Title
Three Essays in Applied Microeconomics

Author
Papich, Sarah

Publication Date
2024

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Three Essays in Applied Microeconomics

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Economics

by

Sarah Papich

Committee in charge:

Professor Richard Startz, Chair
Professor Heather Royer
Professor Kelly Bedard

June 2024

The Dissertation of Sarah Papich is approved.

Professor Heather Royer

Professor Kelly Bedard

Professor Richard Startz, Committee Chair

April 2024

Three Essays in Applied Microeconomics

by

Sarah Papich

"Economics is what economists do." -Jacob Viner

Acknowledgements

I am very grateful for the support and guidance I have received from many amazing people as I worked on this dissertation. First, I would like to thank my committee: Dick Startz, Heather Royer, and Kelly Bedard. I feel very lucky to have had the opportunity to learn from these three excellent economists. The time, support, and guidance you have given me throughout the Ph.D. program has been invaluable. I especially want to thank Dick, my committee chair. Thank you for being so generous with your time, for encouraging me and pushing me to be better, and for responding to my emails so quickly. I am so grateful that you agreed to be my chair.

Next, I am grateful for my colleagues and friends at UCSB. Allegra and Michael, thank you for listening to so many version of my job market presentation, giving feedback on my paper, and providing so much moral support in the past year. I would also like to thank my friends Kite, Lei, Toshio, Roberto, Jaime, Micah, and Kaili for supporting me throughout this process and making the past five years a lot more fun.

I would also like to thank Martin Boileau and Francisca Antman for encouraging me to pursue a Ph.D. when I was an undergraduate student and for helping me apply. Your mentorship and support was truly life-changing, and I would not be where I am today without you.

Finally, I would like to thank my friends outside UCSB, my family, and Danny. Danny, thank you for always believing in me and for your endless love and support throughout this very difficult process. I could not have gotten here without you.

Curriculum Vitæ
Sarah Papich

Education

2024	Ph.D. in Economics (Expected), University of California, Santa Barbara.
2020	M.A. in Economics, University of California, Santa Barbara.
2019	B.A. in Economics (*summa cum laude*) and Political Science, University of Colorado Boulder.

Publications

Forthcoming	"Marijuana Legalization and Fertility." *American Journal of Health Economics*
January 2024	"Do Democracy Vouchers Help Democracy?" *Contemporary Economic Policy*

Working Papers

Effects of Buy Now, Pay Later on Financial Well-Being

Work Experience

April 2022-Present	Senior Research Analyst, UCSB Economic Forecast Project
January-December 2021	Research Assistant, Professor Alisa Tazhitdinova
September-December 2021	Head Teaching Assistant, Econometrics
September 2020-July 2021	Teaching Assistant, Econometrics

Fellowships and Awards

University of California, Santa Barbara

2019-2020	Regents Fellowship
2019-2024	Mortimer Andron Fellowship
Winter 2022	Research Quarter Fellowship
Winter 2022	Best Second Year Paper Award
May 2022	California Policy Lab Seed Grant
Summer 2022	California Policy Lab Graduate Fellowship

University of Colorado Boulder

2015-2019	GPA: 4.0
2015-2019	Joseph A. Sewall Award
2015-2019	Michael A. Hoza Award
2018	Center for Western Civilization, Thought and Policy Scholarship
2018	Sieglinde Talbott Haller Endowed Economics Scholarship
2018	Christopher S. Randall Scholarship
2019	Chancellor's Recognition Award
2019	Val B. and Helen W. Fischer Scholarship for Academic Superiority in Social Sciences

Presentations

2021	Southern California Graduate Conference in Applied Economics
2022	European Economic Association Conference
2022	Pacific Premier Bank Regional Meeting
2023	Western Economic Association International Conference

Publications Outside Economics

2017	The Impacts of Outreach Efforts by Research Computing at the University of Colorado Boulder (joint with Shelley Knuth, Peter Ruprecht and Thomas Hauser) *Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact*

Other Activities

2023-2024	Visiting Student, University of California San Diego
2021-Present	Newsletter Contributor, Laboratory for Aggregate Economics and Finance
2021-2022	UCSB Economics Mentor Program Leader
2020	Organizing Committee Member, Southern California Graduate Conference in Applied Economics

Abstract

Three Essays in Applied Microeconomics

by

Sarah Papich

This dissertation consists of three essays in applied microeconomics. While the topics vary, the three papers are united in their use of causal inference techniques and their relevance to policy: each paper either evaluates effects of an existing policy or examines whether new policies are needed for consumer protection.

