j \delta_j^S) \tag{A.8}$$

$$\lambda x = 1 + \mu \tag{A.9}$$

And $\lambda = \frac{1+\mu}{x}$. Thus, household demands:

$$q_k^P = \frac{\alpha_k^P}{1+\mu} x \tag{A.10}$$

$$q_k^S = \frac{\mu \alpha_k^S}{1+\mu} x \tag{A.11}$$

$$Q_j = \frac{\delta_j^P + \mu \delta_j^S}{1+\mu} x \tag{A.12}$$

And the following conditions hold:

$$\frac{\partial q_k^P}{\partial \mu} = -\frac{\alpha_k^P}{(1+\mu)^2}x \tag{A.13}$$

$$\frac{\partial q_k^S}{\partial \mu} = \frac{\alpha_k^S}{(1+\mu)^2}x \tag{A.14}$$

$$\frac{\partial Q_j}{\partial \mu} = \frac{\delta_j^S - \delta_j^P}{(1+\mu)^2}x \tag{A.15}$$

These imply:

1. The private consumptions of P are decreasing in μ.

2. The private consumptions of S are increasing in μ.

3. Household consumption in public commodity j increases iff S "cares more" about commodity than P does, in the sense that $\delta_j^P > \delta_j^S$.

It's natural to interpret these terms of marginal willingness to pay. These are given for any public good j by:

$$MWP_j^i = \delta_j^i \frac{\sum_k q_k^i}{Q_j} \tag{A.16}$$

where $\sum_k q_k^i$ is the conditional sharing rule, or the amount of private consumptions consumed taking public consumption as given. This is the maximum amount i would be willing to pay to acquire an additional unit of consumption good j, if the amount was to be withdrawn from i's consumption of private good.

Note that the condition $\delta_j^S > \delta_j^P$ is *not* equivalent to S' MWP being larger than P's. It implies that

$$\frac{\partial MWP_j^S}{\partial \sum_k q_k^P} > \frac{\partial MWP_j^P}{\partial \sum_k q_k^S} \tag{A.17}$$

The MWP of S must be more income sensitive than that of P.

A.2 The Limited Commitment Model Prediction

Consider the properties of an efficient self-enforcing consumption maths when spouses' participation constraints bind but they stay married:

$$V_t^{P,M}(\omega_{\mathbf{t}}) = u\Big(c_t^{P*}(\omega_{\mathbf{t}})\Big) + \beta E\Big[V_{t+1}^{P,M}(\omega_{\mathbf{t+1}}|\omega_{\mathbf{t}})\Big] \tag{A.18}$$

$$V_t^{S,M}(\omega_{\mathbf{t}}) = u\Big(c_t^{S*}(\omega_{\mathbf{t}})\Big) + \beta E\Big[V_{t+1}^{S,M}(\omega_{\mathbf{t+1}}|\omega_{\mathbf{t}})\Big] \tag{A.19}$$

$$V_t^{P,M}(\omega_{\mathbf{t}}) \geq V_t^{P,D}(\omega_{\mathbf{t}}) \tag{A.20}$$

$$V_t^{S,M}(\omega_{\mathbf{t}}) \geq V_t^{S,D}(\omega_{\mathbf{t}}) \tag{A.21}$$

This problem can be reformulated as a Lagrangian problem. The couple solves:

$$\begin{aligned}\mathcal{L}^{*,M} =& \theta_t^P u\Big(c_t^{P*}(\omega_{\mathbf{t}})\Big) + \theta_t^S u\Big(c_t^{S*}(\omega_{\mathbf{t}})\Big) + \beta E\Big[V_{t+1}^{M}(\omega_{\mathbf{t+1}}|\omega_{\mathbf{t}})\Big] \\ &+ \lambda_t^P\Big\{u\Big(c_t^{P*}(\omega_{\mathbf{t}})\Big) + \beta E\Big[V_{t+1}^{P,M}(\omega_{\mathbf{t+1}}|\omega_{\mathbf{t}})\Big] - V_t^{P,D}(\omega_{\mathbf{t}})\Big\} \\ &+ \lambda_t^S\Big\{u\Big(c_t^{S*}(\omega_{\mathbf{t}})\Big) + \beta E\Big[V_{t+1}^{S,M}(\omega_{\mathbf{t+1}}|\omega_{\mathbf{t}})\Big] - V_t^{S,D}(\omega_{\mathbf{t}})\Big\}\end{aligned}$$

where $V_{t+1}^M(\omega_{\mathbf{t+1}}) = \theta_{t+1}^P V_{t+1}^{P,M}(\omega_{t+1}) + \theta_{t+1}^S V_{t+1}^{S,M}(\omega_{t+1})$, and λ_t^P and λ_t^S are the Lagrangian multiplier associated with each spouse's sequential participation constraint.

Combining the first order condition with respect to c_t^{P*} and c_t^{S*} leads to the key prediction of this model that the ratio of the marginal utilities of consumption has one-to-one relationship to the slope of the Pareto frontier (γ_t):

$$\begin{aligned}\frac{u'\Big(c_t^{P*}\Big)}{u'\Big(c_t^{S*}\Big)} &= \frac{\theta_t^S + \lambda_t^S}{\theta_t^P + \lambda_t^P} \\ &= \gamma_t\end{aligned}$$

In other words, the couple's consumption allocation in the household is determined by the slope of the Pareto frontier, which can be entirely characterized by the spouses' bargaining power. Figure A-14 illustrates the economic intuition graphically. Each panel plots the primary earner's expected lifetime value of staying married (y-axis) against the secondary earner's expected lifetime utility of staying married (x-axis). The red dashed lines denote the spouse's outside option, and the first quadrant of ellipse represents the Pareto frontier. Any consumption allocation along this Pareto frontier is a feasible allocation, but the position on this Pareto frontier (red dot) is determined by the ratio of the marginal utilities of consumption.

I discuss two cases: when the participation constraint does and does not bind. First, consider the case when the secondary earner's participation constraint does not bind. This case is illustrated in panel a by the fact that the existing resource allocation in period 1, $E[V_1^{i,*}]$, sits in the non-negative orthant created by spouses' best outside options. In this case, the improvement in secondary earner's outside option expands the Pareto frontier and shifts the location of efficient resource allocation outward. However, since secondary earner's participation constraint does not bind, the couple continues the initial resource allocation plan. This is shown in panel b– the slope of the Pareto frontier is unchanged in period 2. Note that the value of the spouses' best outside options intersect in the interior of the Pareto frontier, implying that there is still gains from marriage even after the change in secondary earner's outside option.

Now, consider the case when the secondary earner's participation constraint binds such that the value of her outside option expands to the point where the initial resource allocation plan no longer sits in the non-negative orthant created by spouses' best outside options. The binding constraint triggers bargaining between spouses and increases the secondary earner's decision power by λ^S. This is shown in panel c. The improvement in the secondary earner's decision power makes the slope of the Pareto frontier steeper by tilting resource allocation toward her and thus reducing her marginal utility, $u'(c_2^{S,\omega})$. Figure d shows that this moves the location of resource allocation plan along the Pareto frontier to the new point, $E[V_2^{i,**}]$, where the secondary earner is indifferent from staying married with the new allocation plan or divorcing to take her outside option.

Comparing the ratio of marginal utilities in the case when secondary earner's outside option does not bind, $\frac{u'(c_2^{p,\omega})}{u'(c_2^{s,\omega})} = \frac{\theta^S}{\theta_P} = \widehat{\gamma}_t$, to the case when it does bind, $\frac{u'(c_2^{p,\omega})}{u'(c_2^{s,\omega})} = \frac{\theta^S+\lambda^S}{\theta_P} = \widetilde{\gamma}_t$, reveals how bargaining power determine the consumption allocation in the household. If $\lambda^S > 0$, then $\widetilde{\gamma}_t > \widehat{\gamma}_t$. This is only possible when consumption allocated to secondary earners increases.

A.3 Appendix Figures and Tables

Figure A-1: Timeline

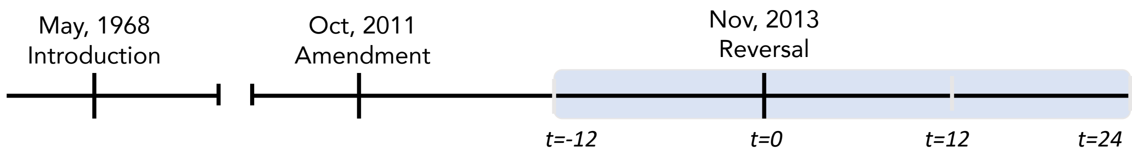

Notes: This figure shows the timeline of the changes to the Truth-In-Lending Act (TILA) Section 150. The area highlighted in blue – 12 months before and 24 months after the reversal – denote my sample period. This study does not consider the 2011 amendment because of data limitation.

Figure A-2: The Change in Reported Monthly Income on
Primary Earners' Credit Card Applications

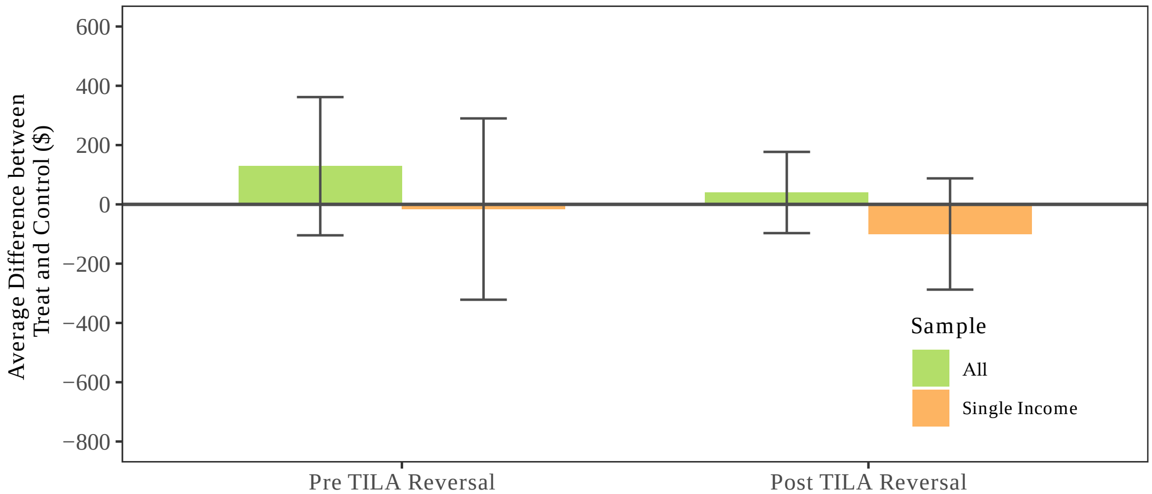

Notes: This figure shows the average difference in the monthly income reported on primary earners' credit card applications between the treated and the control group. The difference is obtained by regressing reported monthly income on the treatment dummy. The whiskers denote 90 percent confidence intervals.

Figure A-3: Spanish vs. British Colonization of the Americas

Notes: This figure illustrates how geographic differences in the marital property system in the U.S. can be traced back to the Spanish versus British Colonization in Northern America in the $17-19^{th}$ century. The community property system in the U.S. has its roots in the Spanish Civil Law while the equitable distribution (initially title-based) system has its roots in the English Common Law. See Table A.1 for details. Source: Academic Dictionaries and Encyclopedias (2022).

Figure A-4: Consumption Measure Validation

(a) Broad vs. Narrow Measure Validation: Wife

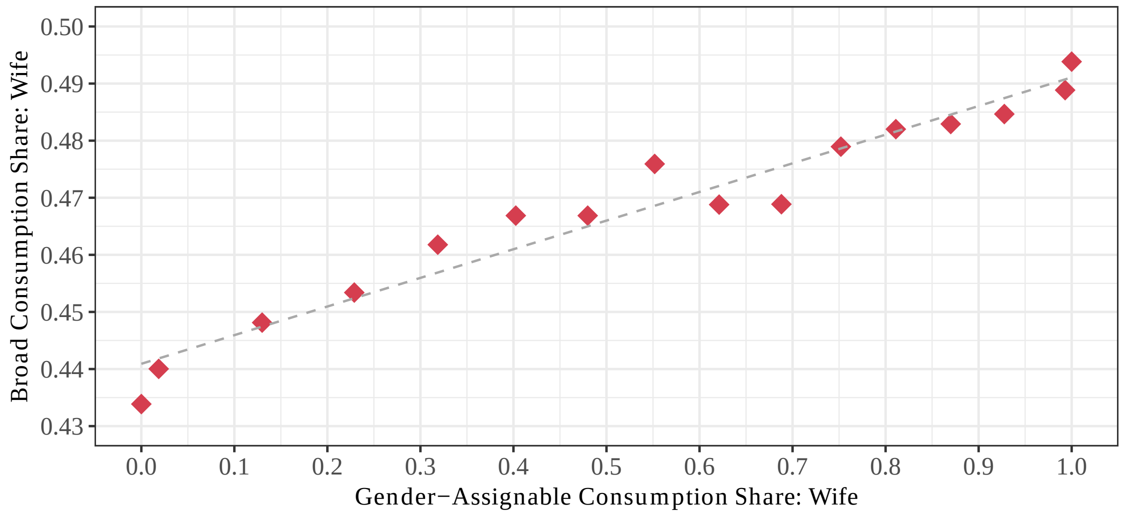

(b) Broad vs. Narrow Measure Validation: Husband

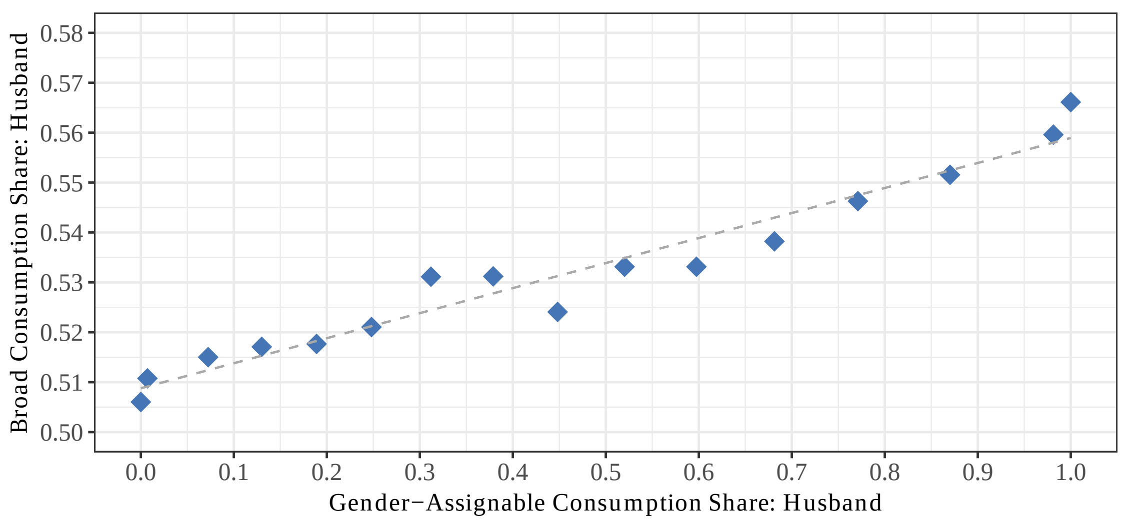

Notes: This Figure illustrates the validity of the assumption that "spenders are consumers" when using the broad consumption measure. Figure A-4a plots the wives' average monthly consumption share using the broad measure against the average monthly consumption share using the gender-assignable measure. If the broad consumption measure is a poor proxy for consumption because wives don't necessarily buy their own gender assignable goods, the slope of this figure would be 0. The positive slope illustrates that "spenders are consumers" is a reasonable proxy for consumption. Figure A-4b shows the same plot for the husbands.

Figure A-5: Household Checking Account Structure Types

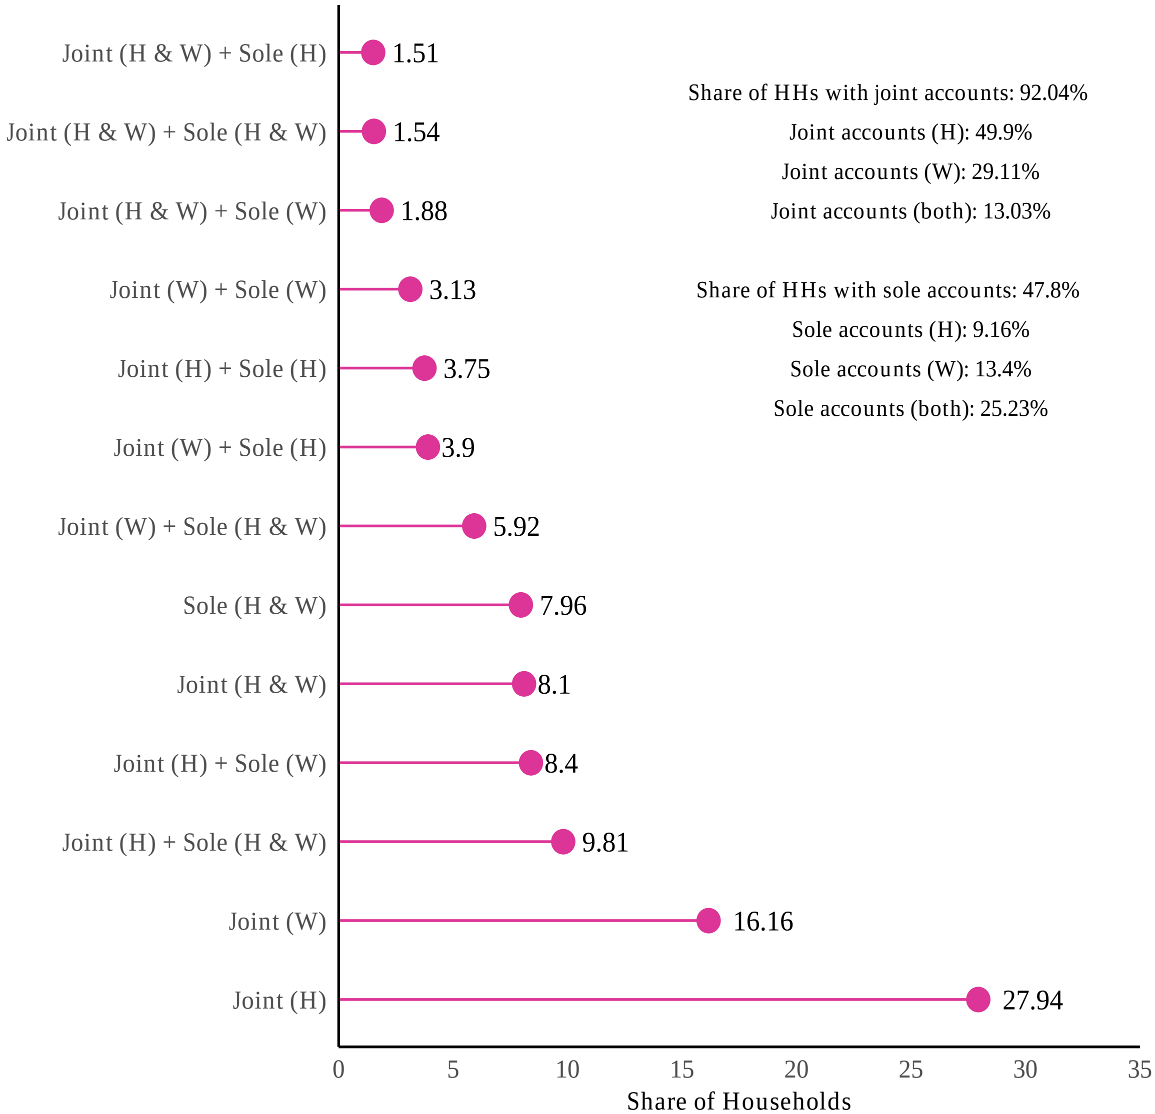

Notes: This figure reports the share of households that hold each type of checking account structure in my sample. The account structure types are mutually exclusive and the shares sum to 100. "Joint" and "Sole" denote the type of checking account, and "H", "W", or "H & W" in parenthesis denote whether the primary account holder of each account is the husband, the wife, or both because they have multiple accounts with each spouse as the primary account holder. For example, the line at the bottom shows that roughly 28% of households in my sample only have a joint checking account where the husband is the primary account holder.

Figure A-6: Household Credit Card Account Structure Types

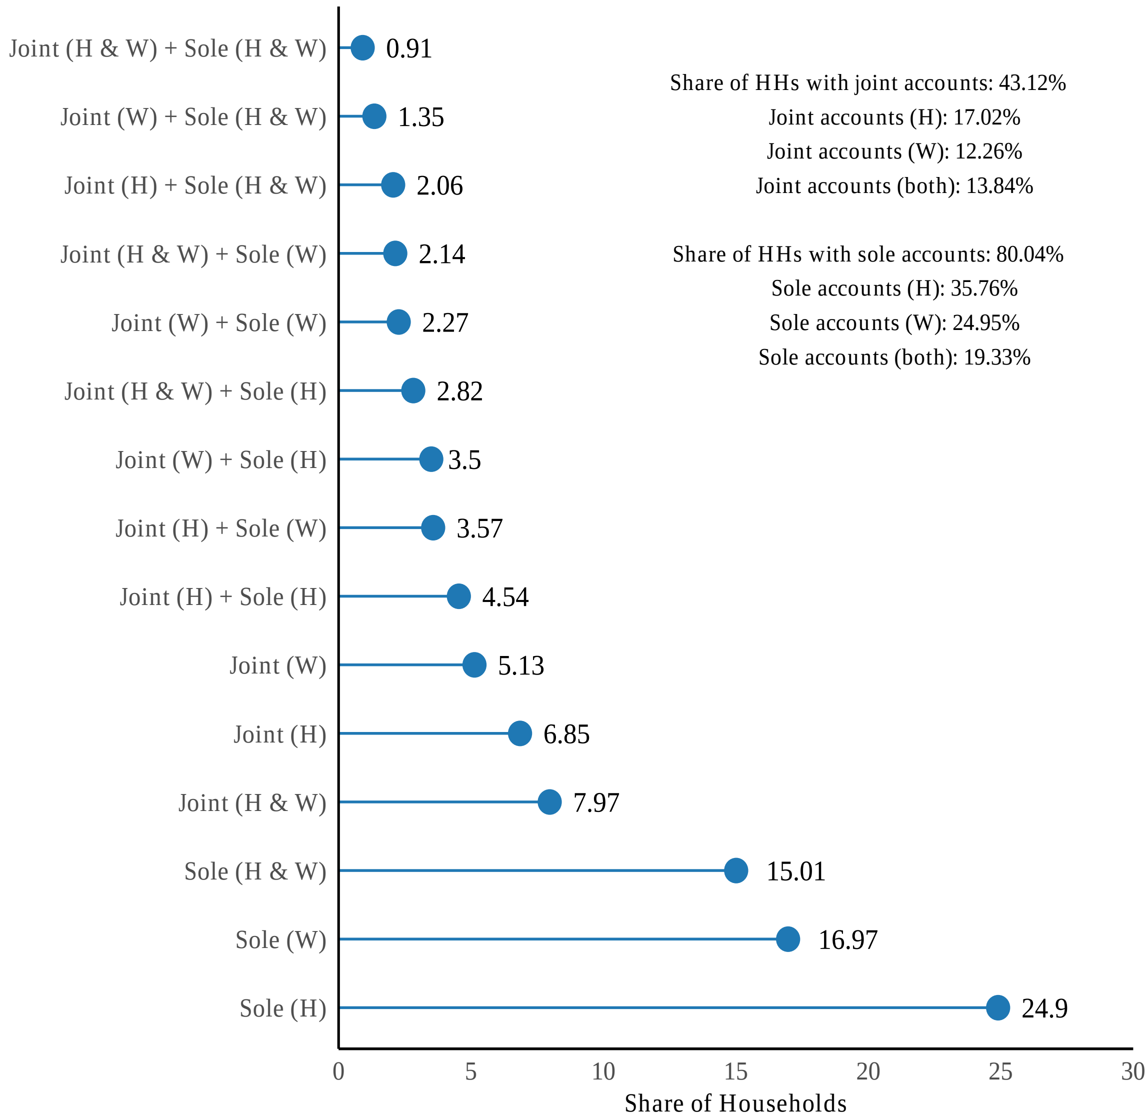

Notes: This figure reports the share of households that hold each type of credit card account structure in my sample. The account structure types are mutually exclusive and the shares sum to 100. "Joint" and "Sole" denote the type of credit card account, and "H", "W", or "H & W" in parenthesis denote whether the primary account holder of each account is the husband, the wife, or both because they have multiple accounts with each spouse as the primary account holder. For example, the line at the bottom shows that roughly 25% of households in my sample only have a sole credit card account where the husband is the primary account holder.