The first essay examines the effects of access to Buy Now, Pay Later (BNPL) on financial well-being. Many American consumers have limited access to credit, raising the question of whether an increase in credit access would make them better off. Fully rational individuals would use an increase in credit access to smooth consumption, yet real consumers may make financial mistakes by accumulating debts they cannot repay. I study the effects of making BNPL accessible to American consumers, including those who otherwise have limited access to credit. This paper provides the first causal evidence of how access to BNPL affects severe measures of financial distress and credit scores. Using credit bureau data and a two-way fixed effects identification strategy that exploits geographic and temporal variation in availability of BNPL at a large retailer, I find that access to BNPL reduces financial distress arising from late or missed debt payments. The total amount past due decreases by 2.4% and the number of current delinquencies decreases by 0.2%. Heterogeneity analysis reveals that these effects are strongest among consumers with "fair" credit scores, the second-lowest credit score category. I also find that BNPL access increases credit scores by an average of 1.6 points and increases use of non-BNPL credit. These results suggest that access to BNPL reduces financial distress

rather than causing consumers to accumulate unsustainable debts.

The second essay studies how public financing for political campaigns affects political participation and campaign contributions. Seattle's Democracy Vouchers program provides a unique form of public financing for political campaigns in which voters decide how to allocate public funding across candidates. This paper is the first to study the effects of public financing for political campaigns on political participation. I estimate that the Democracy Vouchers program increases voter turnout by 4.9 percentage points, suggesting that public financing programs can increase political participation. I also find that campaigns become more reliant on small contributions. For city council candidates, dollars from small contributions under $100 increase by 156% while dollars from large contributions over $250 decrease by 93%.

The third essay examines how legalizing marijuana affects fertility. State-level marijuana legalization has unintended consequences, including its effect on fertility. Marijuana use is associated with behaviors that increase fertility as well as physical changes that lower fertility. In this paper, I provide the first causal evidence of the effects of recreational marijuana legalization on birth rates using a difference-in-differences design that exploits variation in marijuana legalization across states and over time. The main result is that legalizing recreational marijuana decreases a state's birth rate by an average of 2.78%. Heterogeneity analysis shows that the largest decrease in the birth rate occurs among women close to the end of their child-bearing years. I find suggestive evidence of increases in days of marijuana use per month and in the probability of being sexually active. Together, these findings show that the physical effects of marijuana use have the dominant effect on fertility. Finally, I examine the effects of medical marijuana legalization on fertility and find a smaller, statistically insignificant decrease in the birth rate, which is consistent with the smaller increase in marijuana use that results from medical legalization.

Contents

Chapter 1

Effects of Buy Now, Pay Later on Financial Well-Being

1.1 Introduction

Do consumers benefit from an increase in access to credit? Of American adults who applied for credit in 2022, 30% were denied or approved for less credit than they requested [1]. Standard economic theory dictates that these credit-constrained consumers would be better off if they could borrow money to spread out large expenses and weather income shocks. Yet lenders are often unwilling to extend them credit out of fear of moral hazard: rather than simply using loans to smooth consumption, these particular consumers may fail to repay their debts. If credit access were increased and consumers used it to accumulate more debt than they could repay, they could be left worse off than if they had not been given access to credit. Determining the conditions under which these credit-constrained consumers can benefit from credit access is key to improving their financial well-being.

In this paper, I study how increasing credit access affects financial well-being in the context of "Buy Now, Pay Later" (BNPL), a new type of small loan that allows consumers to split a purchase from a retailer into installments. BNPL loans charge zero or very low interest and are accessible to many consumers, including those with limited access to traditional forms of credit. These loans have become very popular: a 2021 survey found that 44% of Americans have used BNPL [2], and $24 billion worth of loans were issued in the United States by the five largest BNPL providers in 2021 alone [3].