Figure A-7: Household Checking Account Structure Types by Treatment

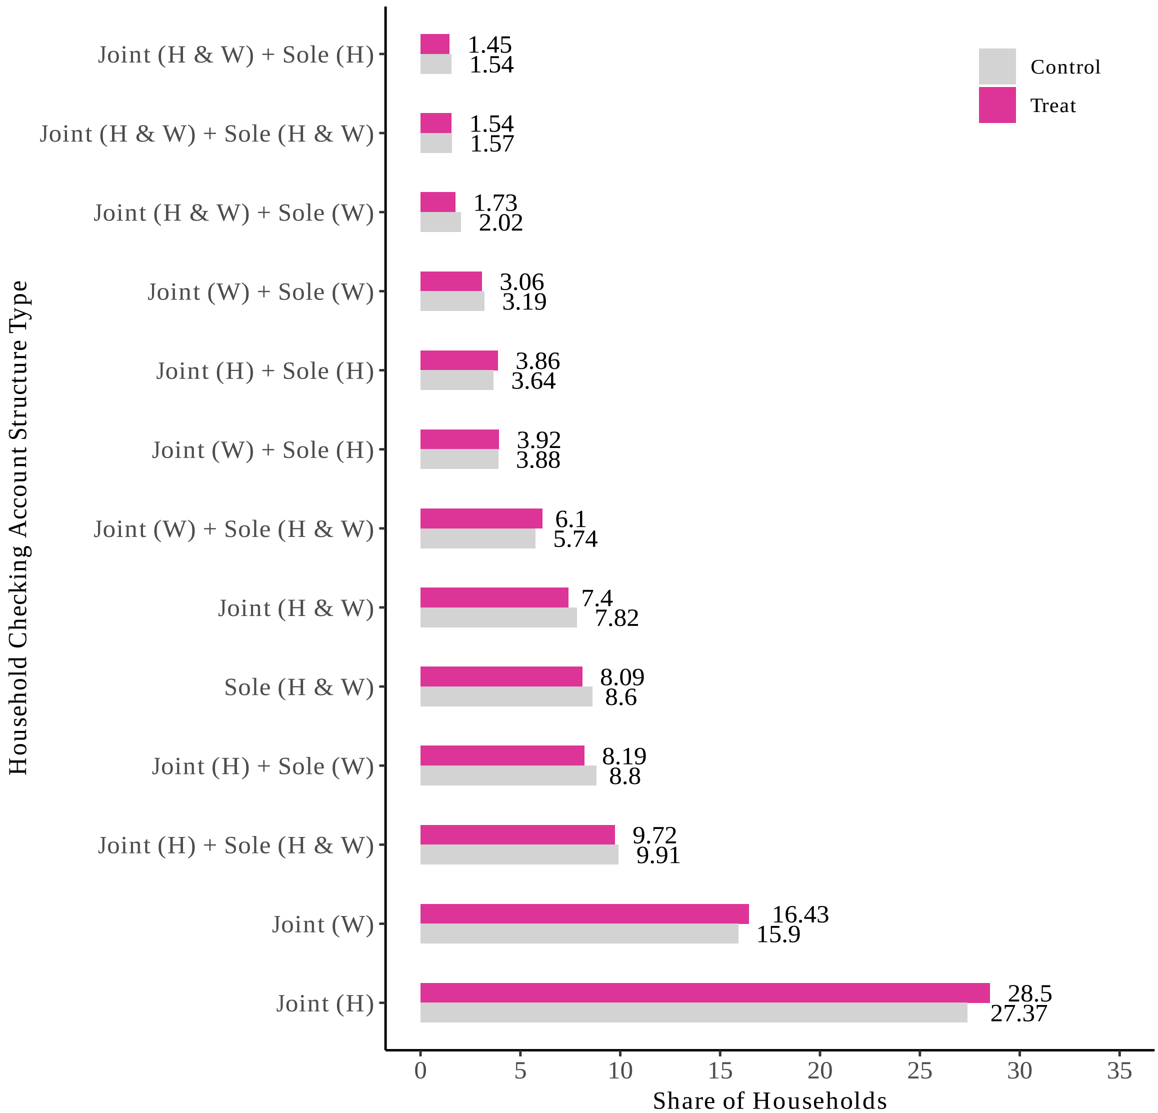

Notes: This figure shows the share of households that hold each type of checking account structure by treatment in my sample. See Figure A-5 for detailed description.

Figure A-8: Household Credit Card Account Structure Types by Treatment

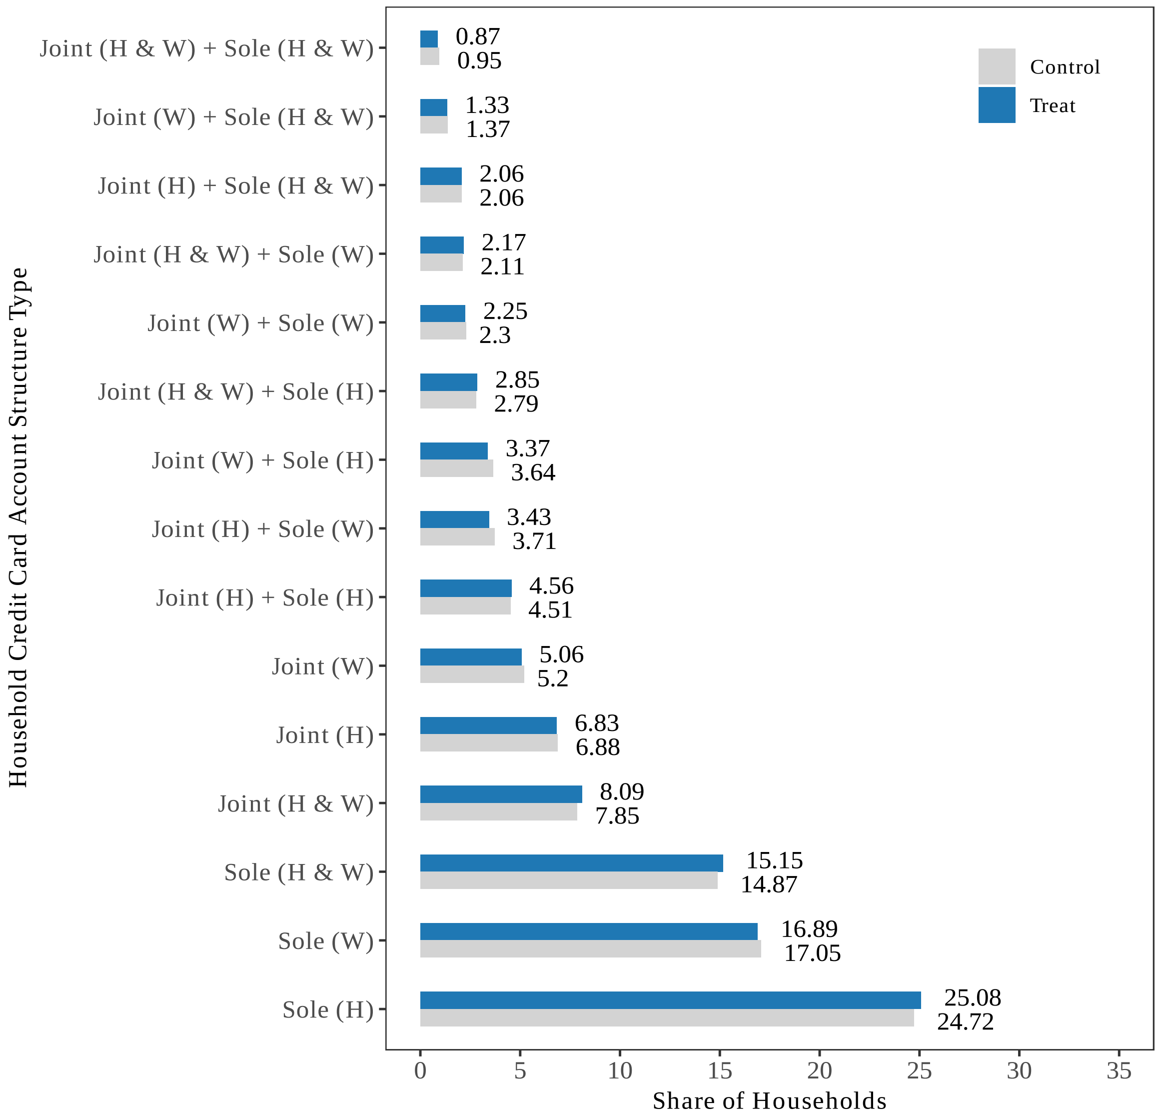

Notes: This figure shows the share of households that hold each type of credit card account structure by treatment in my sample. See Figure A-6 for detailed description.

Figure A-9: Broader Sample:
Within-Household Credit and Consumption Gaps

(a) Share of accessible in the Household

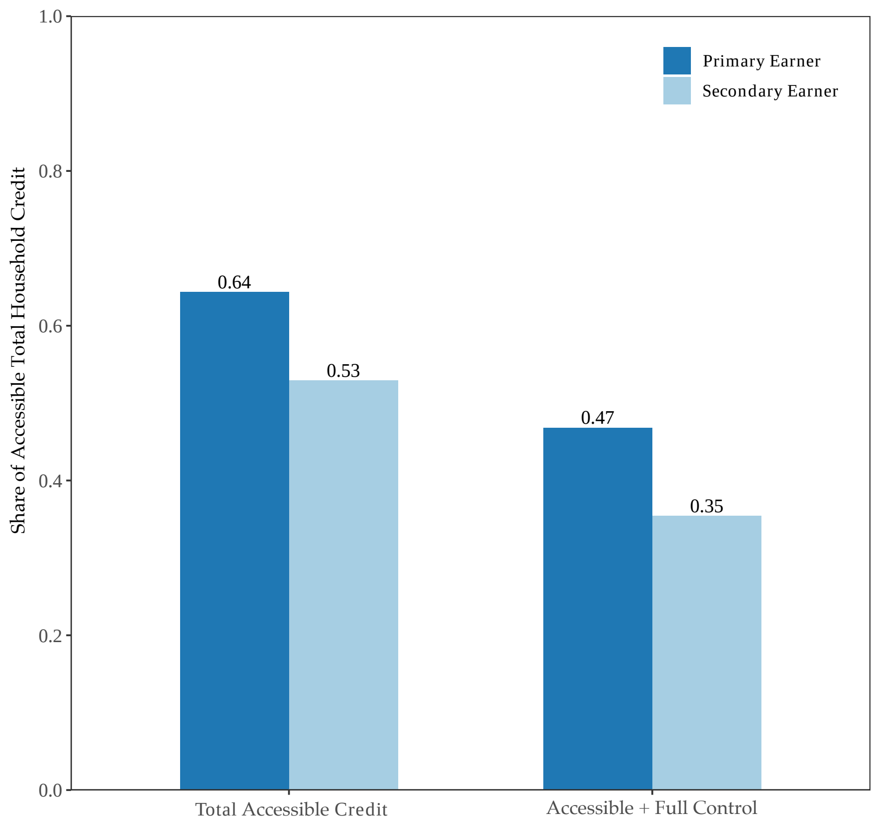

(b) Secondary Earners' Within-Household Consumption Share
by Credit Share Bin

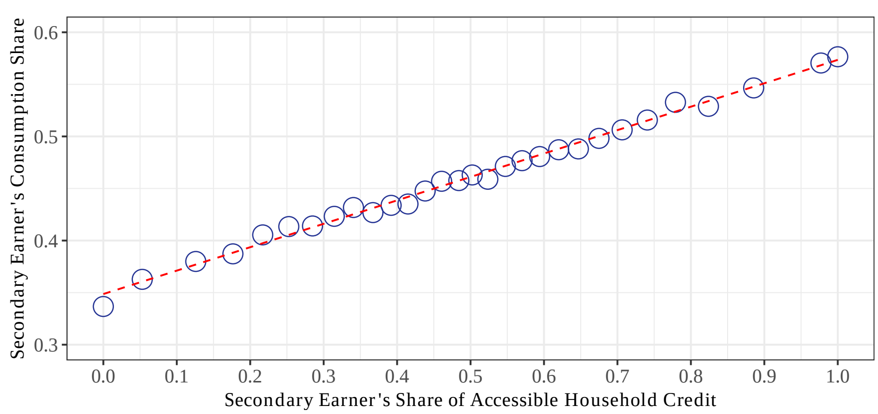

Notes: These figures replicate Figures 1-4 and 1-5 using a broader sample of 138,276 households that include households where secondary earners had credit card accounts at the beginning of my sample period.

Figure A-10: Changes in Secondary Earners' Share of Household Spending by Spending Category

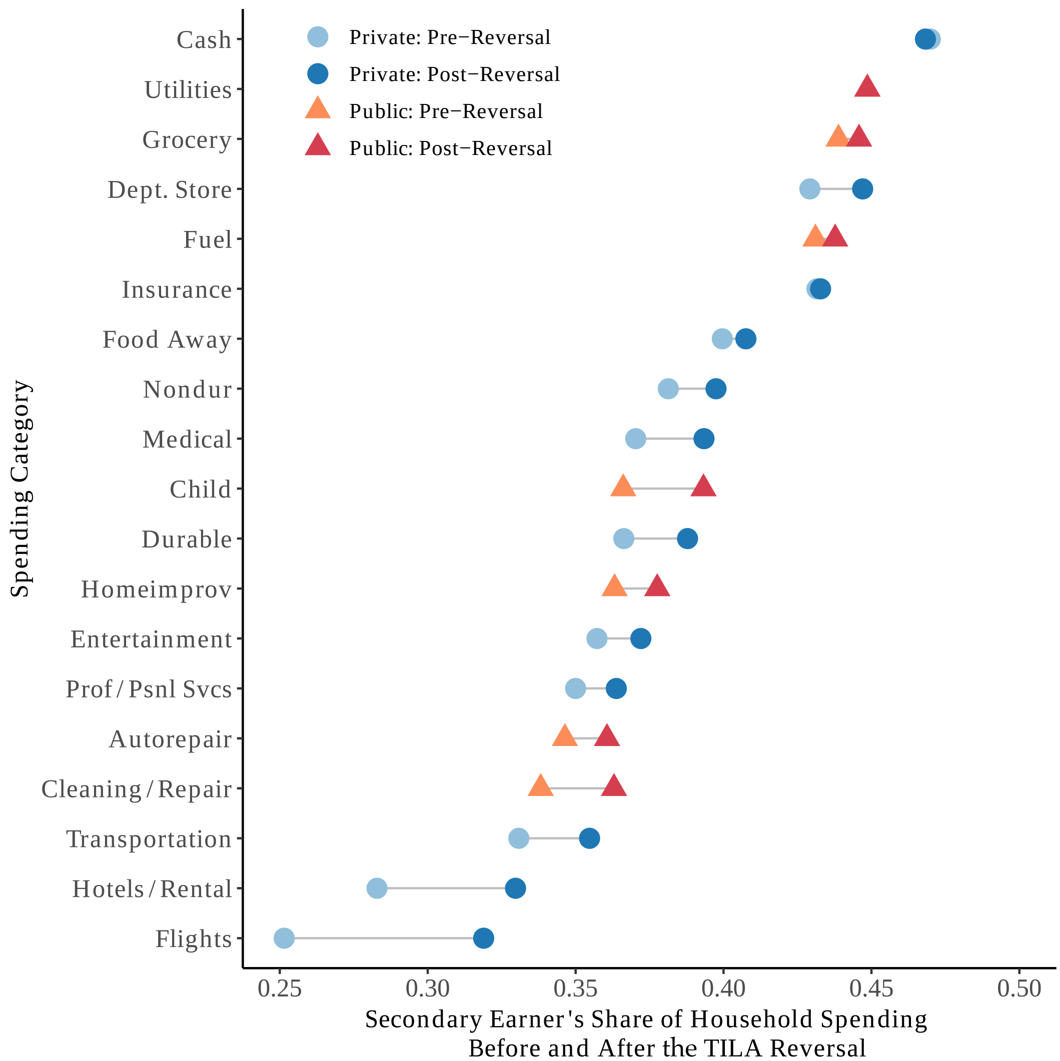

Notes: This figure plots Column 7 of Table A.5, which reports the change in secondary earners' average monthly spending share by spending category. The sample is limited to the treated group.

Figure A-11: Effect of the Reversal on Secondary Earners'
Credit Limit: Parametric

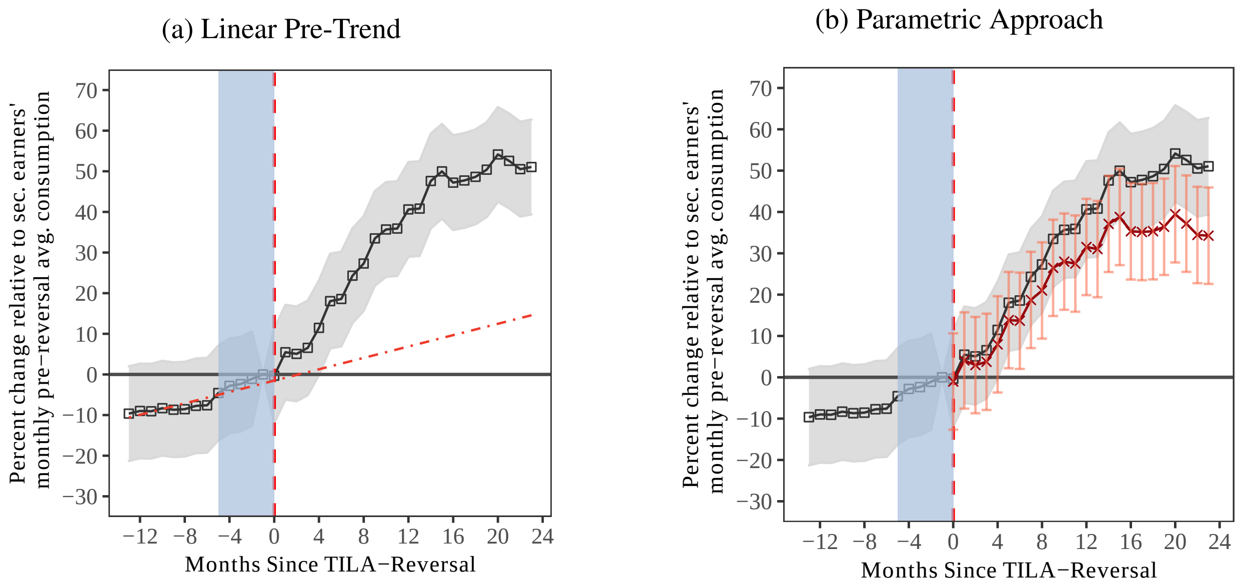

Notes: Figure a provides a visual assessment of the functional form assumption (linear) of pretend in event time. This pretend is driven by the CFPB allowing credit card issuers to start adopting the new income collection standard during the phase-in period (shaded in blue). Figure b superimposes the estimated parametric coefficients on the nonparametric coefficients shown in Figure 1-7. The parametric estimates are obtained by estimating:

$$Y_{h,t} = \alpha_h + \gamma_t + \sum_{s>t^{post}-1} \beta_s (Treat_h \times 1_{s=t}) + \lambda \cdot t \cdot Treat_h + \epsilon_{h,t} \tag{A.22}$$

which only keeps month by treatment fixed effects for post periods while estimating a linear pretend in event time interacted with treatment off the variation in the pre period.

Figure A-12: Decomposition of the Change in Secondary Earner Consumption

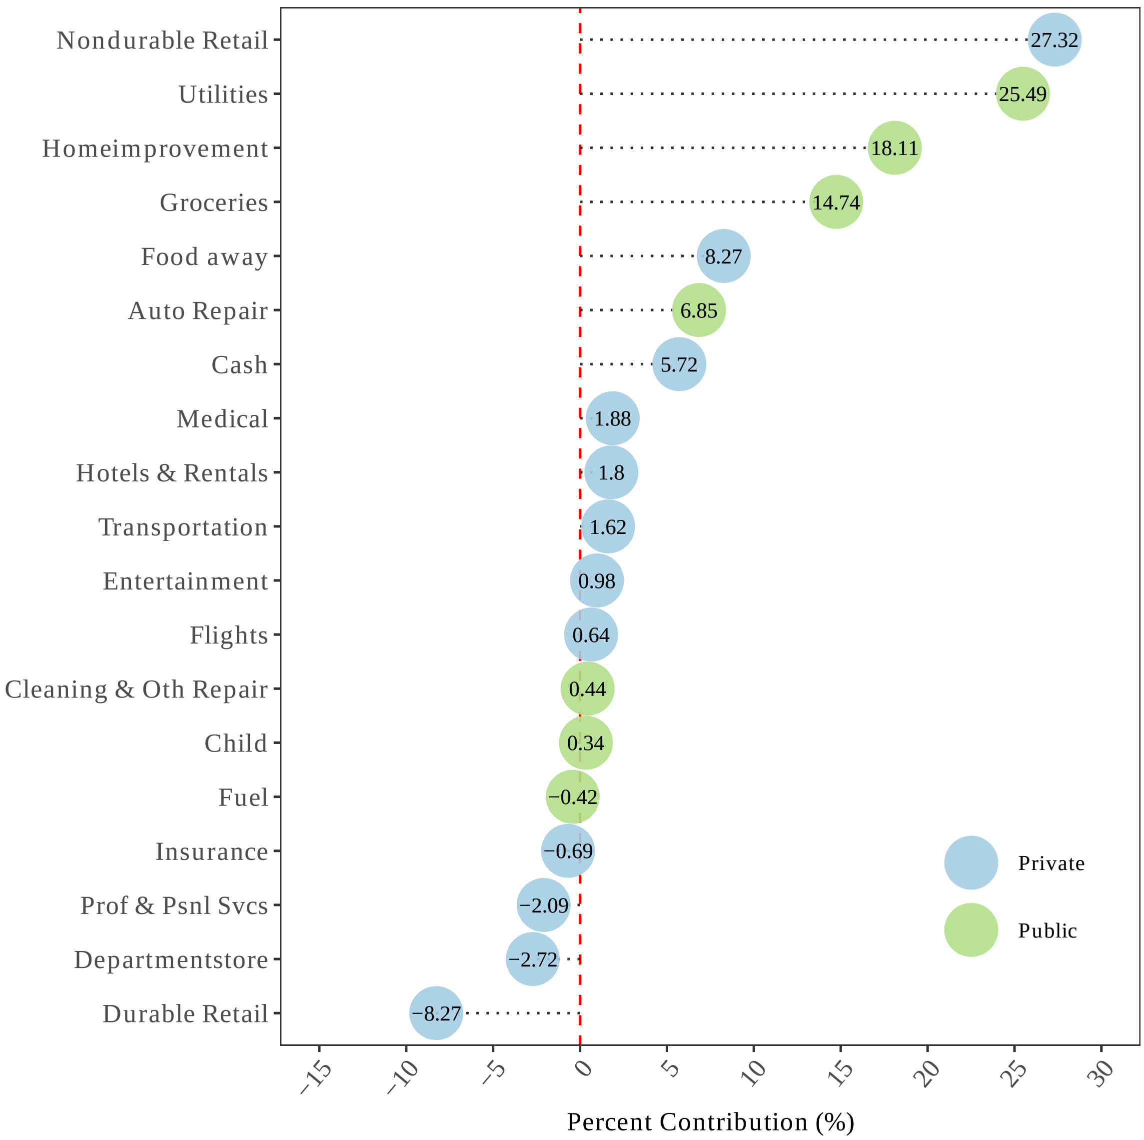

Notes: This figure decomposes the change in secondary earner consumption into detailed spending categories. The number shown in each bubble denotes how much each category contributes to the overall consumption effect, and the numbers sum to 100. For example, spending on nondurable retail explains 30% of the total increase in secondary earner consumption.