This paper provides the first causal evidence of how access to BNPL affects credit scores and severe financial distress arising from late or missed debt payments, such as delinquency and bankruptcy. I introduce a model to characterize borrowing and repayment decisions in a world with both BNPL access and traditional lenders. For empirical analysis, I use variation in BNPL availability at Walmart to identify the causal effect of BNPL access on financial well-being. Walmart is the largest retailer in the United States, with $430.8 billion in sales in 2020 [4], and was an early adopter of BNPL. Walmart made BNPL available in its stores in February 2019 through a partnership with the BNPL provider Affirm. This partnership increased BNPL access more for consumers living close to Walmart stores than for those living far away, creating geographic and temporal variation in BNPL access. I use a two-way fixed effects identification strategy to measure the causal effect of increasing BNPL access for consumers living near Walmart when the partnership went into effect.

I use longitudinal data from the University of California Consumer Credit Panel (UC-CCP), which is provided by one of the three major credit bureaus in the United States and includes 9.9 million observations from June 2014 through March 2022. I use three outcomes to measure the effects of BNPL access on debt-related financial distress: total amount past due; number of current delinquencies; and the number of debt-related court proceedings, which include bankruptcy filings. BNPL loans are not included in the data because they were not reported to credit bureaus during the study period. Therefore, the results show how access to BNPL affects financial distress arising from non-BNPL loans that are reported to credit bureaus. Experiencing financial distress on traditional loans is very costly for consumers, as these financial distress measures are reported to credit bureaus and can lead to long-term difficulties accessing future credit. I also examine how BNPL affects credit scores and use of non-BNPL credit.

The main finding is that BNPL access reduces financial distress. I estimate that the total amount past due decreases by 2.4% and the number of current delinquencies decreases by 0.2%. The estimated effect on debt-related court proceedings is negative, statistically insignificant and close to zero. Heterogeneity analysis by pre-treatment credit score category reveals that evidence of decreased financial distress is strongest among consumers with "fair" credit scores, the second-worst category. Consumers in the worst category of "poor" credit scores are less affected, likely because many of them do not qualify for BNPL. Consumers with the highest credit scores, in the "exceptional" range, are also unaffected. These consumers likely qualify for other forms of credit with very low interest rates and have little to gain from using BNPL, while consumers with "fair" credit typically have low credit limits, pay high interest rates on traditional forms of credit, and qualify for BNPL. Consumers with "fair" credit scores therefore have the most to gain from using BNPL. Although traditional lenders consider consumers with "fair" credit scores to be highly at risk of failing to repay loans, these consumers benefit the most from BNPL access in terms of reduced financial distress.

I find that BNPL access improves credit scores by 1.6 points, likely due to decreased financial distress. This effect is positive but economically small, as credit scores range from 300 to 850. I also estimate that BNPL access increases consumers' use of traditional credit: total open balances increase by 4.3%, and I find suggestive evidence of a 2.1% increase in the number of open credit cards. Increases in use of traditional credit could result from consumers repaying BNPL loans with credit cards and from consumers gaining better access to traditional credit as financial distress decreases and credit scores improve.

I use two placebo tests to assess whether financial distress would have decreased among consumers living near Walmart even if BNPL had not been introduced. The first placebo test repeats the main analysis using only data from Iowa and West Virginia, because these were the only two states in which Walmart did not introduce BNPL in February 2019. The second placebo test repeats the main analysis with the treated group defined as individuals living near Target instead of near Walmart, as BNPL was not available at Target when it became available at Walmart. Both placebo tests provide evidence that the main results capture the effect of increased access to BNPL, rather than a trend over time among individuals living near Walmart that would have happened in the absence of BNPL.

The results are robust to the use of different distance cutoffs to define a consumer as living near Walmart. A second robustness check repeats the analysis with a propensity-score trimmed sample, because Walmart stores are not located randomly and areas near Walmart stores are not demographically identical to those that are farther away. A third robustness check only uses data from before the COVID-19 pandemic to address the concern that changes in consumer behavior during the pandemic could limit external validity. The results of all three robustness checks are similar to the main results. Consumer surveys suggest that individuals under 40 are most likely to use BNPL [5, 6]. I repeat my analysis with a sample of only individuals under age 40 and find that all estimates have the same direction as the main results, but five out of six are larger in magnitude in the restricted sample. Stronger effects among the individuals who are most likely to use BNPL support the assumption that access to BNPL drives the main results.