Figure A-13: Changes in Financial Situations After Divorce by Gender

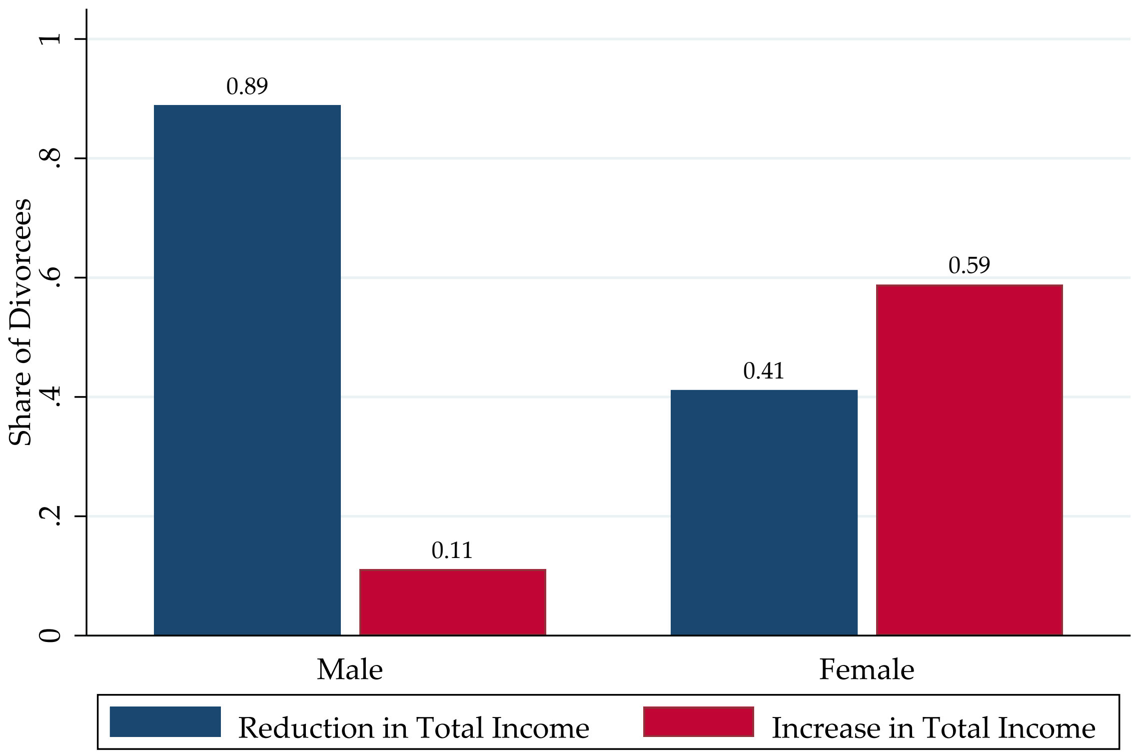

Notes: This figure shows the share of divorced individuals that experience a reduction (blue) or an increase (red) in total income relative to when they were married by gender using the 2012 Health and Retirement Survey (HRS). For example, 89% of male divorcees experienced a reduction in total income after divorce. Post-divorce total income includes labor income, social security benefits, veteran's benefits, pension, life insurance, and other lump-sum settlements. Post-divorce income excludes alimony because it is not reported in the HRS.

Figure A-14: Changes in the secondary earner's outside option
and allocation of resources

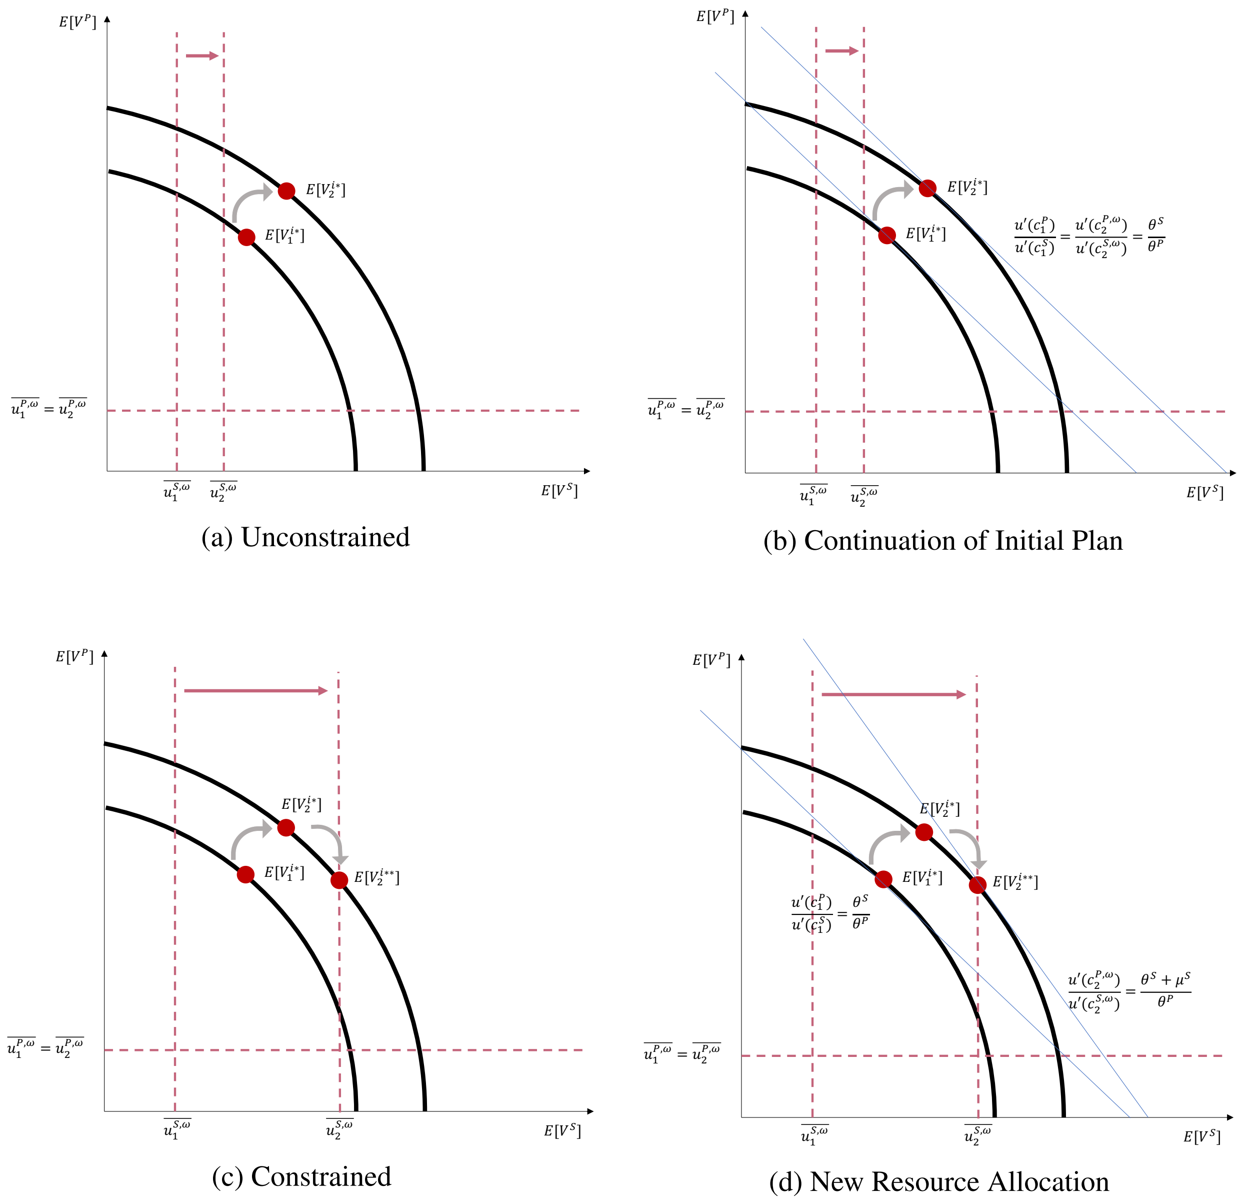

(a) Unconstrained

(b) Continuation of Initial Plan

(c) Constrained

(d) New Resource Allocation

Notes: This figure illustrates potential household responses to changes in the secondary earner's outside option. The y-axis plot the primary earner's expected utility and the x-axis plot the secondary earner's expected utility. Curved black lines show the Pareto frontier and the red points at the tangency of the Pareto frontier indicate the location of efficient intrahousehold allocation of resources. Red dashed lines indicate spouses' respective outside options and blue lines trace the slope of the Pareto frontier. This figure considers cases when only the secondary earner's outside option changes. Top figures a and b illustrate the case when the secondary earner's participation constraint does not bind. Bottom figures c and d illustrate the case when the improvement in secondary earner's outside option makes the participation constraint bind. This figure builds on Chiappori and Mazzocco (2017).

Table A.1: The Origin and the Civil Law Foundation of the U.S. Marital Property System

	Community Property (1)	Equitable Distribution (2)
	A. Background	
System	Spanish Civil Law	English Common Law
Foundation	Visigothic/Roman 653 AD	Anglo-Saxon/Norman 871 AD
	B. Characteristics	
Partnership	Spouses as equal partners	Spouses as one person in law
Legal status for women	Married woman as a separate judicial entity	Legal oneness of husband and wife
Marital property	Ownership-based	Initially title-based; now equitably distributed
Putative spouse doctrine	Recognized	Not recognized
Relationship between husband and wife	Civil contract between a man and a woman	Principle of covenant

Notes: Panel A reports the legal origin of the U.S. marital property system. Panel B reports key characteristics that differ between community property and equitable distribution system. Partnership refers to how the relationship between spouses is viewed under each system. Legal status for women refers to whether each system recognizes a married woman as a separate judicial entity, apart from her husband. Community property did not recognize the common law principle that the legal existence of the wife was merged into that of her husband (i.e., coverture, or the concept that dictated a woman's subordinate legal status during marriage because a woman's legal existence as an individual was suspended under "marital unity"). Marital property refers to how properties are treated under each legal system. Putative spouse doctrine refers to recognition of "putative" spouse, or a person who believes in good faith that he or she has a valid marriage, even though they do not. This concept is also known as "deemed marriages" and recognized under the Social Security program in the U.S. Relationship between husband and wife refers to whether respective law systems relied on the principle of covenant (i.e., more permanent marriage) to characterize the marital relationship. This note draws heavily from Newcombe (2011).

Table A.2: Sample Representativeness

	Benchmark Mean (1)	Sample Mean (2)
Head of Household Age (years)	55	44.31
Share of Double Income Households	0.53	0.54
Total Income ($)	83,413	118,729
Annual Consumption ($)	62,015	88,068
Public ($)	18,765	29,153
Private ($)	43,250	58,915
Expenditure to Income Share	0.74	0.74
Public to Expenditure Share	0.30	0.33
Private to Expenditure Share	0.70	0.67

Notes: This table compares the representativeness of my analysis sample described in Section 1.4.1 to external benchmarks from the Consumer Expenditure Survey (CEX) Table 3424 (i.e., consumer units of two people) for 2014 and Bureau of Labor Statistics (BLS). Column 1 reports annual average household characteristics in external benchmarks. The CEX excludes households that earn less than $20,000 to make the benchmark sample more comparable to my sample, which limits analysis to households to earn at least $17,000 (2013 U.S. poverty threshold for two-member household). Statistics are re-weighted by population share in each income bin. Column 2 reports annual average household characteristics for 2014-2015 in my sample. "Head of Household Age" shows the "age of reference person" in Column 1 and the oldest member in the household in Column 2. Total income includes labor, capital, business, retirement income, other income, and government transfers, including child support. Public expenditures reported Column 1 include spending on maintenance, repairs, other expenses; utilities, fuels; household operations; miscellaneous household equipment; laundry and cleaning supplies; other household products; household textiles; floor coverings; food at home; other vehicle expenses; and children. Private expenditures reported in Column 1 include all other spending that is not public spending. Spending categorization for Column 2 is reported in Table A.4.

Table A.3: Summary Statistics of Account Ownership Structure and Payment Choice for Married Individuals

	Mean (1)	Median (2)
Number of:		
Checking Accounts	1.51	1
Debit Cards	1.46	1
Credit Cards	4.03	3
Has Credit Cards	0.84	1
Use Cash	0.93	1
Primarily Obtain Cash from ATM	0.55	1
Checking Accounts Shared with a Spouse:		
Primary Account	0.73	1
Secondary Account	0.28	0
Own Primary Residence	0.82	1

Notes: This table reports account ownership structure and payment choice statistics for married individuals using the 2020 Survey of Consumer Payment Choice (SCPC). The financial accounts considered in the survey are those that belong to the survey respondent or jointly with the respondent's spouse. It excludes accounts that are only held by the respondent's spouse. "Use Cash" reports the share of respondents that used cash as a payment method in the last 12 months. "Checking Accounts Shared with a Spouse" reports the share of respondents who share their primary or secondary checking account with their spouse. "Own Primary Residence" asks whether the survey respondent or the respondent's spouse is a home owner.

Table A.4: Detailed Spending Categories

Category	**Type**	**Examples**
Department Store	Private	Department stores
Entertainment	Private	Theater, travel agency, tourist attraction, cruise lines, golf course, recreational camps
Flights	Private	Various airline companies
Hotels/Rentals	Private	Hotels, inns, resorts
Insurance	Private	Insurance premiums, direct marketing insurance service
Medical	Private	Ambulance services, dentists, doctors and physicians, chiropractors, optometrists, nursing and personal care facilities.
Transportation	Private	Cabs, bus lines, passenger railways, airports, parking lots, transportation svcs
Food Away	Private	Caterers, eating places and restaurants, fast food restaurants
Durable Retail/Misc	Private	Equipment, appliances, electronics, furniture, donation, organization, membership
Nondurable Retail/Misc	Private	Stationary, office supplies, duty free store, discount store, book store
Cash	Private	ATM withdrawals
Prof/Personal Services	Private	Consulting, legal, funeral services, tax preparations, advertising, tailors, mending
Auto Repairs/Parts	Public	Car washes, paint shops, automobile and truck dealers, vehicle supplies and new parts, car sales, services, repairs
Fuel	Public	Service stations, automated fuel dispensers
Utilities	Public	Utility service, electric, gas, sanitary and water, cable, telecommunication services
Groceries	Public	Grocery stores and supermarkets
Home improvement	Public	Florists, hardware supplies, home supply warehouse stores, building materials, glass stores, wall paper stores, garden supply stores
House keeping/repairs	Public	Cleaning, maintenance, repairs, heating, roofing
Child	Public	Child care services, children's and infant's wear stores

Notes: This table reports examples of detailed spending types included in each spending category.

Table A.5: Average Change in Secondary Earners' Consumption Share by Finer Spending Category (Treated Only)

	Pre-Reversal			Post-Reversal			Difference
	Household Mean (1)	Sec. Earner Mean (2)	Share Col 2 / Col 1 (3)	Household Mean (4)	Sec. Earner Mean (5)	Share Col 5 / Col 4 (6)	Col 6 - Col 3 (7)
				A. Private Consumption ($)			
Flights	43.90	11.04	25.15	57.41	18.31	31.89	6.74
Hotels & Rental	108.63	30.73	28.29	144.46	47.63	32.97	4.68
Durable Retail	63.96	21.16	33.08	82.75	29.35	35.47	2.39
Medical	83.48	30.92	37.03	106.82	42.02	39.34	2.31
Transportation	468.59	171.65	36.63	547.49	212.33	38.78	2.15
Department Store	58.48	25.10	42.92	70.02	31.30	44.70	1.78
Nondurable Retail	1,607	612.77	38.13	1,906	757.45	39.75	1.61
Entertainment	136.53	48.77	35.72	175.08	65.14	37.20	1.48
Food Away	183.82	64.34	35.00	237.33	86.33	36.37	1.38
Professional Services	323.89	129.42	39.96	396.47	161.58	40.75	0.80
Insurance	283.08	122.17	43.16	329.39	142.56	43.28	0.12
Cash	738.67	347.11	46.99	753.28	352.72	46.82	-0.17
				B. Public Consumption ($)			
Child	15.06	5.51	36.61	20.29	7.98	39.32	2.71
Home Cleaning/Repair	53.48	18.09	33.82	67.43	24.47	36.30	2.48
Autorepair	572.86	208.04	36.32	681.53	257.34	37.76	1.44
Home Improvement	253.18	87.70	34.64	299.99	108.18	36.06	1.42
Fuel	448.78	196.96	43.89	508.38	226.64	44.58	0.69
Groceries	279.15	120.33	43.11	264.74	115.88	43.77	0.67
Utilities	427.62	191.82	44.86	484.71	217.46	44.86	0.01

Notes: This table reports the average household and secondary earner spending on each spending category before and after the 2013 TILA reversal. The sample is limited to the treated group households only. Columns 3 and 6 report secondary earners' average monthly spending share and Column 7 reports the change in the share.

Table A.6: Extensive Margin: Credit Card Opening and Closing

	Card Holders			All Sample		
	Baseline	Single Income	Sec. Earner Older	Baseline	Single Income	Sec. Earner Older
	(1)	(2)	(3)	(4)	(5)	(6)
	A. Secondary Earners' Sole Credit Card Accounts					
Credi Card Opening	-.036	-.17	.08	-.004	-.03	.02
	(.096)	(.141)	(.156)	(.017)	(.025)	(.028)
Credit Card Closing	.002	-.01	-.02	.001	-.001	-.003
	(.013)	(.02)	(.023)	(.002)	(.004)	(.004)
	B. Joint Credit Card Opening					
Accounts held by Secondary Earners	-.01	-.02	-.03	.002	.01	.01
	(.01)	(.02)	(.02)	(.01)	(.015)	(.017)
Accounts held by Primary Earners	.02	.01	-.01	-.001	-.004	.003
	(.02)	(.03)	(.03)	(.011)	(.016)	(.017)
Number of Observations	455,157	211,916	171,117	2,577,970	1,198,437	944,585
	C. Pre TILA-Reversal Mean					
Card Opening Probability (%)	2.1	2.1	2.0	0.4	0.4	0.4
Card Closing Probability (%)	0.01	0.01	0.01	0.00	0.00	0.00
Sec Earner Joint Account Opening (%)	0.03	0.04	0.03	0.12	0.12	0.13
Prim Earner Joint Account Opening (%)	0.07	0.06	0.05	0.14	0.14	0.14

Notes: This table reports difference-in-differences regression estimates. The dependent variables in Panel A include indicators for secondary earners' sole credit card account opening and closing; and in Panel B include indicators for joint credit card opening rates by accounts held by secondary earners or primary earners. Panel C reports pre-reversal average of the outcome variables. All specifications include household and time (month-year) fixed effects. Standard errors are clustered at the state-level and reported in parentheses. Reported coefficients are multiplied by 100 for readability. β can be interpreted as a percentage point change in secondary earners' credit card opening or closing probabilities. The first three columns restrict the sample to households where secondary earners eventually opened sole credit card accounts during my sample period (i.e., The "Card Holders" sample), and the last three columns use all sample, including households where secondary earners never opened credit card accounts (i.e., The "All Sample"). Within each sample, Columns 2 and 5 further restrict the sample to single-income households where primary earners are breadwinners, and Columns 3 and 6 restrict the sample to those where secondary earners are older than primary earners. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.7: Secondary Earners' Other Credit Card Outcomes

Secondary Earner Outcomes	Card Holders			All Sample		
	Baseline	Single Income	Sec. Earner Older	Baseline	Single Income	Sec. Earner Older
	(1)	(2)	(3)	(4)	(5)	(6)
	A. Difference-in-Differences Estimates					
Annual Percentage Rate	.003	.09	-.28	.003	.09	-.28
	(.06)	(.09)	(.202)	(.062)	(.091)	(.202)
Credit Card Balance	5.99 ***	11.45 ***	8.49 ***	2.28 ***	4.63 ***	2.95 ***
	(.595)	(.914)	(.981)	(.181)	(.277)	(.303)
Credit Card Utilization	-.43 ***	-.33 *	-.25	-.052	.008	-.142 **
	(.14)	(.2)	(.22)	(.034)	(.049)	(.059)
Other Bank Card Payments	.99 *	1.65 **	-.84	.95 ***	1.79 ***	.72 **
	(.52)	(.82)	(.83)	(.219)	(.341)	(.354)
Number of Observations	455,157	211,916	171,117	2,577,970	1,198,437	944,585
	B. Pre TILA-Reversal Mean					
Annual Percentage Rate (%)	4.59	4.11	4.88	4.59	4.11	4.88
Credit Card Balance ($)	351	318	353	677	621	663
Credit Card Utilization (%)	5.93	5.71	5.75	4.28	4.05	4.29
Other Bank Card Payments ($)	334	290	325	364	320	343

Notes: This table reports difference-in-differences regression estimates. The dependent variables include secondary earners' annual percentage rates (APR) on their sole credit card accounts, end-of-billing-cycle credit card balance, credit card utilization rates, and credit card payments to other financial institutions. Credit card balance and card payments to other financial institutions are scaled by secondary earners' pre-reversal average monthly consumption. Thus, β can be interpreted as a percentage point change in the APR or credit card utilization rates (first and third outcomes) or as a percent change in secondary earners' credit card balance or card payments to other banks relative to their pre-reversal consumption. All specifications include household and time (month-year) fixed effects. Standard errors are clustered at the state-level and reported in parentheses. Reported coefficients are multiplied by 100 for readability. The first three columns restrict the sample to households where secondary earners eventually opened sole credit card accounts during my sample period (i.e., The "Card Holders" sample), and the last three columns use all sample, including households where secondary earners never opened credit card accounts (i.e., The "All Sample"). Within each sample, Columns 2 and 5 further restrict the sample to single-income households where primary earners are breadwinners, and Columns 3 and 6 restrict the sample to those where secondary earners are older than primary earners. Panel B reports pre-reversal average of the outcome variables. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.8: Household Financial Decision-Making:
Other Credit Card Outcomes and Long-Term Effects

Household Outcomes	Card Holders			All Sample		
	Baseline	Single Income	Sec. Earner Older	Baseline	Single Income	Sec. Earner Older
	(1)	(2)	(3)	(4)	(5)	(6)
	A. Other Outcomes					
Revolving Balance Utilization Rate	-0.19	-0.49 **	-0.08	0.07	-0.35 ***	0.00
	(0.15)	(0.21)	(0.24)	(0.06)	(0.08)	(0.1)
Average APR	0.07	0.11	-0.08	0.00	0.00	0.01
	(0.05)	(0.07)	(0.07)	(0.01)	(0.02)	(0.02)
Number of Observations	455,157	211,916	171,117	2,577,970	1,198,437	944,585
	B. Long-Term Effects					
Delinquency	0.09	-0.04	0.03	0.00	-0.05	0.03
	(0.06)	(0.07)	(0.08)	(0.02)	(0.03)	(0.04)
Overdraft	0.08	0.01	0.08	0.02	-0.01	0.00
	(0.05)	(0.08)	(0.09)	(0.02)	(0.03)	(0.03)
High-interest loans	-0.01	0.07	0.14	0.03	0.06 **	0.03
	(0.06)	(0.08)	(0.1)	(0.02)	(0.03)	(0.03)
Debt prioritization	1.32 **	2.21 **	1.39	0.51	1.00 **	1.13 **
	(0.65)	(0.94)	(1.07)	(0.34)	(0.51)	(0.56)
Number of Observations	303,239	141,196	114,008	1,717,422	798,401	629,309
	C. Pre-Reversal Mean					
Revolving Balance Utilization	15.47	15.45	15.63	23.43	23.45	24.24
Average APR of HH Cards	12.00	11.94	12.42	13.83	13.81	13.95
Delinquency	0.32	0.29	0.28	0.65	0.64	0.68
Overdraft	0.73	0.81	0.77	0.52	0.54	0.55
High-interest loans	1.05	1.08	1.21	0.77	0.70	0.91
Debt prioritization	80.64	80.40	80.05	81.72	81.21	80.66

Notes: This table reports difference-in-differences regression estimates. Panel A examines other household credit card outcomes, such as revolving balance utilization rates and average annual percentage rate (APR) on credit cards available to households. Panel B examines long-run effects of the TILA reversal by dropping the first year of the post-period data for the same outcomes reported in Table 1.6. Debt prioritization analysis is limited to households with at least two credit card accounts. Roughly 55% (20%) of households sampled in the first (last) three columns have multiple credit cards. All specifications include household and time (month-year) fixed effects. Standard errors are clustered at the state-level and reported in parentheses. Reported coefficients are multiplied by 100 for readability. The first three columns restrict the sample to households where secondary earners eventually opened sole credit card accounts during my sample period (i.e., The "Card Holders" sample), and the last three columns use all sample, including households where secondary earners never opened credit card accounts (i.e., The "All Sample"). Within each sample, Columns 2 and 5 further restrict the sample to single-income households where primary earners are breadwinners, and Columns 3 and 6 restrict the sample to those where secondary earners are older than primary earners. Panel C reports pre-reversal average of the outcome variables. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.9: Effect of the TILA Reversal on Secondary Earner Consumption and Consumption Share Using Gender-Assignable Measure

Secondary Earner Outcomes	Gender-Assignable Measure Baseline (1)		Debit Cards (2)		Include Cash (3)		Include Non-Clothing (4)	
	A. Difference-in-Differences							
Consumption	6.05	***	4.61	***	2.49	***	5.10	***
	(0.66)		(0.68)		(0.32)		(0.39)	
Consumption Share	4.42	***	5.03	***	1.77	**	5.98	***
	(0.59)		(0.73)		(0.72)		(1.29)	
Number of Observations	163,412		163,412		163,412		163,412	
	B. Pre-Reversal Mean							
Consumption ($)	10.4		7.7		279		43.9	
Consumption Share (%)	51.2		52.3		48.0		50.7	

Notes: This table replicates Table 1.4 using the gender-assignable consumption measure described in Section 1.4.2. The sample is restricted to households where secondary earners eventually opened sole credit card accounts during my sample period (i.e., The "Card Holders" sample). Since spending on gender-assignable goods (i.e., clothing) is infrequent, the estimates are obtained from aggregating data to quarterly and by scaling the outcome variables by the average pre-reversal group mean. Specifically, secondary earners' consumption (share) is scaled by the average pre-reversal monthly consumption (share) within each treatment group. The quarterly estimates are converted back to monthly rates to facilitate comparison to other tables. Columns 1 and 2 use gender-assignable measures based on spending on clothing only. Columns 2 uses a gender-assignable measure constructed using spending on debit cards only. Column 3 uses a measure that includes each spouse's cash withdrawals to the measure used in Column 1. Column 4 uses a broader gender-assignable measure that includes broader spending categories (in addition to clothing) that were shown to be associated with gender-specific income shocks in existing studies, such as alcohol, gambling, and tobacco for men (Duflo and Udry, 2004); and hair or nail salons, spas, or jewelry for women. All specifications include household and time (month-year) fixed effects. Standard errors are clustered at the state-level and reported in parentheses. Reported coefficients are multiplied by 100 for readability.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.10: Measurement Robustness

Secondary Earner Outcomes	Broad Measure: Include Other Cards (1)		Net of Travel (2)		Net of Cash (3)		Debit Cards (4)	
	A. Difference-in-Differences							
Consumption	8.90	***	13.12	***	22.10	***	4.96	***
	(0.77)		(1.)		(1.43)		(0.74)	
Consumption Share	4.19	***	4.89	***	6.54	***	1.25	***
	(0.28)		(0.33)		(0.38)		(0.22)	
Number of Observations	462,896		462,896		462,896		462,350	
	B. Pre-Reversal Mean							
Consumption ($)	2,892		2,464		2,197		1,803	
Consumption Share (%)	45.4		45.3		44.9		48.3	

Notes: This table examines the sensitivity of baseline estimates to using alternative "broad" consumption measure used in Table 1.4. The outcomes are scaled by secondary earners' pre-reversal monthly mean of each outcome. The sample is restricted to households where secondary earners eventually opened sole credit card accounts during my sample period (i.e., The "Card Holders" sample). Column 1 uses a measure that includes payments to other credit companies to the broad measure. Column 2 uses a measure that excludes spending on work-related (travel) expenses, such as spending on flights, hotels/lodging, and transportation. Column 3 uses a measure that excludes cash withdrawals from the broad measure. Columns 4 uses a consumption measure constructed using spending on debit cards only. All specifications include household and time (month-year) fixed effects. Standard errors are clustered at the state-level and reported in parentheses. Reported coefficients are multiplied by 100 for readability. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.11: Specification Robustness

Secondary Earner Outcomes	Baseline (1)		No Controls (2)		Only HH f.e. (3)		Only Time f.e. (4)		State Trends (5)		Quarterly Spec (6)	
Credit Limit	40.45	***	40.51	***	40.35	***	40.59	***	36.80	***	35.60	***
	(1.69)		(1.91)		(1.76)		(1.85)		(1.83)		(2.71)	
Consumption Share	4.97	***	5.04	***	5.04	***	5.04	***	4.53	***	4.90	***
	(0.34)		(0.34)		(0.34)		(0.34)		(0.37)		(0.49)	
Number of Observations	453,910		453,910		453,910		453,910		453,910		163,246	
Household f.e.	✓				✓				✓		✓	
Time f.e.	✓						✓		✓		✓	
State-specific trends									✓			
Quarterly specification											✓	

Notes: This table examines the sensitivity of baseline estimates to using alternative specifications. The sample is restricted to households where secondary earners eventually opened sole credit card accounts during my sample period (i.e., The "Card Holders" sample). The dependent variables are: (i) secondary earners' sole credit card limits scaled by their average pre-reversal monthly consumption; and (ii) secondary earners' consumption share in the household scaled by their average pre-reversal monthly consumption share. Column 1 reports my baseline DiD estimates also reported in Tables 1.3 and 1.4. Column 2 excludes household and time fixed effects. Column 3 only includes household fixed effects, and Column 4 only includes time fixed effects. Column 5 includes state-specific linear time trends in addition to the baseline specification used in Column 1. Column 6 reports estimates obtained from aggregating data to quarterly. The quarterly estimates are converted back to monthly rates to facilitate comparison to other columns.*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.12: Sample Robustness

	Other Samples				Placebo		
	No Joint Checking	No Joint Any	Has Joint	Any HH with Cards	No Changes in Access	Only Using Pre-Period	Primary Earner
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	A. Secondary Earner Outcomes						
Credit Limit ($)	45.75***	42.91***	40.51***	2.49***	–	5.14	
	(4.32)	(4.61)	(1.81)	(0.67)	–	(3.77)	
Consumption Share (%)	7.84***	6.22***	4.84***	1.12***	0.19	0.63	
	(1.05)	(1.09)	(0.36)	(0.14)	(0.21)	(0.56)	
	B. Primary Earner Outcome						
Credit Limit ($)							0.57
							(0.92)
Number of Observations	53,735	47,224	407,933	1,726,062	851,489	151,918	455,157
Credit Limit ($)	626	659	902	7,421	0	877	3,294
Consumption Share (%)	44.6	45.3	45.1	46.9	38.6	45.1	

Notes: This table examines the sensitivity of baseline estimates to using alternative samples. The outcome variables are secondary earners' sole credit card limit scaled by their pre-reversal average monthly consumption (first row); secondary earners' consumption share scaled by their pre-reversal average monthly consumption share (second row); and primary earners' sole credit card limit scaled by their pre-reversal average monthly consumption (third row). Column 1 restricts the sample to households without joint checking accounts. Column 2 restricts the sample to households without any joint accounts (i.e., checking or credit). Column 3 restricts the sample to households with any joint accounts (i.e., checking or credit). Column 4 uses a broader sample of households including those where secondary earners had credit card accounts at the beginning of my sample period. Column 5 reports estimates for households with no change in credit because no household member opens or closes credit card accounts. Column 6 only uses the pre-treatment periods and sets treatment date to be March 2013. Column 7 analyzes primary earners' sole credit card limit scaled by their pre-reversal average monthly consumption using the card holder sample. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.13: Strategic Behavior

	Card Holders						All Sample					
Revolving Balance Utilization Rates	Baseline (1)		Single Income (2)		Sec. Earner Older (3)		Baseline (4)		Single Income (5)		Sec. Earner Older (6)	
	A. Accounts Held by Secondary Earners											
Sole Accounts	-1.14	**	-1.79	***	-1.38	*	-0.53	**	-1.03	***	-1.39	***
	(0.45)		(0.68)		(0.74)		(0.25)		(0.38)		(0.4)	
Joint Accounts	0.48	***	-0.02		1.19	***	0.31		0.68	**	0.65	*
	(0.11)		(0.16)		(0.19)		(0.22)		(0.31)		(0.34)	
	B. Accounts Held by Primary Earners											
Sole Accounts	-1.52	***	-3.53	***	-2.59	***	-1.17	***	-2.98	***	-1.77	***
	(0.45)		(0.7)		(0.78)		(0.28)		(0.43)		(0.47)	
Joint Accounts	-0.03		0.59	***	0.84	***	0.58	***	1.43	***	0.51	
	(0.15)		(0.22)		(0.26)		(0.21)		(0.31)		(0.37)	
Number of Observations	32,143		13,872		11,949		63,089		27,814		24,041	
	C. Pre-Reversal Revolving Balance Utilization Mean											
Sec. Earner Sole Account	1.54		1.18		1.41		0.78		0.59		0.70	
Sec. Earner Joint Account	0.08		0.10		0.07		6.59		6.40		7.07	
Prim. Earner Sole Account	5.15		5.65		4.62		6.32		6.72		5.62	
Prim. Earner Joint Account	0.31		0.29		0.32		6.13		6.05		7.03	

Notes: This table presents difference-in-differences estimates using the sample of 1,620 households with multiple credit cards that carry the same interest rate to examine the strategic behavior of couples in the absence of price effects. Since interest rates are the same across all credit card accounts in the household, spouses should be indifferent between which credit cards they use unless they have strategic motives. The outcomes are monthly revolving balance utilization rates on different types of credit cards. Panel A reports estimates on sole or joint credit card accounts where secondary earners are primary account holders, and Panel B reports estimates on those accounts where primary earners are primary account holders. β can be interpreted as a percentage point change in average monthly revolving balance utilization rates. All specifications include household and time (month-year) fixed effects. Standard errors are clustered at the state-level and reported in parentheses. Panel C reports pre-reversal average revolving balance utilization rates.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.14: Comparison between Model and Data

	Model (1)	Data (treated) (2)	External Benchmark (3)
Labor Income	6,091	6,092	
Consumption	6,812	6,289	
Net Assets	3,776	3,965	
Revolving Debt	4,212	4,120	
Share of Revolvers	0.321	0.43	
Share Double Income	0.497	0.56	0.53
Probability of Divorce	0.404	—	0.44

Notes: This table compares average monthly household-level outcomes generated in the model and observed in the data. Column 2 reports statistics using the treated group only. Net assets and revolving debt in Column 1 refer to a_t when a_t is positive (net assets) and negative (borrowing), while they refer to checking account liquid balance and revolving credit card debt in Column 2. Share of revolvers represent the share of households that borrow in Column 1 and the share of households with positive revolving debt in Column 2. The share of double income and the probability of divorce in Column 3 are from the BLS and CDC, respectively.

Table A.15: Model Outcomes

	Pre-Reversal (1)	Post-Reversal (2)
Share Borrow	0.18	0.40
Share Working (Sec. Earner)	0.50	0.49
Probability of Divorce	0.40	0.41

Notes: This table reports the average monthly outcomes generated in the model before and after the TILA reversal.

Bibliography

Academic Dictionaries and Encyclopedias. 2022. "New Spain." *Webpage: https://en-academic.com/dic.nsf/enwiki/11563836.*

Chiappori, Pierre-André, and Maurizio Mazzocco. 2017. "Static and Intertemporal Household Decisions." *Journal of Economic Literature*, 55(3): 985–1045.

Duflo, Esther, and Christopher Udry. 2004. "Intrahousehold Resource Allocation in Côte D'Ivore: Social Norms, Separate Accounts, and Consumption Choices." *NBER Working Paper*, , (10498).

Newcombe, Caroline Bermeo. 2011. "The Origin and Civil Law Foundation of the Community Property System, Why California Adopted It and Why Community Property Principles Benefit Women." *University of Maryland Law Journal of Race, Religion, Gender, and Class*, 11(1): 1–38.

Appendix B

Appendix for Chapter 2

B.1 Conceptual Framework

I present a simple conceptual framework that highlights why education spending might lead to business responses of self-employed households. To fix ideas, I outline a parsimonious model that highlights the link between education spending and business outcomes in the presence of liquidity constraints. The goal of this section is to illustrate that education spending can affect self-employed households' labor margins, or the households' decision to grow the business using personal funds or exit to become wage-earners. The model heavily draws from Evans and Jovanovic (1989) and Holtz-Eakin, Joulfaian and Rosen (1994).

A small business owner h with entrepreneurial ability θ_h operates a business that generates gross receipts of $R_h = \theta_h f(k_h)\epsilon$. The decreasing returns to scale production function $f(\cdot)$ uses capital k_h (i.e., operating expense) as its only input[1], and there is an idiosyncratic component to production $\epsilon \sim N(1, \sigma^2)$ with mean 1 and finite variance. A business owner has available personal assets A_h that earn gross return r. Therefore, any remaining assets after purchasing business input $A_h - k_h$ can earn capital income. Then a business owner's net income is given by the sum of gross receipts and capital income generated by unused personal income: $Y_h^E = \theta_h f(k_h)\epsilon + r(A_h - k_h)$.[2] If a business owner does not have enough personal assets to buy capital inputs (i.e., liquidity

[1] I abstract away from labor input because more than 85% of the businesses in my sample are non-employer firms.

[2] By definition, $k_h - A_h$ is the amount of capital financed by borrowing.

constrained), she can borrow from financial markets, but only up to a point $k_h \leq \Phi(A_h)$. The borrowing capacity $\Phi(A_h)$ is imposed by imperfect financial markets, and the size of the capacity increases in the owner's personal assets such that $\Phi'(A_h) > 0$.

Under this condition, a business owner chooses her optimal capital k_h^* by maximizing the expected net business income

$$\max_{k_h} \mathbb{E}\Big[\theta_h f(k_h)\epsilon + r(A_h - k_h)\Big]$$
$$\text{s.t. } k_h \leq \Phi(A_h)$$

The maximization leads to three possible cases. The first is the unconstrained case, where the expected marginal product of capital is equal to the gross return r, or $\theta_h f'(k_h^*) = r$. In this case, k_h^* rises with the entrepreneur's ability. The second case is when liquidity constraints do not bind, but the expected marginal product of capital falls below r because of low entrepreneurial ability θ_h. If a business owner has an outside opportunity as a wage earner and has sufficiently low θ_h, such that the expected net income of running a business is lower than being a wage earner, $Y_h^E \leq Y_h^W = w_h + rA_h$, she exits. Finally, when the liquidity constraint binds, she sets $k_h^* = \Phi(A_h)$.

The solutions imply that capital input is sensitive to changes in personal assets for liquidity constrained households. For constrained households, capital input increases as A_h increases because $\frac{\partial k_h^*}{\partial A_h} = \Phi'(A_h) > 0$, while it does not change for unconstrained households: $\frac{\partial k_h^*}{\partial A_h} = 0$. Thus, k_h^* becomes a function of A_h, in addition of r and θ_h, which in turn affects the revenue of the firm

$$R_h = \theta_h f(k_h^*)\epsilon \equiv R(\theta_h, A_h, r, \epsilon)$$

This result highlights a potential reason why education spending can affect business outcomes. All else equal, entrepreneurial households with college-entering dependents have less personal funding to meet business capital demands: $\tilde{A}_h = A_h - E_h$, where $\tilde{A}_h \leq A_h$ denotes personal assets for households with education payment obligations, E_h. Therefore, the decrease from A_h to $\tilde{A}_h$ moves business owner's capital stock farther away from the optimal level, which in turn leads to lower business expenses, revenues, and higher exit probabilities.

B.2 Placebo and Diagnostics

The 2SLS identification hinges on the assumption that the instruments have a clear effect in the first stage. The strong first stage estimates reported in table 2.3 and the visual evidence in figure 2-1 of a sudden rise in spending on education at the 18 year-old-mark support this assumption. Another way to confirm the validity of my instruments is to conduct a placebo test of whether education spending increases for households with children of non-college entry age. If the interactions between quarter-to-quarter transition dummies and college-entry age dummy are valid predictors of spending on education that arises from sending kids to school, I should not detect strong first stage effects for households that do not have college-entering kids.

Table B.4 presents a placebo test that restricts the sample of self-employed households to those with children aged between 15 and 17. I compare the first stage and reduced form estimates of households with 15 year olds to 16 and 17 year olds. If the 2SLS assumptions hold, the business outcomes should not respond to a child's being 16 or 17 (i.e., not statistically different from the outcomes of the households with 15 year olds). Column 1 shows that education spending rises with a child's age. However, the magnitude of this increase is small relative to the baseline estimate of 40 log points reported in table 2.3. Moreover, the increase in education spending is not large enough to induce large business spending response.

In addition to placebo, I run additional diagnostic tests to confirm the validity of my instruments (i.e. $Z_{h,t}$ =Quarter Transitions $\times$ $\mathbb{1}$(College-going age)). Table B.5 reports statistics from these tests based on the 2SLS results reported in table 2.3. The weak instruments diagnostics tests the null hypothesis that all instruments are weak. I reject this null, confirming that my instruments are strong. The Wu-Hausman test for endogeneity confirms whether the Instrumental Variable approach is the appropriate empirical strategy– i.e., that education spending is indeed correlated with the error term in equation 2.1 and thus there is a need for instruments. I reject the null that $H_0 : Cov(Ed, \eta) = 0$, confirming the existence of endogeneity and the need for instruments. Overall, my estimates pass a placebo test as there is no first stage effect for households with dependents that are not yet college-age. A battery of diagnostic tests further confirm the validity of my instruments and identification strategy.

B.3 Appendix Figures and Tables

Figure B-1: Age Distribution of First-Year Students

Notes: This figure reports the age distribution of first-time, full-time, first-year students in the U.S who enrolled in 2015. Age is reported as of the last day of the enrollment year (December 31, 2015). The sample is based on over 10 million students surveyed by the Cooperative Institutional Research Program at the Higher Education Research Institute (CIRP HERI) at UCLA (Eagan, Stolzenberg, Ramirez, Aragon, Suchard and Rios-Aguilar, 2016).

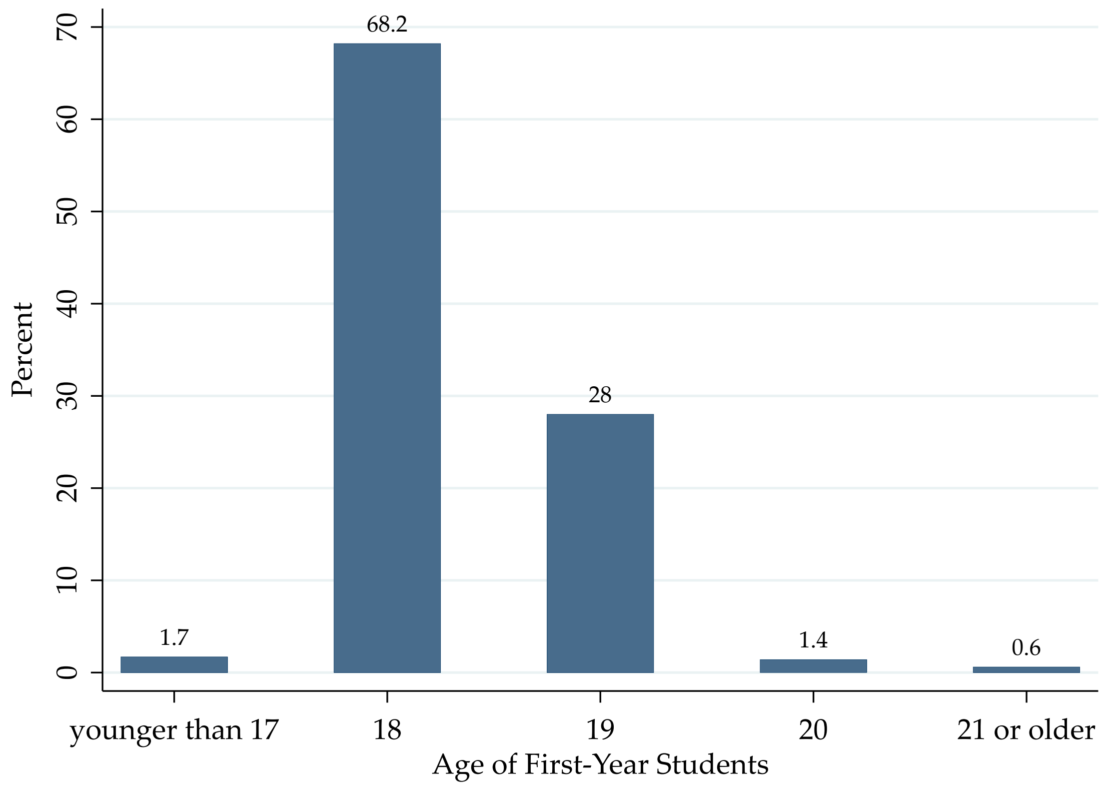

Figure B-2: Business Investment and Exit Rates by Average Growth Rates

Notes: This figure plots β_a of equation 2.4, which captures the effect of self-employed households having an a-year old dependent on business spending on machineries and exit rates by businesses with different growth propensities. Business growth rates are calculated as the average year-over-year revenue growth *before* a child turns 18 years old. Self-employed households are grouped into quartile bins by average growth rates. Bin 1 includes firms that has the lowest pre-18 average growth rates and bin 4 includes those with the highest pre-18 growth rates. The sample is restricted to self-employed households that ever had 18 or 19 year olds during the sample period. All regressions control for the age of the business and its owner, the number of dependents in a household, business industry, state of residence, and employer status of a business. Whiskers show 95% confidence intervals.