Overall, I find that BNPL access improves financial well-being by reducing financial distress. While this finding is consistent with descriptive evidence that BNPL users do not have higher delinquency rates than the general population [7], they contrast with news coverage of BNPL, which has largely focused on anecdotal evidence that it can be used to accumulate debts that consumers struggle to repay [8, 9, 10, 11]. The improvement in repayment outcomes likely results from an increased ability to smooth consumption, especially for consumers without access to other low-interest forms of credit. When these consumers experience a negative income shock or make an unusually large purchase, BNPL access can enable them to avoid missing payments on traditional loans. These results are an important consideration for policymakers as they decide whether to increase regulation of BNPL to protect consumers. BNPL is an area of interest for policymakers: the Consumer Financial Protection Bureau opened an inquiry into BNPL in December 2021, and the House of Representatives held a hearing to gather information about BNPL in November 2021.

This paper contributes to a broader literature on the effects of access to credit, including payday loans and credit cards. The payday lending literature provides mixed evidence on whether access to high-interest payday loans benefits [12, 13, 14, 15, 16, 17, 18, 19, 20, 21] on consumers. The most closely related paper in this literature is [22], which uses an individual's distance from the nearest payday lender to estimate effects on financial hardship and finds that payday loan access increases the probability of struggling to pay bills. I use a similar identification strategy to estimate the effects of an increase in access to BNPL. BNPL shares two important characteristics with payday loans: both are for small amounts and are accessible to consumers with limited access to other forms of credit, but not to the very poorest consumers.

The literature on credit card borrowing provides evidence that present bias can lead to welfare-reducing over-borrowing by unsophisticated consumers [23, 24] and can cause many consumers to deviate from their planned repayment schedules [25]. Present bias also has the potential to lead BNPL users to over-borrow. However, BNPL and credit cards have two key differences: BNPL loans have lower interest rates and BNPL sets a strict, short-term payment schedule, while consumers have discretion over how much time to spend paying back credit card debt as long as they meet the minimum payment [26]. [27] find that the minimum payment on a credit card bill serves as an "anchor", exerting outsize influence on consumers' repayment choices.

This paper also makes three contributions to the emerging economic literature on BNPL. First, I use data from a credit bureau to provide a comprehensive picture of how BNPL access affects use and repayment of traditional forms of credit. Previous work on BNPL has used transaction-level data on spending [28, 29] or descriptive evidence from credit cards only [30]. My finding that use of traditional credit increases is consistent with evidence in the literature that BNPL payments are often made by credit card [30] and that BNPL access increases retail spending [28]. Second, I provide the first causal evidence of how BNPL access affects severe measures of financial distress and credit scores, and I find improvements that contrast with the previous literature. Prior work has found detrimental effects of BNPL on mild measures of financial distress, including overdraft fees [28, 29], low balance fees [28], credit card interest payments and late fees [29]. My results answer the question raised in previous work [30, 28, 29] of whether BNPL access leads to a "debt spiral" that could culminate in bankruptcy: I directly examine effects on delinquencies and bankruptcy and find no evidence of this phenomenon.

1.2 Background

1.2.1 BNPL

Consumer surveys provide information about which individuals are most likely to use BNPL. Consumers with low credit scores or short credit histories are especially likely to use BNPL, due to difficulties qualifying for other low-interest forms of credit [5]. BNPL is most popular with adults under 40 [5, 6]. Individuals who have recently had a baby, fallen ill or lost a job are more likely to use BNPL [31]. 89% of BNPL users have at least one credit card, and 73% have an auto loan [7]. BNPL may be appealing to consumers who mistrust mainstream financial institutions, as well as workers with irregular incomes such as gig economy workers [32]. BNPL is typically used for small purchases, with the average transaction costing $200 [33]. 74% of BNPL users make their payments on time [26], an encouraging sign that most consumers are not spending more than they can afford to repay.