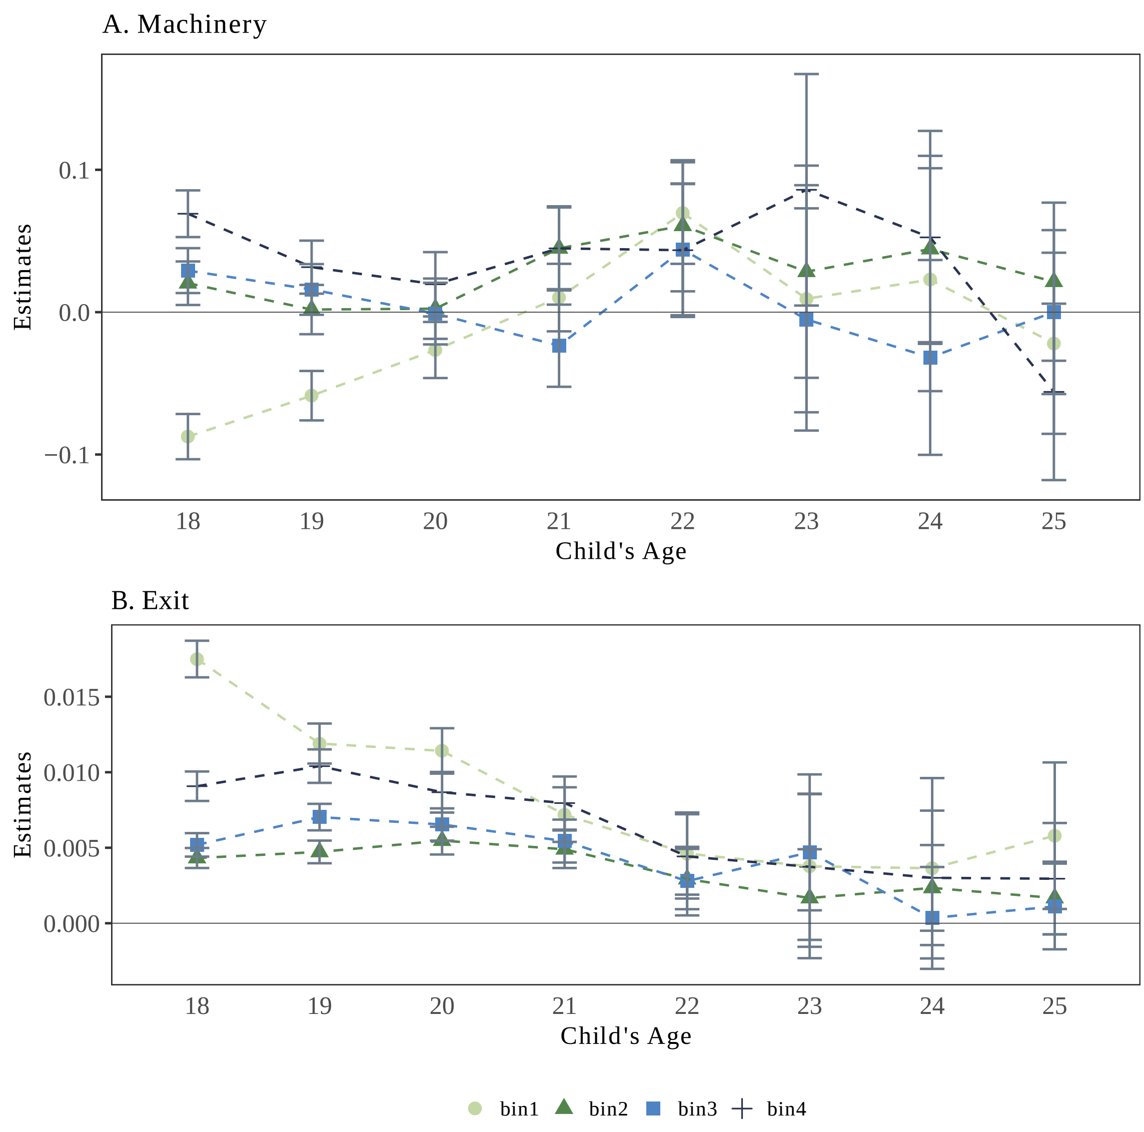

Figure B-3: Household Expenditures by Average Growth Rates

Notes: This figure plots β_a of equation 2.4, which captures the effect of self-employed households having an a-year old dependent on household spending on durable goods and services *net* of education expenditure by businesses with different growth propensities. Business growth rates are calculated as the average year-over-year revenue growth *before* a child turns 18 years old. Self-employed households are grouped into quartile bins by average growth rates. Bin 1 includes firms that has the lowest pre-18 average growth rates and bin 4 includes those with the highest pre-18 growth rates. The sample is restricted to self-employed households that ever had 18 or 19 year olds during the sample period. All regressions control for the age of the business and its owner, the number of dependents in a household, business industry, state of residence, and employer status of a business. Whiskers show 95% confidence intervals.

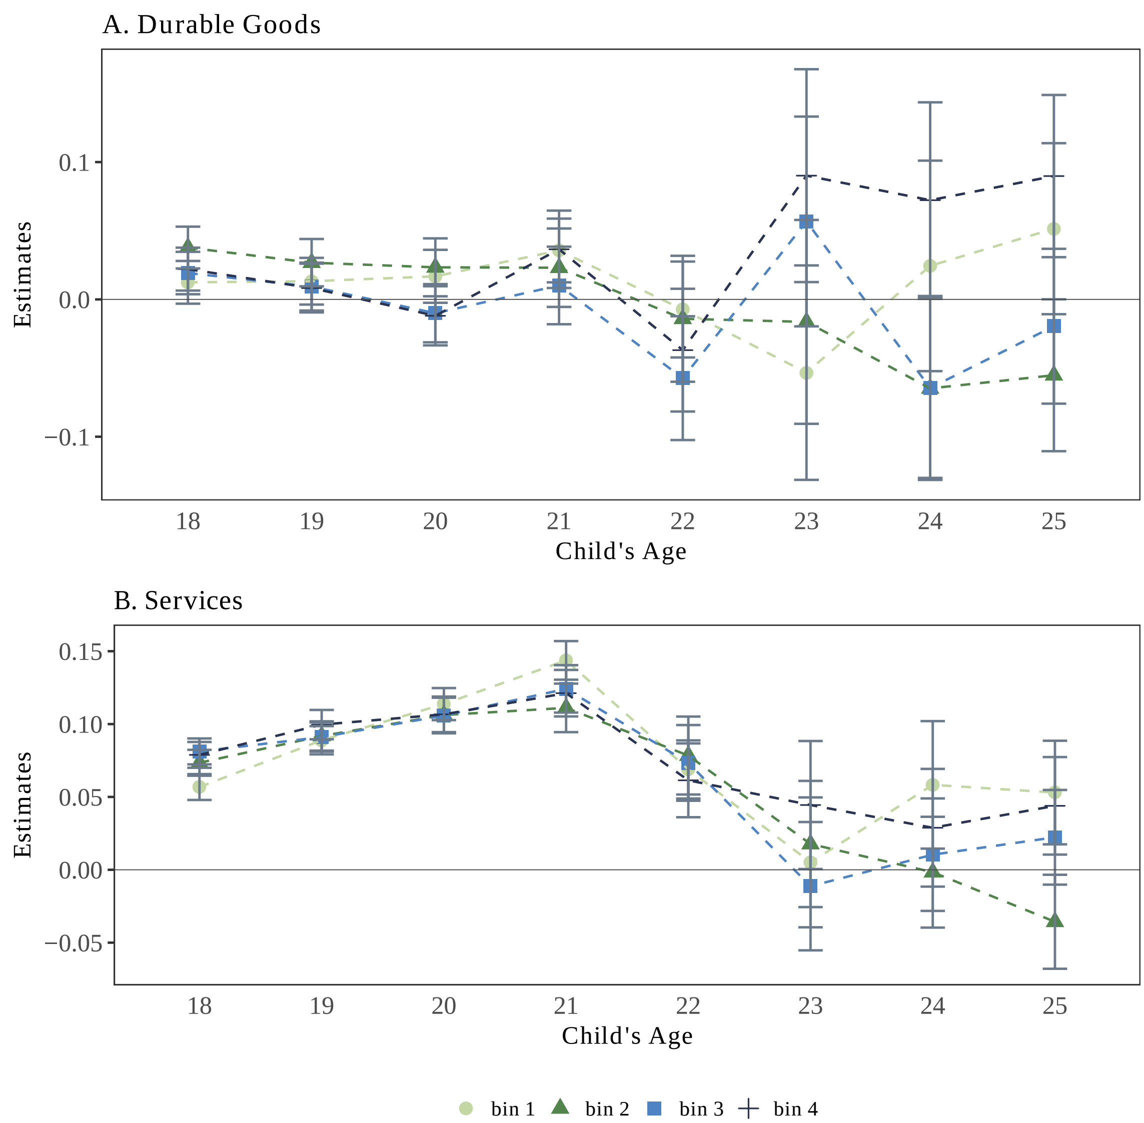

Figure B-4: The Link between Career Choice and Business Production

Notes: This figure plots the estimates reported in table 2.5 by households that do and don't exit from self-employment. Whiskers show 95% confidence intervals. Solid fitted lines are estimated from local regressions, and 95% confidence bands of the fitted lines are shown in grey.

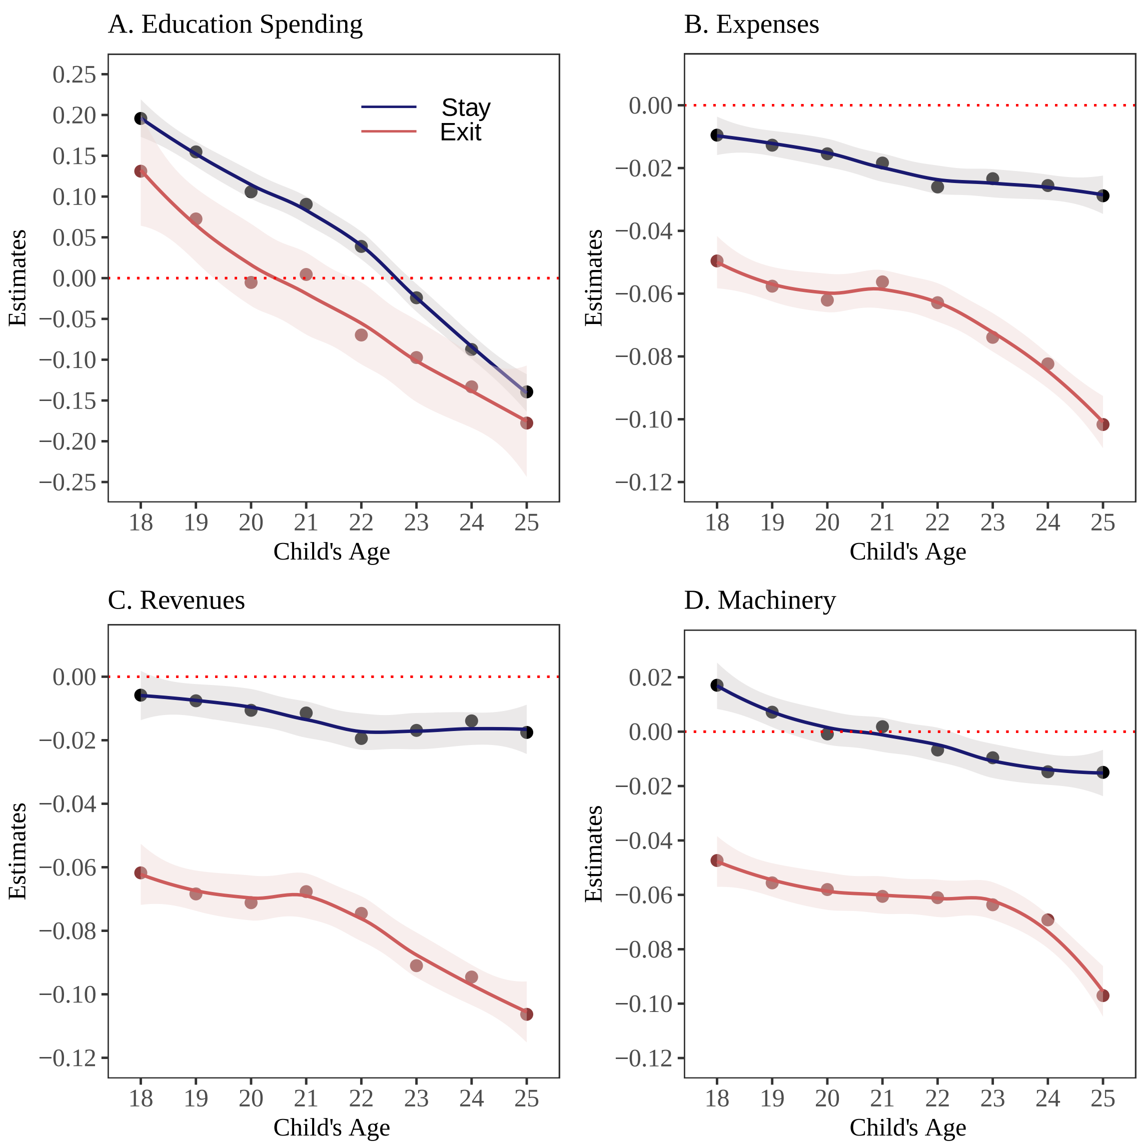

Figure B-5: The Link between Career Choice and Household Consumption

Notes: This figure plots the estimates reported in table 2.6 by households that do and don't exit from self-employment. Whiskers show 95% confidence intervals. Solid fitted lines are estimated from local regressions, and 95% confidence bands of the fitted lines are shown in grey.

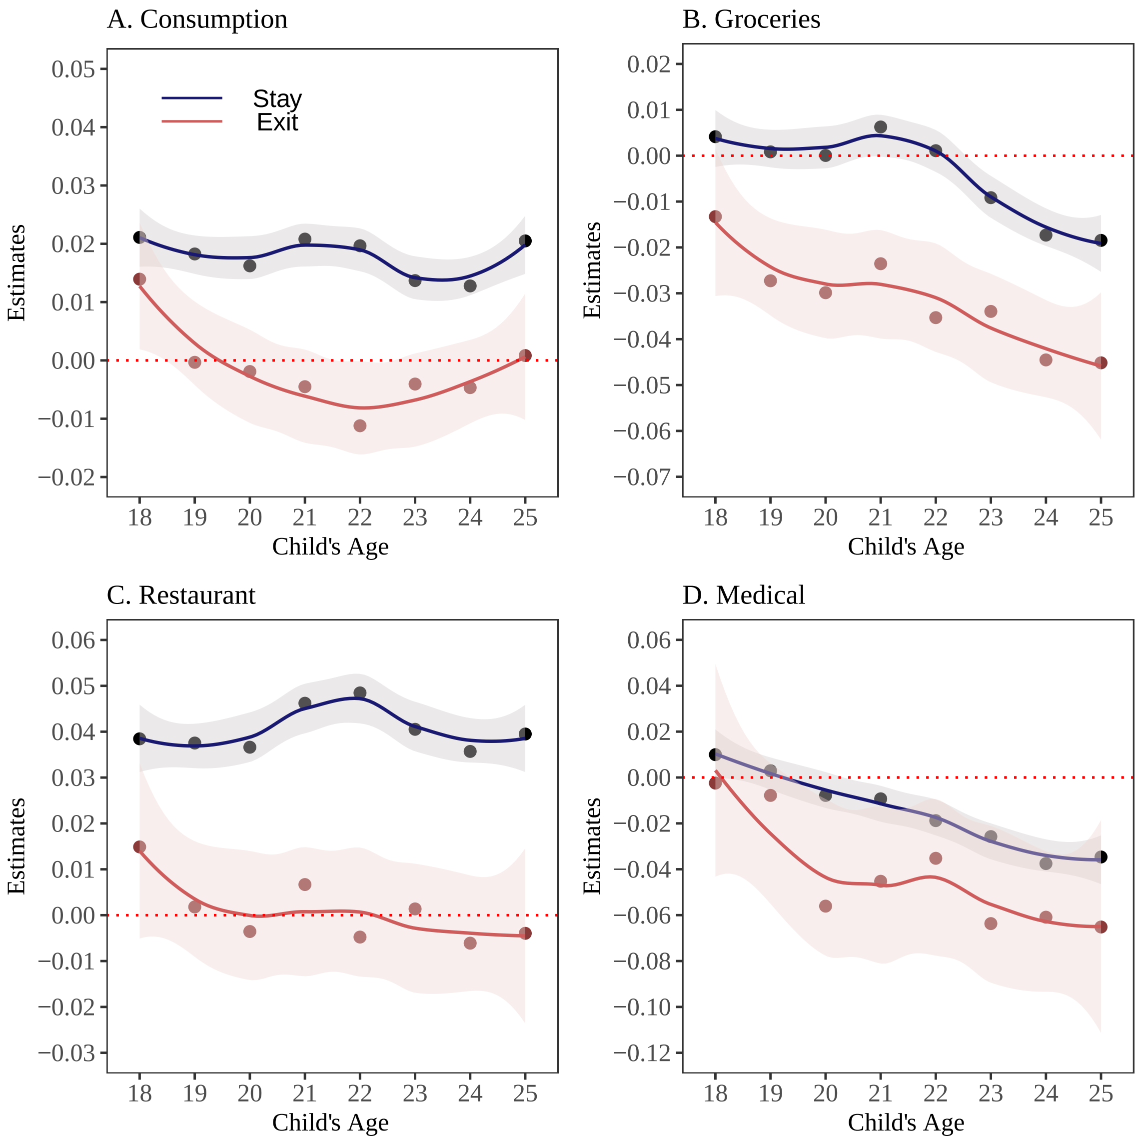

Figure B-6: Exit Rates over a Child's Age Profile

Notes: This figure plots the exit rates from self-employment over a child's age profile. The top panel plots the share of households that exit in each age bin conditional on exiting. The bottom panel plots the cumulative distribution of exit rates. Business exit is inferred from account closures or prolonged inactivity of business checking accounts in the sample period I analyze (2012 Q4 to 2018 Q2). The shaded area in grey indicates the period when the dependents are most likely to be enrolled in college (18-23).

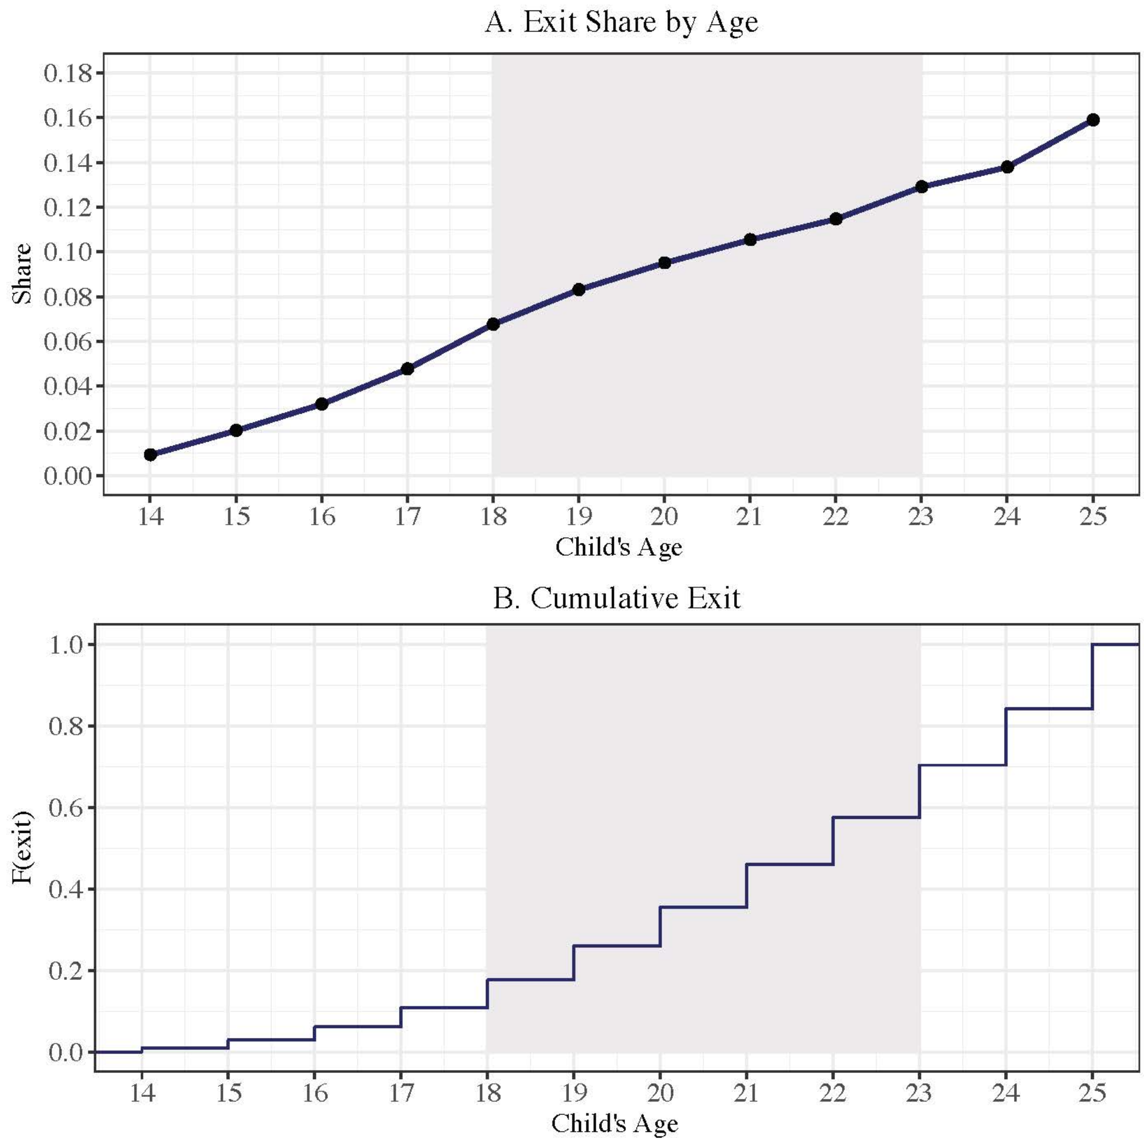

Table B.1: Tuition Payment Types

Notes: This table reports the breakdown of tuition payment types made to a mid-sized research university in the Northeast for the fiscal years 2013 and 2018. Check payments include any paper checks made from students' or parents' bank or 529 account. Other category includes tuition paid through scholarships or lockbox transfers. Wire transfers are all other online payment types excluding check or other payment types.

	Fiscal Year 2013		Fiscal Year 2018	
	% of Transactions	% of Dollars ($)	% of Transactions	% of Dollars ($)
	(1)	(2)	(3)	(4)
Wire	87%	78%	89%	81%
Check	5%	10%	6%	17%
Other	8%	12%	5%	3%

Table B.2: Comparing Near-College Entering and College-Entering Households

Notes: This table reports the average characteristics of self-employed households with near-college entering dependents (ages 15-17) that make up the control group and those with college-entering dependents (ages 18-19) that make up the treated. Parent's age refers to the oldest member in a household and the child's age refers to the oldest child in a household. Net consumption refers to household consumption net of education spending.

	Control Mean (1)	Treated Mean (2)
Number of Family Members	3.2	3.4
Number of Dependents	1.2	1.3
Parent's Age	50.2	51.9
Child's Age	15.9	18.5
Business Years in Operation	5.6	6.1
Share of Employer Firms	0.17	0.16
Share of Households with 529	0.04	0.04
Business Expenses ($)	66,095	64,399
Business Revenues ($)	77,743	75,631
Net Consumption ($)	16,619	16,710

Table B.3: Reduced Form Effects on Business Outcomes

Notes: This table reports the estimated coefficients β_a of the regression equation 2.4. The outcomes are scaled by each household's pre-18 baseline levels if a household exists in both pre-18 and post-18 time period. Otherwise, the scaling factor is the household-specific average of the outcome. The estimates capture the effect of the dependent's age being a on business outcomes of the self-employed households. Standard errors are clustered at the household level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.10

	Outcomes Relative to Pre-18 Levels							
Child's Age	Business Expenses		Business Revenues		Machinery		Exit	
	(1)		(2)		(3)		(4)	
18	-0.013	***	-0.011	***	0.011	***	0.002	***
	(.002)		(.002)		(.003)		(.00)	
19	-0.017	***	-0.014	***	0.001		0.003	***
	(.002)		(.002)		(.003)		(.00)	
20	-0.020	***	-0.017	***	-0.006	*	0.004	***
	(.002)		(.002)		(.003)		(.00)	
21	-0.022	***	-0.017	***	-0.004		0.004	***
	(.002)		(.002)		(.003)		(.00)	
22	-0.030	***	-0.025	***	-0.012	***	0.004	***
	(.002)		(.002)		(.003)		(.00)	
23	-0.028	***	-0.024	***	-0.015	***	0.005	***
	(.002)		(.002)		(.003)		(.00)	
24	-0.031	***	-0.022	***	-0.020	***	0.005	***
	(.002)		(.002)		(.003)		(.00)	
25	-0.036	***	-0.027	***	-0.024	***	0.005	***
	(.002)		(.002)		(.003)		(.00)	
Number of Observations	1,649,212		1,649,212		1,649,212		1,649,212	

Table B.4: Placebo

Notes: This table reports first stage and reduced form β_t restricting the sample to self-employed households with younger dependents (ages 15 - 17). The specification includes interactions of quarter transition dummies with age dummies that equals 1 if a dependent's age is 16 or 17 years old. All regressions control for baseline covariates reported in table 2.3. Standard errors are reported in parenthesis. *** p<0.01, ** p<0.05, * p<0.10

Quarter Transitions $\times\ \mathbb{1}(\text{Age} = a)$	First Stage Education spending		Reduced Form Expense		Revenue		Machinery		Exit	
	(1)		(2)		(3)		(4)		(5)	
Q1 to Q2 × 16	0.025		-0.002		0.003		-0.013		0.000	
	(.025)		(.006)		(.007)		(.019)		(.001)	
Q2 to Q3 × 16	0.052	**	-0.001		0.000		0.006		0.000	
	(.022)		(.005)		(.006)		(.016)		(.001)	
Q3 to Q4 × 16	0.059	***	-0.003		-0.001		0.012		0.001	
	(.021)		(.005)		(.006)		(.016)		(.001)	
Q4 to Q1 × 16	0.021		-0.006		-0.002		-0.037	*	0.000	
	(.026)		(.006)		(.007)		(.020)		(.001)	
Q1 to Q2 × 17	0.168	***	-0.005		-0.001		-0.024		0.002	**
	(.024)		(.006)		(.007)		(.018)		(.001)	
Q2 to Q3 × 17	0.232	***	-0.004		-0.005		-0.006		0.000	
	(.020)		(.005)		(.006)		(.015)		(0.001)	
Q3 to Q4 × 17	0.307	***	-0.012	**	-0.015	**	0.001		0.002	***
	(.020)		(.005)		(.006)		(.015)		(.001)	
Q4 to Q1 × 17	0.145	***	-0.004	**	0.000		-0.019		0.001	
	(.024)		(.006)		(.007)		(.018)		(.001)	
Number of Observations	212,947		212,947		212,947		212,947		212,947	

Table B.5: 2SLS Diagnostic Tests

Notes: This table reports statistics from 2SLS diagnostic tests that validate the assumptions behind my instruments based on the 2SLS results in table 2.3. My instruments are Quarter Transitions $\times$ $\mathbb{1}$(College-going age). Weak instruments diagnostics use the first stage F-test to validate the relevance of my instruments. Wu-Hausman tests the consistency of the OLS estimates under the assumption that the IV is consistent.

	Business Expense		Business Revenues		Exit	
	statistics	p-value	statistics	p-value	statistics	p-value
	(1)	(2)	(3)	(4)	(5)	(6)
Weak Instruments	744.61	0.00	744.61	0.00	744.61	0.00
Wu-Hausman	350.84	0.00	208.10	0.00	2.74	0.00

Table B.6: Variable Descriptions

Variable Name	Definition
Firm	A collection of business checking accounts linked to the same owner.
Education Spending	Any payments to post-secondary institutions (e.g., tuition, room, and board), testing service agencies, student loan servicing companies, and savings to 529 accounts. The measure does not include paper checks to post-secondary institutions, room and board expenses if a student lives off-campus, cost of health or dental insurance for a child, discretionary financial support that a family provides for a child, or additional off-the-book borrowing from friends or family.
Business Expense	Operating expense out of business checking accounts. Excludes financial transactions, such as transfers between accounts or fee payments, that are unlikely to capture the actual cost of operating the business. It does not include spending on education.
Business Revenue	Operating revenues into business checking accounts. Excludes financial transactions, such as transfers between accounts or fee reversals, that are unlikely to capture the actual revenues incurred from operating the business.
Exit	A closure or inactivity of business accounts. Inactivity is defined as having less than $500 in outflows and less than 10 transactions for 3 out of 12 consecutive months.
Household Consumption	Any durable and non-durable spending from personal checking accounts. Consumption includes goods (e.g., groceries, fuel, home improvement, etc), services (e.g., restaurants, doctor's visits, air fares, etc), other uncategorizable bill payments using Paypal or wire transfers, utilities (e.g., phone bills, internet, cable), non-housing debt payments (auto, personal, or student loans), housing debt (HELOC, mortgages), and credit card payments.
Non-business income	Any income that does not come from running a business. Includes labor income (e.g., direct deposit or payroll), capital income (e.g., investment income from pensions or annuity accounts), government benefits (e.g., transfers or tax refunds), and other miscellaneous direct deposits. Transfers from business to household accounts are excluded.

Bibliography

Eagan, Kevin, Ellen B. Stolzenberg, Joseph J. Ramirez, Melissa C. Aragon, Maria R. Suchard, and Cecilia Rios-Aguilar. 2016. "The American Freshman: Fifty-Year Trends, 1966-2015." Los Angeles: Higher Education Research Institute, UCLA.

Evans, David S., and Boyan Jovanovic. 1989. "An Estimated Model of Entrepreneurial Choice under Liquidity Constraints." *Journal of Political Economy*, 97(4): 808–827.

Holtz-Eakin, Douglas, David Joulfaian, and Harvey S. Rosen. 1994. "Sticking it Out: Entrepreneurial Survival and Liquidity Constraints." *Journal of Political Economy*, 102(1): 53–75.

Appendix C

Appendix for Chapter 3

C.1 Administration of the 7(a) loan program

This section provides additional detail on the administration of SBA loan programs. The SBA oversees various assistance programs, such as the Lending Programs, Entrepreneurial Development Programs, and Federal Contracting and Assistance Programs, which provide loan guarantees to small businesses. The maximum loan size limit is capped at $5 million, and the use of proceeds ranges widely from traditional term loans to debt refinancing. Since there is no formal limit as to how much SBA loans a given lender can underwrite, the Office of Credit Risk Management monitors lender performance and oversees the growth of loan portfolios of banks.

While loan maturity depends largely on borrower's ability to repay, loans for working capital, machinery, and equipment have a maturity of up to five to ten years while loans for real estate have a maturity of up to 25 years. Lenders and borrowers can negotiate the interest rate, but it may not exceed the maximum rate set by the SBA. The maximum interest rates are based on a loan amount and maturity such that they decrease in loan amount and increase in loan maturity within two tiered maturity groups defined by a seven-year maturity mark.

A new lender that is not familiar with the SBA loan submission process uses the General Program (GP). Under this program, the lender submits a full application requesting SBA guarantee to the Loan guarantee Processing Center (LGPC). The more experienced SBA lenders are given the "delegated" lender status. Experienced lenders that have met certain performance standards are eligible to use the Certified Lender Program (CLP). Under the CLP, a lender undergoes the same application process as non-delegated lenders, but the SBA expedites the loan processing and services. The most experienced lenders use the Preferred Lender Program (PLP). PLP lenders have the authority to process, service, or close any SBA loans without SBA's prior approval.

There are benefits and costs associated with becoming an SBA lender. A key benefit is that the SBA guarantee helps lenders mitigate credit risks while allowing them to expand their customer base by serving borrowers who may not meet the conventional lending requirements. From a regulatory perspective, since the risk weight of guaranteed loans is lower than for unguaranteed loans, the 7(a) guarantee lowers a lender's risk-weighting for meeting the Basel II capital requirements. SBA loans also have the potential to receive Community Reinvestment Act (CRA) consideration if the loans meet the definition of "loans to small business."

The costs for lenders include a one-time guarantee fee, annual ongoing servicing fee for each loan approved and disbursed, and other applicable fees associated with ongoing SBA oversight, late payment, or packaging and other services. The lender is required to submit the one-time guarantee fee with the loan application for loans with maturities of 12 months or less, and within 90 days of the date of the loan's approval for loans with maturities exceeding 12 months. This guarantee fee is based on the loan maturity and the guaranteed portion of the loan.[1] Lenders could pass-through this one-time guarantee to borrowers, and borrowers in turn may use loan proceeds to pay the guarantee fee in the initial disbursement. The annual ongoing servicing fee is set at the time of loan approval and based on the outstanding principal balance of the guaranteed portion of the loan. In fiscal year 2018, this fee is set to 0.55% of the outstanding balance of the SBA's share. Note that this cost structure could incentivize the lenders to not always charge the maximum allowable interest rates and guarantee rate on loans to reduce the amount of fees paid to SBA.

Table C.2 reports the industry breakdown of the borrowers that receive SBA loans. In our sample, small businesses in accommodation and food services industry receive SBA loans most

[1]For any short-term loans with maturities of 12 months or less, the fee is 0.25% of the guaranteed portion of the loan. For loans with longer maturities, loans of $150,000 or less require 2%; loans of amount greater than $150,000 but less than $700,000 require 3%; loans of amount greater than $700,000 but less than $1 million require 3.5%; and loans of size greater than a million require 3.75% of the guaranteed portion of the loan.

frequently (i.e., 18% of all loans), and the top 10 industries make up nearly 90% of all loans originated to small businesses. Small businesses in accommodation and food services industry is over-represented in the SBA data when compared to the industry composition of small businesses at the national level, where businesses in this industry only make up 8% of all small businesses. On the other hand, businesses in professional services and construction are under-represented in the SBA sample. In other industries, SBA industry composition line up well with the industry composition at the national level.

C.2 Calculation of Γ

We calculate Γ_i as the expected guarantee subsidy for a loan, net of expected charge-offs, guarantee fee payments, and guarantee reimbursements using the following steps:

1. Run the logistic regression $\pi_i = \alpha + \beta D_i + e_i$, where π_i is an indicator that a loan was charged off and D_i is the loan size. We use all loans in the estimation sample, pooling across multiple years.

2. Predict back $\hat{\pi}_i$, the expected charge-off probability, for each loan in the sample. Since the specification in step one did not have year effects, $\hat{\pi}_i$ is not time-varying. This means that all variation in Γ_i over time comes from policy-driven changes in the guarantee rates.

3. Calculate the expected reimbursement rate (as a percentage of loan principal) as the guarantee reimbursement rate for that particular loan multiplied by the expected charge-off probability: $\gamma_i * \hat{\pi}_i$. The reimbursement rate varies by loan size and by year.

4. The yearly fee is expressed as the amount paid *yearly* as a percent of loan principal. We therefore multiply it by the term of the loan to convert it to the same units as the one-time fee and the expected charge-off probability. We add together the converted yearly fee and one-time fee to get the net fees paid on the loan: σ_i. These fees vary by loan size and by year.

5. We then calculate Γ_i for each loan in the sample as: $\Gamma_i = \gamma_i * \hat{\pi}_i - \sigma_i$.

C.3 FDIC Statistics on Depository Institutions

This section describes the FDIC Statistics on Depository Institutions (SDI) data used in the paper, and our construction of shares. The SDI data records the total number and amount of small business loans outstanding at a quarterly level per institution, and further splits small business lending into categories of loan size and purpose. We specifically look at small business commercial and industrial loans. The FDIC SDI statistics will *include* SBA lending by a particular institution– therefore when combined with the SBA data they allow us to calculate the bank-specific "share of small business lending that is through the SBA". We observe the yearly stock of loans outstanding in the SDI data, and the yearly “flow" of SBA loans in the SBA data. Therefore, we convert the SDI data into a comparable flow measure, and then calculate the bank-year specific SBA share as follows:

1. From the SDI report we observe the stock of number of small business loans from a bank in a given year.
2. We divide this stock by the average maturity (10 years) to get the approximate flow of small business loans from that bank.
3. Calculate from the SBA data the flow of SBA small business loans in a given year.
4. Calculate the bank-year specific SBA share as $\frac{flow\ of\ SBA\ loans}{flow\ of\ all\ small\ business\ loans}$ in a given year.

This calculation generates a distribution of high to low intensity SBA lenders. Banks that lend primarily through the SBA have less ability to substitute between their SBA and non-SBA portfolios. Therefore if we find a similar response to the guarantee across the SBA share distribution, it is unlikely that the portfolio substitution response has biased our elasticity estimates.

C.4 Data appendix: 2003 Survey of Small Business Finances

The 2003 Survey of Small Business Finances (SSBF) is the fourth survey of U.S. Small businesses conducted by the Board of Governors, and the last wave before the releases of the Small Business

Credit Surveys. The survey collected information on firm and owner characteristics, an inventory of small businesses' use of financial services and of their financial service suppliers, and income and balance sheet information.

The data set for the 2003 Survey of Small Business Finances contains information on 4,240 small businesses that were in operation during December 2003 and at the time of the interview. The interview for most firms took place between June and December in 2004. The reference date for most questions is the date of the interview; the reference date for the income statement and balance sheet information is the date of the firm's most recent fiscal year-end and can range from July 1, 2003 to June 30, 2004. For the 2003 release, the SSBF data set includes five implicates. Each implicate includes 4,240 firms. In total, the entire data set contains 21,200 observations. There are 4,240 firm observations in total. There are in total 225 firms which took loans from a government agency, including the SBA.

Appendix Table C.3 shows the fraction of firms that access credit from more than one source in the past three years. The table indicates that very few firms access credit from more than one source. Appendix figure C-11 shows the fraction of firms by the number of lending institutions dealt with. The fact that many firms deal with many lending institutions, but only borrow from one (typically an SBA lender) is indicative of inability to obtain credit elsewhere.

C.5 Appendix Figures and Tables

Figure C-1: Simulated changes in loan size by n and Γ

Notes: This figure simulates how loan size responds to a varying type of the underlying variance distribution n_i (top) and to Γ for high and low variance distribution of expected returns.

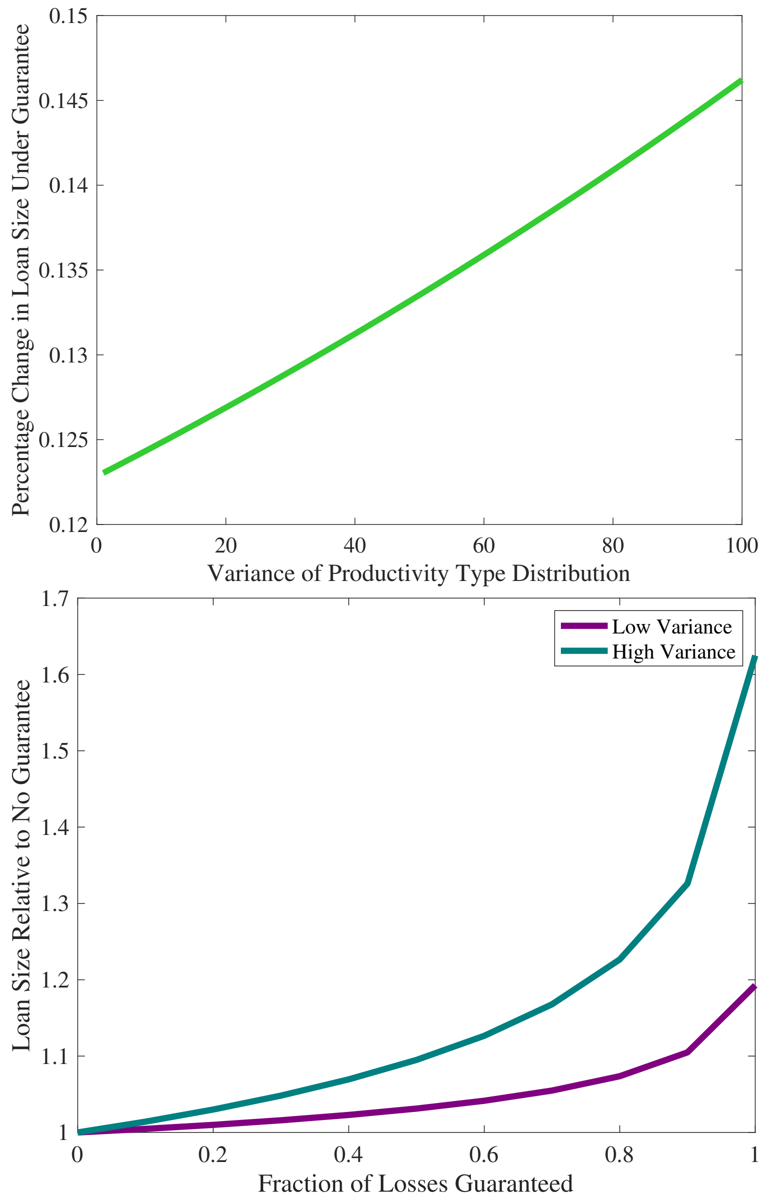

Figure C-2: Bunching at the guarantee notch, wider axis in placebo years

Notes: This figure shows the number of loans made in discrete $2,000 bins across the threshold. The graph includes years 2009 and 2010, when the guarantee notch was eliminated, with an alternative wider axis. Note bunching at round numbers, which is controlled for in the elasticity estimate. Source: SBA.

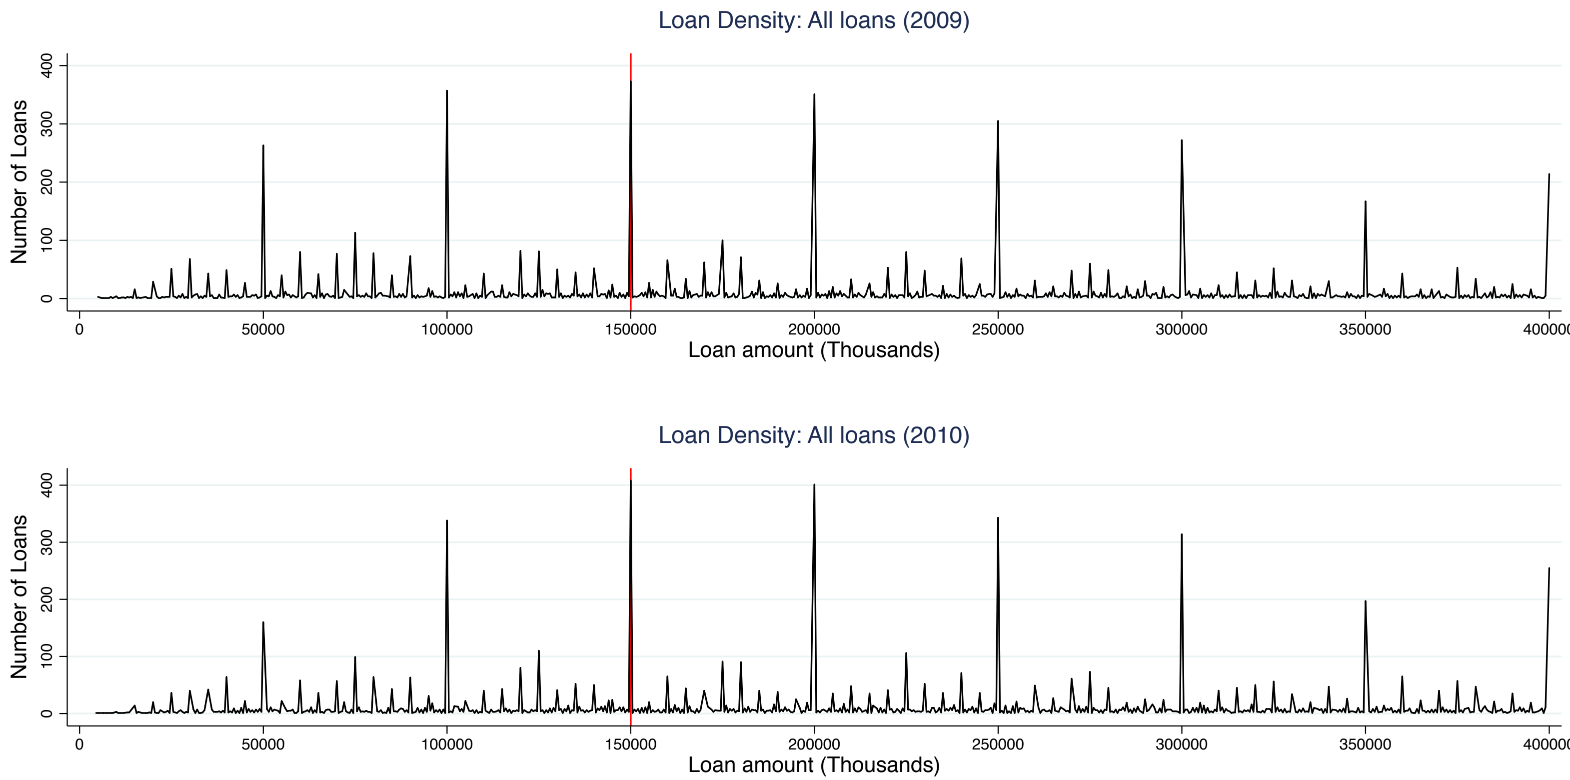

Figure C-3: Three-Year cohort default rates over time

Notes: This figure shows the 3-year cohort default rates over time. We exclude all loans originated after 2015 in this graph to ensure that every loan in the sample has a valid 3-year cohort default rate. Source: SBA.

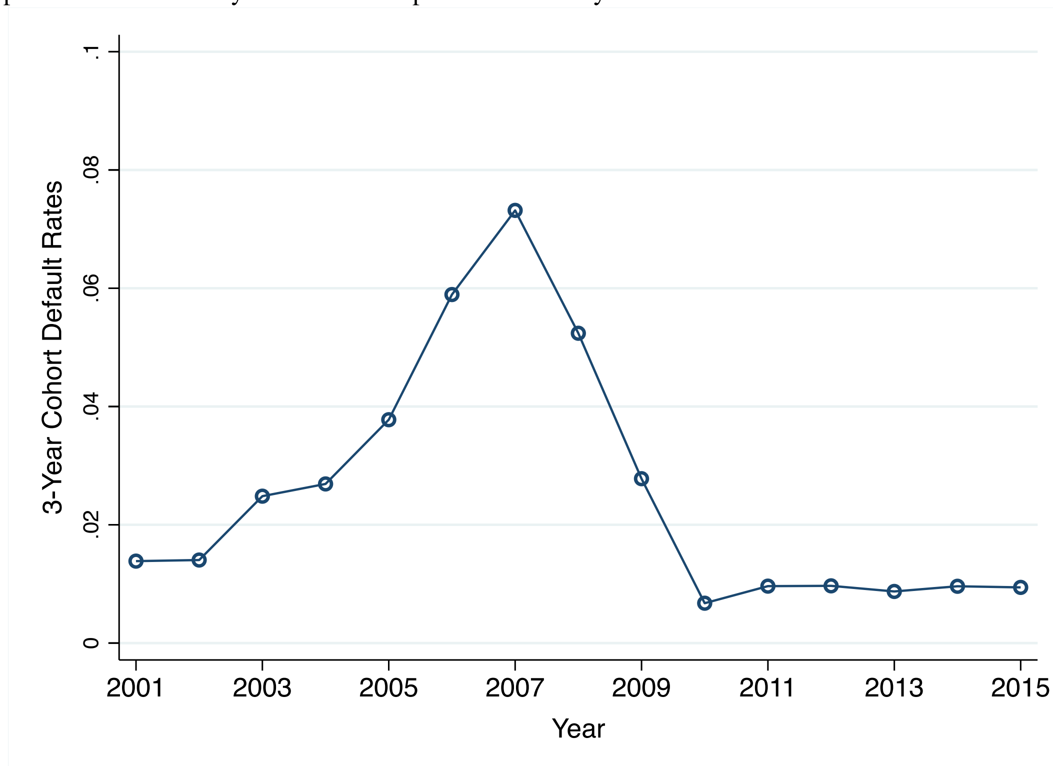

Figure C-4: Guarantees and fees by loan amount

Notes: This figure shows the average expected guarantee fees and reimbursement rate as a percentage of the loan principal amount for discrete 2000 bins across the threshold. The graph pools over all years 2008-2017. Source: SBA.

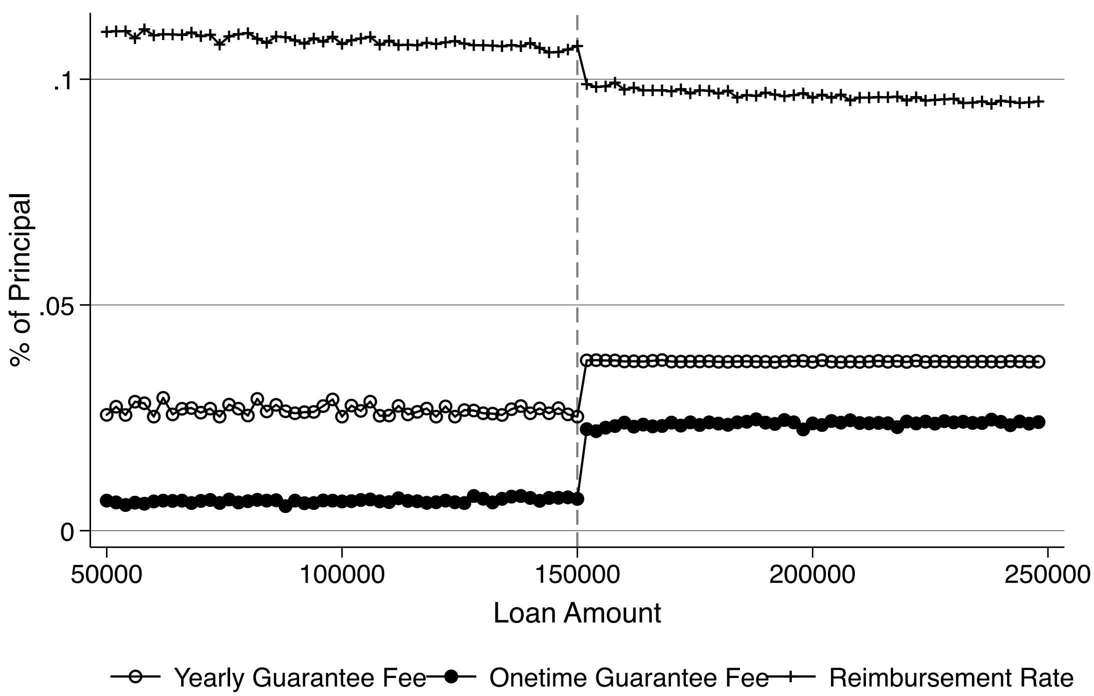

Figure C-5: Observed and estimated loan density for elasticity estimation during placebo years

Notes: This figure plots the observed loan density (black) and the estimated counterfactual density (red) for 2009 and 2010, the "placebo” years, when the guarantee notch did not exist. This allows us to directly test the fit of our estimated counterfactual distribution against years when there was no discontinuity at $150,000. For estimation, we restrict the loan size to be between $75,000 to $225,000. The counterfactual is estimated for each notch separately by fitting a 6th-order polynomial with round-number fixed-effects to the empirical distribution using step size of 500, and excluding data around the notch, as specified in equation 3.15. The missing mass at the threshold is measured as the distance between the black and red lines at $150,000.

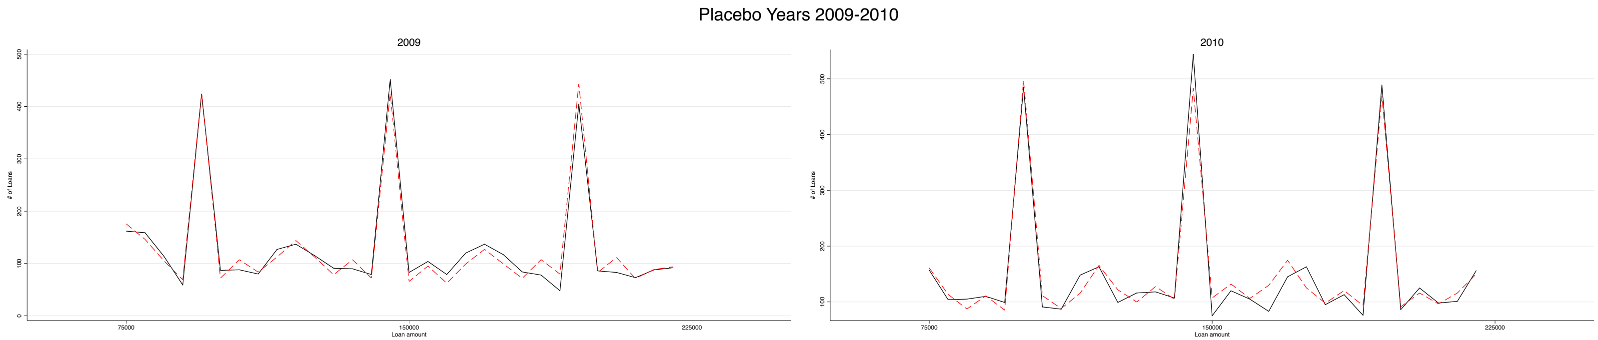

Figure C-6: Bunching at the guarantee notch, wider axis

Notes: This figure shows the number of loans made in discrete $2,000 bins across the threshold. The graph pools over all years 2008-2017 with an alternative wider axis. Note bunching at round numbers, which is controlled for in the elasticity estimate. Source: SBA.

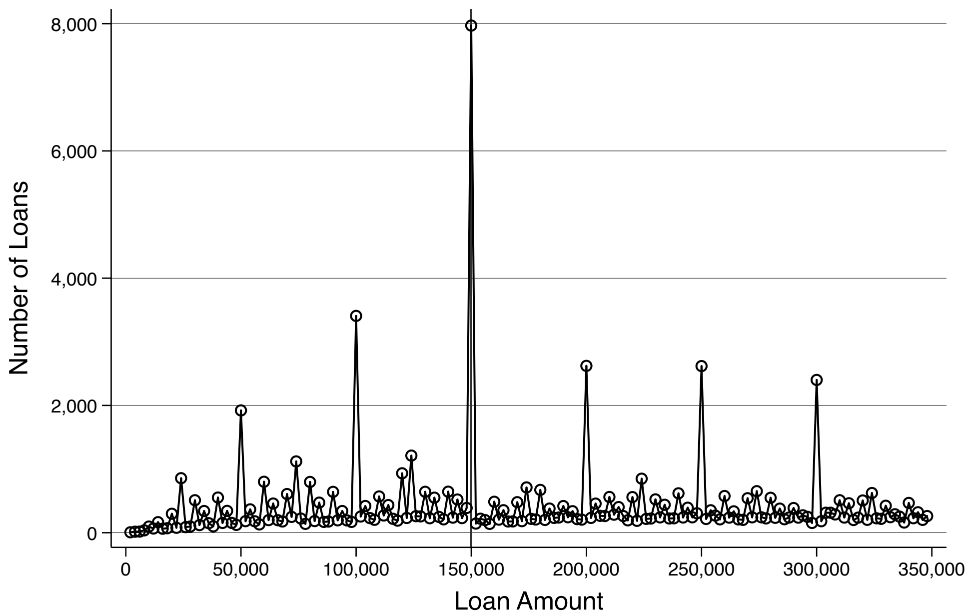

Figure C-7: Percentage of loans at the binding interest rate maximum

Notes: This figure shows the percentage of loans made at the maximum interest rate cap in discrete $2,000 bins across the threshold. The graph pools over all years 2008-2017, absorbing year-month effects and bank fixed effects. Source: SBA.

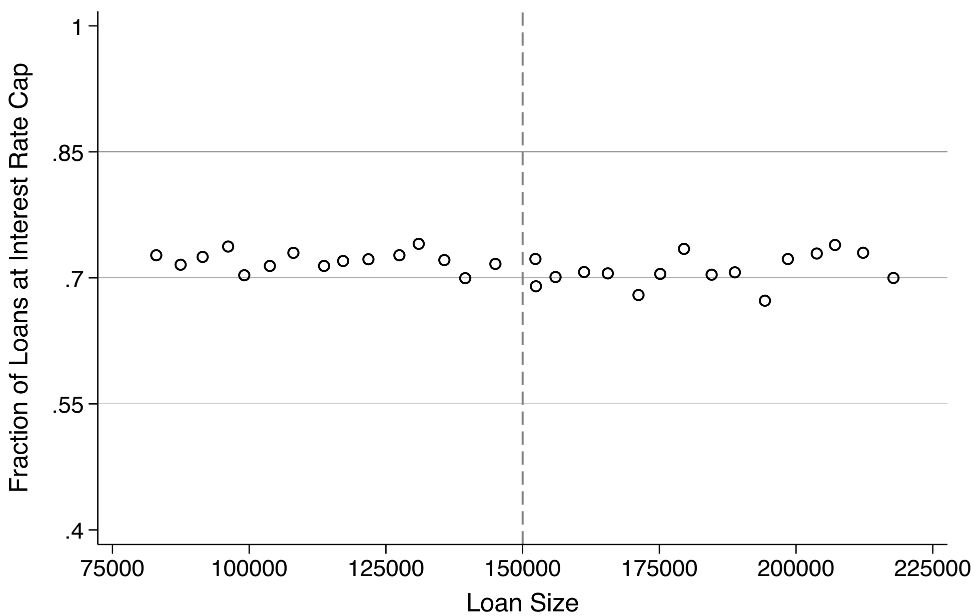

Figure C-8: Other variables at the guarantee notch

Notes: This figure plots the average interest rate, revolving loan percentage, charge-off percentage, and loan term across the threshold. They are normalized with respect to the value of the variable at the threshold. There is no significant difference in initial interest rate, the percentage of revolving loans, the charge-off percentage across the threshold. Note the presence of round number bunching in the bottom right panel. The graph pools over all years 2008-2017, absorbing year-month effects and bank fixed effects. Source: SBA.

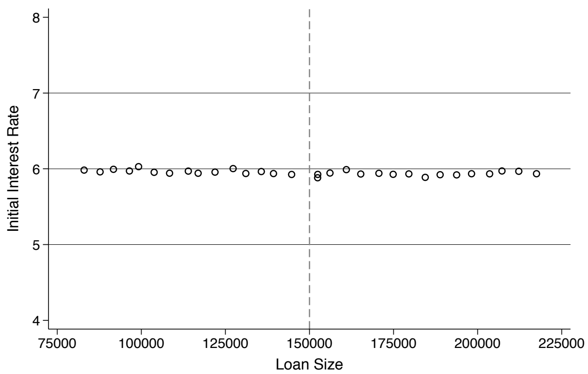

(a) Average interest rate

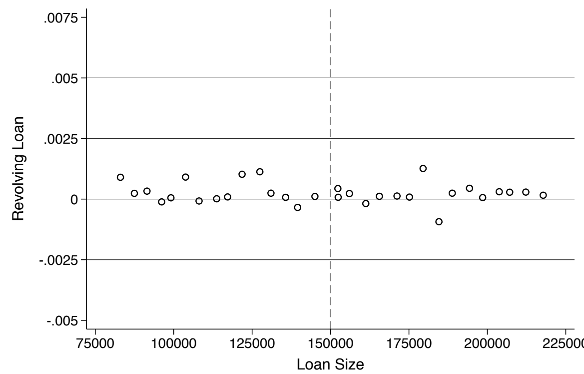

(b) Revolving loan percentage

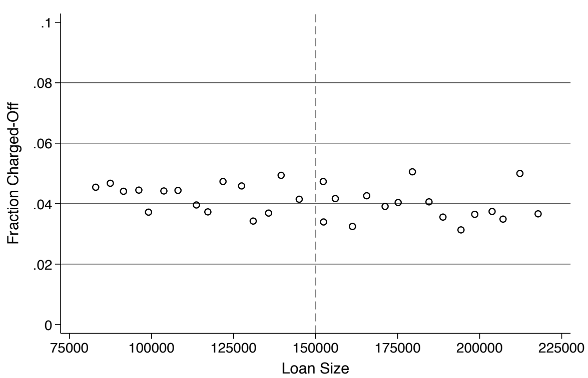

(c) Charge-off percentage

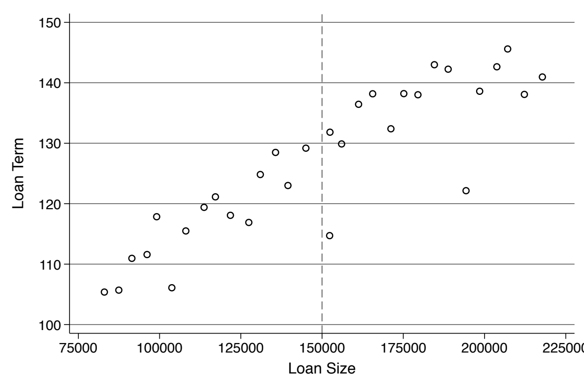

(d) Loan term

Figure C-9: Substitution: lending supply response by year

Notes: This figure plots the coefficient ζ_t on the interaction term between the bank's pre-ARRA propensity to issue SBA loans and year indicators in equation 3.6.2, along with 95% confidence interval. The dependent variables are log of total, non-SBA, or SBA small business loans. The baseline is 2017. Total and non-SBA small business lending are from FDIC SDI and converted into flows as described in appendix C.3. The sample is restricted to loan size between $\$50,000$ and $\$225,000$, and for banks that operated in 2008 such that it has non-missing pre-ARRA exposure. The regressions are estimated at the bank-year level. Standard errors are clustered at the bank-level and reported in parentheses. Source: SBA and FDIC SDI.

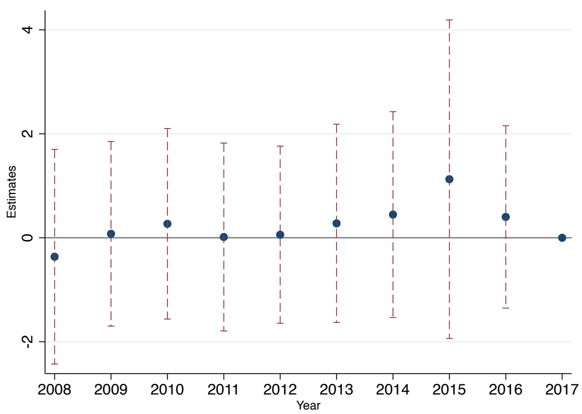

(a) Total small business loans

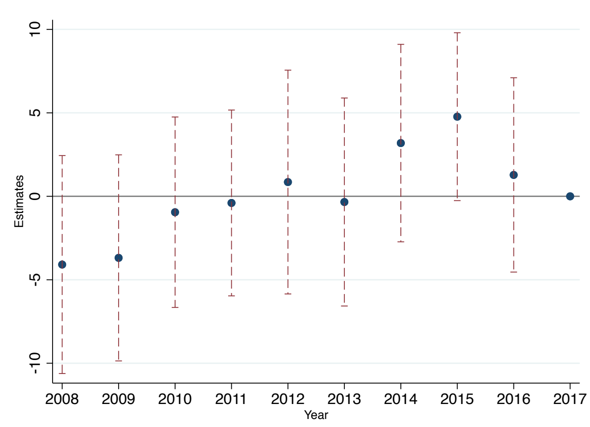

(b) Non-SBA small business loans

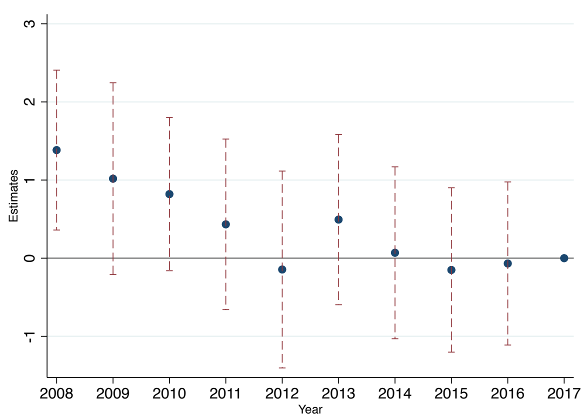

(c) SBA loans

Figure C-10: Estimates varying beginning and ending limit

Notes: This figure shows elasticity estimates varying the starting and ending range. The top panel varies the starting range, while the bottom panel varies the ending range around the $150,000 threshold. A sixth order polynomial is used. Source: SBA.

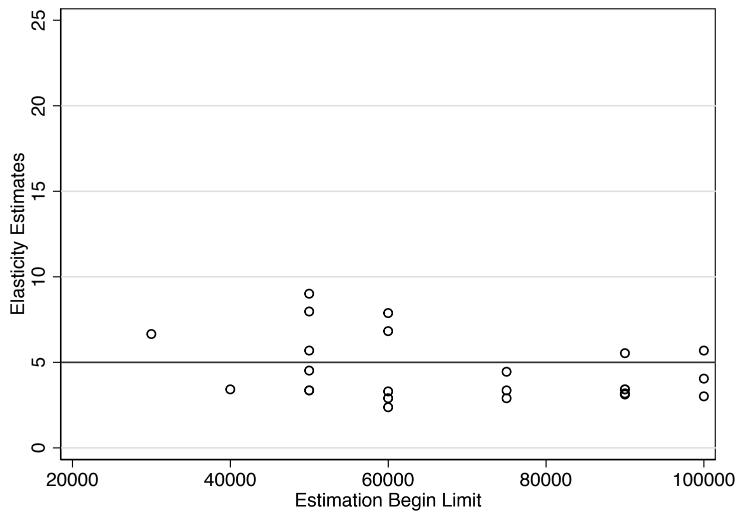

(a) Starting limit

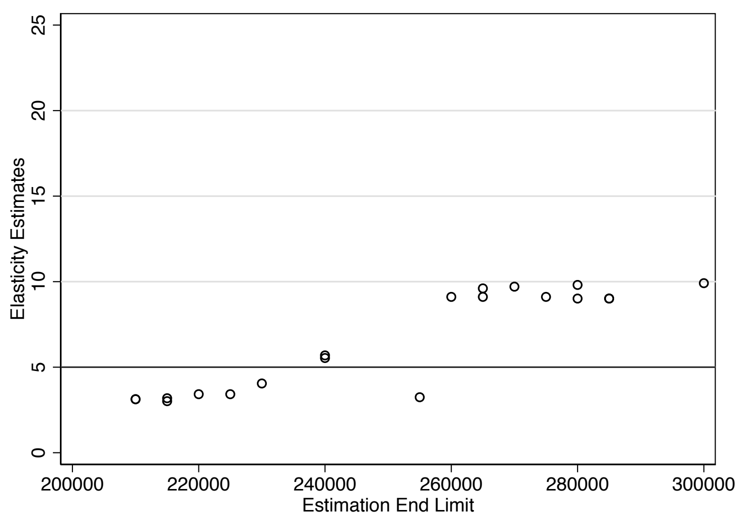

(b) Ending limit

Figure C-11: Number of lending institutions

Notes: This figure shows the fraction of firms by the number of lending institutions that a small business dealt with in the past 3 years. The sample is restricted to firms with a loan from a government agency, including the SBA Source: SSBF.

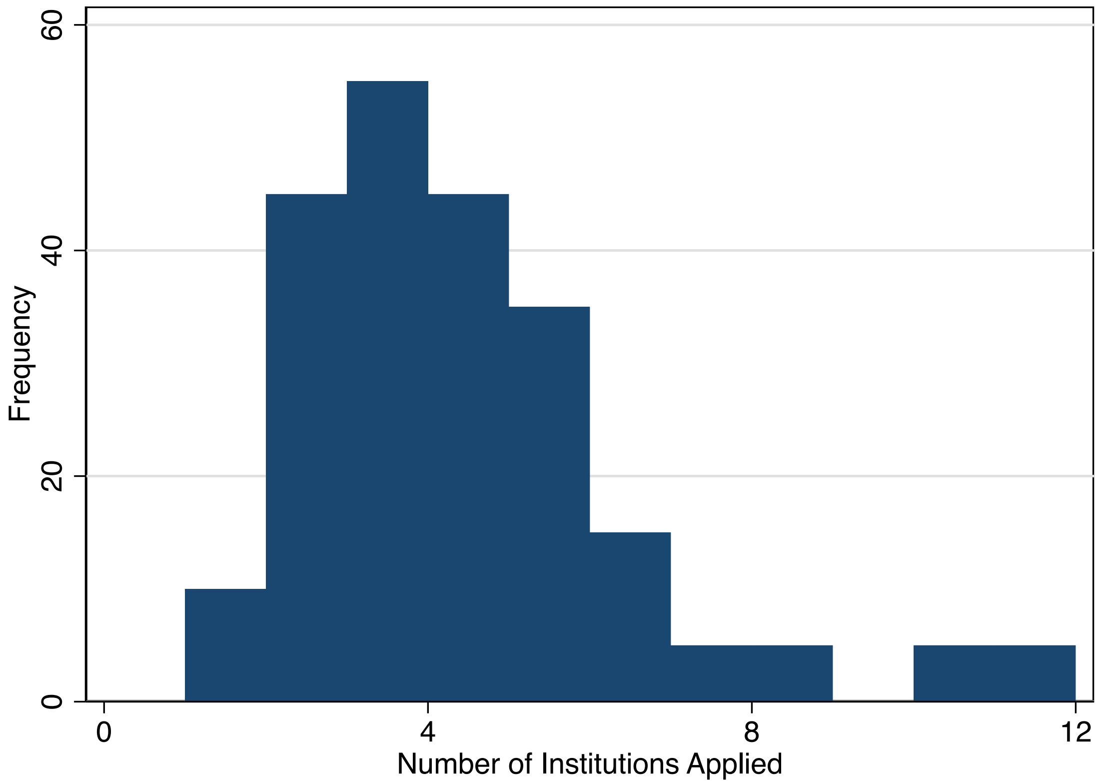

Figure C-12: Observed and estimated loan density for elasticity estimation

Notes: This figure plots the observed loan density (black) and the estimated counterfactual density (red) for each year. We separately show the years in which a notch at the $150,000 threshold existed, and when it did not (2009 and 2010, the "placebo" years). For estimation, we restrict the loan size to be between $75,000 to $225,000. The counterfactual is estimated for each notch separately by fitting a 6th-order polynomial with round-number fixed-effects to the empirical distribution using step size of 500, and excluding data around the notch, as specified in equation 3.15. The missing mass at the threshold is measured as the distance between the black and red lines at $150,000.

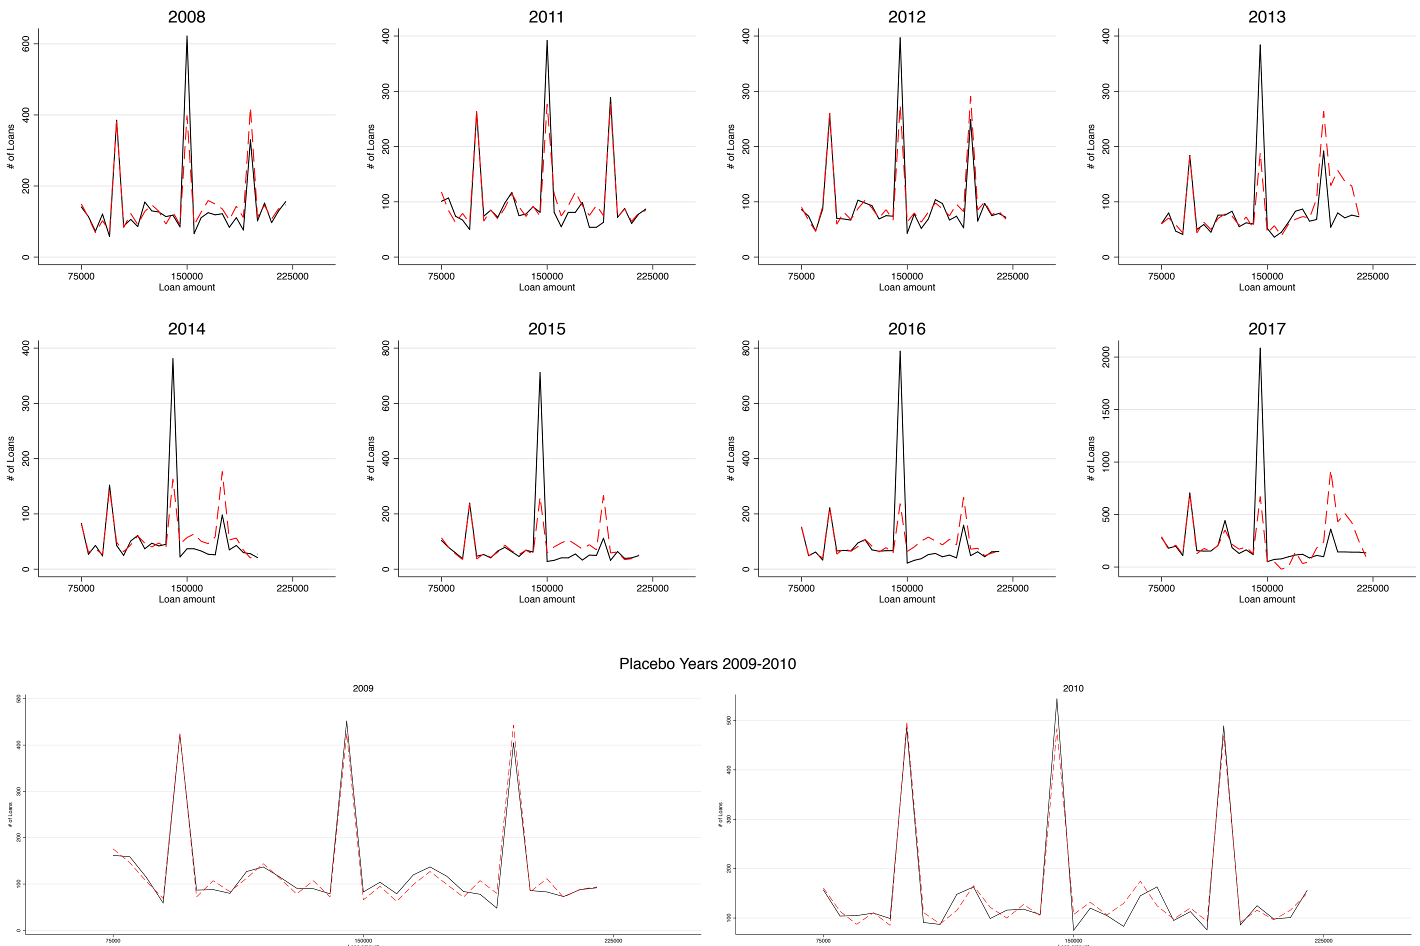

Figure C-13: Observed and estimated loan density by SBA share and market concentration

Notes: This figure plots the observed loan density (black) and the estimated counterfactual density (red) for subsamples of banks with high/low SBA lending shares and in high/low concentration markets. The first row splits the sample into banks who issue fewer than 80% of their small business loans through the SBA (left), and banks with $\geq$ 80% of their small business lending through the SBA (right). The second row splits the sample by banks in regions with fewer than 7 SBA lenders (left), or $\geq$ 7 lenders (right).

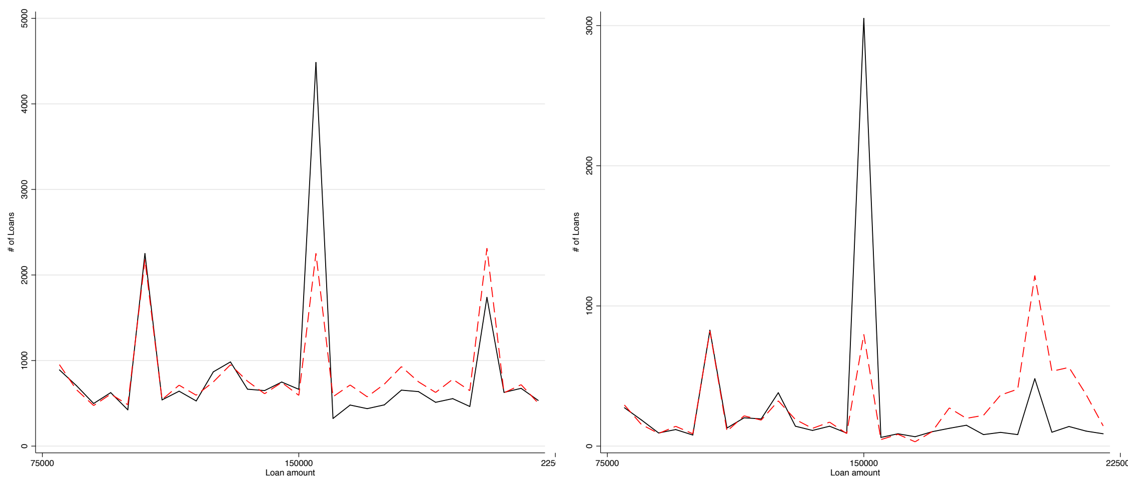

(a) Sample split by SBA share $\leq 80\%$ (left) and $> 80\%$(right)

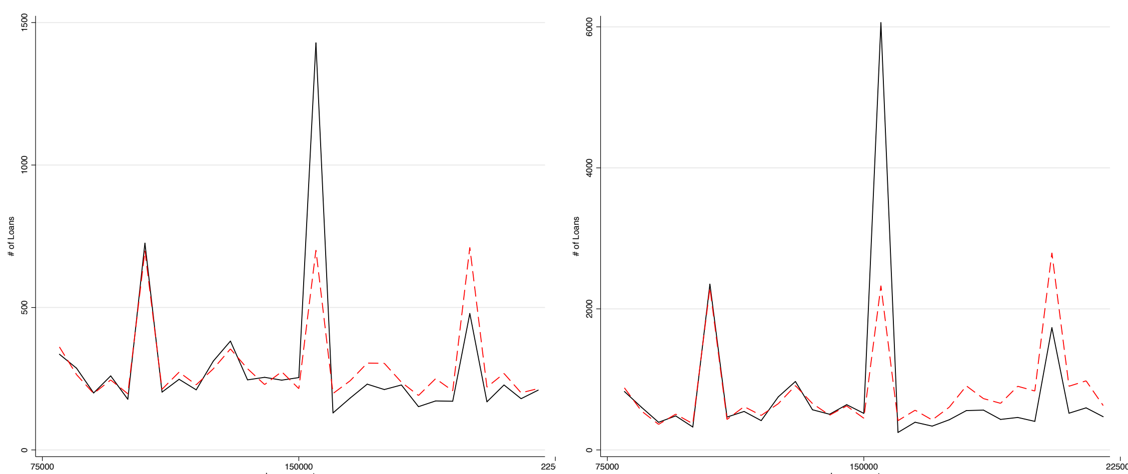

(b) Sample split by regions with ≤ 7 (left) and > 7 (right) unique banks

Table C.1: Variable descriptions

Notes: This table reports the main analysis variables, their definitions, and source.

Variable name	**Definition**	**Source**
Loan amount	Total loan amount in dollars.	SBA
Reimbursed amount	Amount of SBA's loan guarantee.	SBA
Charge-off amount	Total loan balance charged-off (includes guaranteed and non-guaranteed portion of loan.)	SBA
Interest rate	Initial interest rate at the time loan was approved (base rate plus spread.)	SBA
Reimbursement rate	Total guarantee rate for loans. For most years, 85% guarantee for loans of \$150,000 or less; 75% guarantee for loans greater than \$150,000 (up to \$3.75 million maximum guarantee.)	Derived from SBA
Maximum rate	Maximum interest rate a bank can charge a borrower.	SBA. LIBOR from BNY Mellon
Maturity	Length of loan term	SBA
Yearly fee	A yearly fee that a lender must pay to SBA for each loan guaranteed under the 7(a) program. Based on the guaranteed portion of the loan and not the total loan amount. This fee cannot be passed on to the borrower.	SBA
One-Time fee	One-time guarantee fee that a borrower pays the SBA to obtain a loan.	SBA
Average expected guarantee benefit	Predicted guarantee amount as a share of loan principal net of one-time and yearly fees, assuming 100% charge-off.	Derived from SBA
Excess mass	The amount of bunching at the \$150,000 notch computed as the difference between the observed and counterfactual bin counts between the lower limit of the excluded region (d_l) and the threshold (D^T).	Estimated following Kleven and Waseem (2013)
Share of excess mass	Excess mass as a share of the total number of loans in the estimation range.	Estimated

Table C.2: Industry breakdown

Notes: This table reports the industry breakdown of the borrowers that received loans in the full sample. Industries are grouped by NAICS 2-digit sector code. The second and third columns report the number of loans by industry and the share of loans as a fraction of total loans in the SBA sample. The last two columns report the number of small businesses in each industry and their share as a fraction of total number of small businesses in the U.S. The data for the last two columns are obtained from the 2012 Statistics of U.S. Businesses (SUSB) reported by the Census Bureau. "Public Administration" is a newly added NAICS code not represented in the 2012 SUSB data. "N/A" represents missing industry information. Source: SBA and SUSB.

	SBA sample		Population (SUSB)	
Industry	Number of loans	Share	Number of firms	Share
Accommodation and Food Services	35,797	0.180	495,347	0.086
Retail Trade	31,748	0.160	650,749	0.112
Health Care and Social Assistance	23,995	0.121	640,724	0.111
Other Services (excl. Public Admin)	19,939	0.100	667,176	0.115
Manufacturing	17,173	0.086	256,363	0.044
Professional Services	14,729	0.074	772,685	0.133
Construction	10,636	0.053	640,951	0.111
Wholesale Trade	9,194	0.046	315,031	0.054
Admin Support and Waste Management	6,452	0.032	327,214	0.056
Arts, Entertainment, and Recreation	6,403	0.032	114,969	0.020
Real Estate and Rental and Leasing	5,943	0.030	270,034	0.047
Transportation and Warehousing	4,773	0.024	168,057	0.029
Agriculture	3,836	0.019	21,351	0.004
Finance and Insurance	3,231	0.016	234,841	0.041
Educational Services	2,424	0.012	84,503	0.015
Information	1,879	0.009	71,108	0.012
Mining and Gas Extraction	578	0.003	22,149	0.004
Utilities	135	0.001	5,973	0.001
Management	125	0.001	26,819	0.005
Public Administration	18	0.000	0	0.000
N/A	5	0.000	7,104	0.001

Table C.3: Alternative sources of credit

This table reports the fraction of firms with a loan from a government agency, including the SBA, which have multiple sources of different types of credit in the last 3 years. Source: SSBF.

Outcome	Mean
Multiple lines of credit	0.044
Multiple credit related services	0.044
Multiple equipment loans	0.044
Multiple capital leases	0.000
Multiple other loans	0.067
SBA reason for loan	0.022
Observations	225

Table C.4: Effect of guarantee on loan substitution

This table reports δ and ζ from equation 3.18, which capture the effect of guarantee on loan substitution and differential lending supply response for firms with higher propensity to issue SBA loans. The column headers report the dependent variables. Treat equals one for years 2009 and 2010, when the guarantees were increased and equalized on both sides of the threshold as part of the ARRA stimulus. Exposure is a bank-specific share of small business lending that is through the SBA in 2008, which captures a bank's propensity to specialize in SBA lending prior to the ARRA. The outcomes are log of total small business lending, non-SBA small business lending, and SBA lending. Total and non-SBA small business lending are from FDIC SDI and converted into flows as described in appendix C.3. The sample is restricted to loan size between \$50,000 and \$225,000, and for banks that operated in 2008 such that it has non-missing pre-ARRA exposure. The regressions are estimated at the bank-year level. Standard errors are clustered at the bank-level and reported in parentheses. Source: SBA and FDIC SDI.

	Total loans		Non-SBA loans		SBA loans	
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	-.024	-.110	.104	-.233	.006	.067
	(.068)	(.116)	(.354)	(.563)	(.106)	(.144)
Treat $\times$ Exposure	.029	.031	-2.0	-2.01	.525	.539
	(.222)	(.214)	(1.66)	(1.65)	(.292)	(.299)
Number of observations	1,346	1,346	1,346	1,346	1,346	1,346
Bank FEs	X	X	X	X	X	X
Year FEs		X		X		X

Table C.5: Estimates using alternative range and bin size

This table reports estimates of excess mass and the main elasticity estimates, varying the range used from the restriction in the main sample to loans between \$75,000 to \$225,000, use the step size of 500, include round number dummies for multiples of 1,5, 10, 25, and 50 thousand. The estimates are run using a polynomial of degree 6. The change in the guarantee rate ($\Delta\Gamma$) at the threshold for years in which a notch existed is computed as the weighted average of the average expected guarantee benefit as a percentage of the loan principal, where the weights correspond to the number of loans across years 2008, 2011-2017. Source: SBA.

Year	Range	Excess mass	ΔD	Elasticity	Excess Mass	ΔD	Elasticity	Excess mass	ΔD	Elasticity
		Bin size = 100			Bin size = 200			Bin size = 500		
2008	65,000 - 215,000	4,732	60,400	3.842	4.732	60,600	3.863	4,744	61,000	3.919
2011-2017	65,000 - 225,000	4,737	70,100	5.176	4,736	70,200	5.191	4,747	70,500	5.235
	65,000 - 235,000	4,710	55,500	3.244	4,710	55,200	3.209	4,717	55,500	3.244
	65,000 - 245,000	4,732	90,300	8.589	4,733	91,000	8.722	4,745	91,500	8.818
	75,000 - 225,000	4,744	65,100	4.464	4,736	65,400	4.505	4,710	65,500	4.520
	75,000 - 235,000	4,722	56,400	3.350	4,722	55,600	3.256	4,731	56,000	3.303
	75,000 - 245,000	4725	80,100	6.758	4,725	80,200	6.775	4,736	80,500	6.826
	75,000 - 255,000	4,733	65,100	4.464	5,086	80,200	6.775	5,086	80,500	6.826
	85,000 - 225,000	4,733	65,100	4.464	4,733	65,200	4.477	4,745	65,500	4.519
	85,000 - 235,000	4,727	60,300	3.830	4,726	60,400	3.842	4,735	61,000	3.919
	85,000 - 245,000	4,721	80,100	6.758	4,722	80,200	6.775	4,730	80,500	6.826
	85,000 - 255,000	5,085	80,100	6.758	5,080	80,200	6.775	5,081	80,500	6.826
	85,000 - 265,000	5,085	80,100	6.758	5,077	80,200	6.775	5,080	80,500	6.826

Table C.6: Robustness tests on elasticity estimate parameters

This table reports estimates of excess mass and the main elasticity estimates in each year, varying the polynomial and bin size. The top panel denotes the polynomial used, while the bottom panel denotes the bin size. The change in the guarantee rate ($\Delta\Gamma$) at the threshold for years in which a notch existed is computed as the weighted average of the expected guarantee benefit as a percentage of the loan principal, where the weights correspond to the number of loans across years 2008, 2011-2017. Source: SBA.

	Polynomial degree 5			Polynomial degree 6			Polynomial degree 7		
Year	Excess mass	ΔD	Elasticity	Excess mass	ΔD	Elasticity	Excess mass	ΔD	Elasticity
2008	302.73	67,000	8.83	301.99	53,000	5.48	302.17	53,000	5.48
2009	19.31	3,500	-	19.12	3,500	-	19.16	3,500	-
2010	35.41	7,500	-	35.02	7,000	-	34.94	7,500	-
2011	194.49	46,500	4.56	195.37	43,500	3.98	196.40	46,500	4.56
2012	153.07	66,500	10.20	152.68	59,000	8.00	153.46	58,000	7.72
2013	238.84	62,500	4.76	240.31	72,500	6.43	240.03	63,000	4.83
2014	335.74	57,000	1.62	335.73	62,500	1.96	337.77	73,000	2.68
2015	637.79	61,500	1.96	637.69	56,500	1.65	634.36	54,000	1.51
2016	806.67	71,500	3.23	804.95	62,500	2.45	804.33	72,000	3.27
2017	2021.43	64,500	4.78	2031.94	71,000	5.80	2029.94	71,500	5.89

	Bin size = 100			Bin size = 200			Bin size = 500		
Year	Excess mass	ΔD	Elasticity	Excess mass	ΔD	Elasticity	Excess mass	ΔD	Elasticity
2008	304.17	71,700	10.13	302.21	54,200	5.74	301.99	53,000	5.48
2009	21.22	3,100	-	21.42	3,200	-	19.12	3,500	-
2010	35.18	8,100	-	35.03	9,200	-	35.02	7,000	-
2011	192.56	46,100	4.48	193.87	45,600	4.38	195.37	43,500	3.98
2012	147.65	57,300	7.53	149.66	57,800	7.67	152.68	59,000	8.00
2013	231.96	65,000	5.15	232.39	64,200	5.02	240.31	72,500	6.43
2014	331.94	62,700	1.97	331.32	61,200	1.87	335.73	62,500	1.96
2015	638.19	58,100	1.75	637.65	61,200	1.94	637.69	56,500	1.65
2016	794.90	61,100	2.34	800.34	61,200	2.35	804.95	62,500	2.45
2017	2024.25	69,300	5.52	2024.26	70,200	5.67	2031.94	71,000	5.80

Table C.7: Heterogeneity by location and industry characteristics

This table reports estimates of excess mass and the main elasticity estimates by subsamples of loans by borrower's project location and sectors. Bank competition is measured using the Herfendahl-Hirschman Index (HHI) based on the lender's dollar volume lending share in the borrower's project county. High (Low) bank competition refer to counties that are below the 25th percentile (above the 75th percentile) of the distribution of the HHI measure. High (Low) exit industries refer to industries that are above the median percentile of the distribution of average exit rates by industries. The industry exit rates are obtained from the Census Business Dynamics Statistics. The classification for goods- vs. services-producing industries follows the categorization by the Bureau of Labor Statistics. The classification for tradable vs. non-tradable sectors follows the categorization by Mian and Sufi (2014). The estimation restricts the sample to loans to be of size between \$75,000 to \$225,000, uses the step size of 500, and includes round number dummies for multiples of 1,5, 10, 25, and 50 thousand. The degree of the polynomial used in the estimation is denoted in the second column. The change in the guarantee rate ($\Delta\Gamma$) at the threshold for years in which a notch existed is computed as the weighted average of the average expected guarantee benefit as a percentage of the loan principal, where the weights correspond to the number of loans across years 2008, 2011-2017. Source: SBA and FDIC SDI.

Year	Polynomial	Excess mass	ΔD	$\Delta\Gamma$	Elasticity	Excess mass	ΔD	$\Delta\Gamma$	Elasticity
		High bank competition				Low bank competition			
2008, 2011-2017	6	1,533	66,000	0.038	4.588	907	60,500	0.038	3.856
		(13.67)	(4,736)	–	(0.661)	(20.17)	(7,365)	–	(0.963)
	No. obs.	7,316				9,578			
		Low exit industries				High exit industries			
2008, 2011-2017	6	4,193	66,000	0.038	4.588	553	69,500	0.038	5.088
		(47.90)	(5,822)	–	(0.792)	(8.30)	(5,481)	–	(0.687)
	No. obs.	29,855				3,399			
		Goods-producing industries				Services-producing industries			
2008, 2011-2017	6	3,990	66,000	0.038	4.589	756	68,500	0.038	4.943
		(45.56)	(6,163)	–	(0.832)	(12.20)	(5,736)	–	(0.754)
	No. obs.	28,447				4,807			
		Non-tradable				Tradable			
2008, 2011-2017	6	847	56,000	0.038	3.303	328	66,000	0.038	4.588
		(14.36)	(7,468)	–	(0.950)	(6.73)	(4,159)	–	(0.458)
	No. obs.	7,653				2,177			

Table C.8: Estimate split by borrower's industry

This table reports estimates of excess mass and the main elasticity estimates by the top 5 industries in terms of the share of loans as a fraction of total loans in the SBA sample. The estimation restricts the sample to loans to be of size between \$75,000 to \$225,000, uses the step size of 500, and includes round number dummies for multiples of 1,5, 10, 25, and 50 thousand. The degree of the polynomial used in the estimation is 6. The change n the guarantee rate ($\Delta\Gamma$) at the threshold for years in which a notch existed is computed as the weighted average of the average expected guarantee benefit as the percentage of the loan principal, where the weights correspond to the number of loans across years 2008, 2011-2017. Source: SBA.

Industry	Excess mass	ΔD	$\Delta\Gamma$	Elasticity	No. obs.
Accommodation and Food Services	743.052	72,500	0.038	5.537	5,862
Retail Trade	687.019	55,500	0.038	3.245	5,256
Health Care and Social Assistance	360.836	62,500	0.038	4.115	3,354
Other Services (excl. Public Admin)	596.223	70,500	0.038	5.235	4,243
Manufacturing	333.498	67,000	0.038	4.729	2,398

Table C.9: Components of Main Elasticity Estimates

This table lists the main outputs of the bunching estimation routine for each year. For this estimation: Step size = 500, the range was limited to 75,000-225,000, we included round number dummies for multiples of 1,5, 10, 25, and 50 thousand, and we used a polynomial of degree 6. We excluded years 2009 and 2010 when there was no change in the guarantee. D_L refers to the lower bound of the excluded region, D^* is the threshold, D_U is the estimated upper bound of the excluded region, ΔD is the size of the excluded region, B is the excess number of loans estimated at the threshold, and M is the estimated number of missing loans in the excluded region.

Year	D_L	$D^\star$	D_U	ΔD	$\hat{B}$	$\hat{M}$	Step size
2008	149,000	150,000	201,500	52,500	248.39	-335.98	500
2011	149,000	150,000	190,500	41,500	151.81	-190.00	500
2012	149,000	150,000	210,500	61,500	132.64	-167.35	500
2013	149,000	150,000	221,500	72,500	199.91	-366.70	500
2014	149,000	150,000	212,000	63,000	233.02	-269.15	500
2015	149,000	150,000	205,500	56,500	457.83	-516.82	500
2016	149,000	150,000	210,500	61,500	564.04	-562.26	500
2017	149,000	150,000	219,500	70,500	1386.12	-1462.46	500

Bibliography

Kleven, Henrik J., and Mazhar Waseem. 2013. "Using Notches to Uncover Optimization Frictions and Structural Elasticities: theory and Evidence from Pakistan." *Quarterly Journal of Economics*, 128(2): 669–723.

Mian, Atif, and Amir Sufi. 2014. "What Explains the 2007-2009 Drop in Employment?" *Econometrica*, 82(6): 2197–2223.