The low cost, ease of access and consumption smoothing features of BNPL make it attractive to consumers. A survey from TransUnion found that the two most common reasons for using BNPL are the ability to spread out payments over time and the easy application process [7]. Many consumers who would not qualify for a low-interest credit card can qualify for BNPL. Credit cards charge an average interest rate of 16.45%, and rates can be as high as 30% for individuals with lower credit scores [34]. Payday loans charge an average of 391% interest on small, short-term loans that do not require a credit check [35] and are frequently utilized by credit-constrained individuals. BNPL access could also help consumers by enabling consumption smoothing: using BNPL to spread out the cost of a large purchase could reduce month-to-month volatility in a consumer's bills. Biweekly BNPL installments often align with the timing of paychecks [36], which could help consumers avoid late or missed payments.

Offering BNPL has clear advantages for retailers. BNPL increases the average retail ticket size by 30-50% and helps retailers attract new customers [33]. The minimum purchase amount that some retailers require for BNPL could be a factor driving the increased ticket size. The BNPL provider assumes the risk that the consumer will not repay the loan, allowing the retailer to benefit from the higher sales that BNPL generates without the risk of offering loans itself. Retailers have offered other payment-over-time options in the past, including layaway and zero-interest repayment for large purchases when a consumer opens a store credit card. However, BNPL differs substantially from these two alternatives. First, BNPL tends to be used more frequently and for smaller purchases than layaway or store credit cards [36]. Second, layaway requires the customer to pay in full before receiving the item, while BNPL allows the customer to receive the item now and pay later [33]. Third, repayment plans that involve opening a store credit card are only available to customers who would qualify for the credit card, and BNPL is available to many customers who would not qualify.

1.2.2 BNPL at Walmart with Affirm

Many retailers that partner with a BNPL provider specialize in one type of good, such as makeup or home decor. The BNPL partnership between Walmart and Affirm is unique because Walmart sells a wide variety of items and has more customers than any other American retailer. Walmart was also a relatively early adopter of BNPL: only 0.48% of Americans had used BNPL in 2018, the last year before the Walmart-Affirm partnership formed [37]. Walmart's largest competitors, Target and Amazon, did not offer BNPL until 2021. All Walmart stores in the United States made Affirm available in February 2019 except stores in Iowa and West Virginia.

The partnership made BNPL available in Walmart stores for purchases of at least $144. Walmart made Affirm availability salient by advertising it to in-store shoppers with co-branded signs hung from the ceiling and placed in aisles; tags placed on individual items for sale; takeovers of the wall of television screens in the electronics section; and advertisements mailed directly to the homes of Walmart customers. Figure 1.1 shows an example of a sign placed in a Walmart store, Figure 1.2 shows an example of a tag placed on an individual item for sale at Walmart, and Figure 1.3 shows an example of an image

Figure 1.1: Sign Advertising Affirm

shown in a television wall takeover.

Figure 1.2: Tag Advertising Affirm

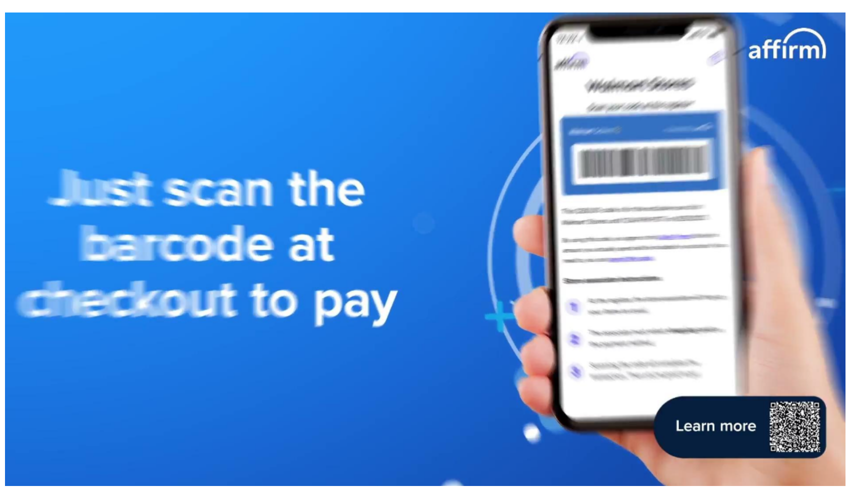

Figure 1.3: Affirm TV Wall Takeover

Walmart pays a transaction fee every time a customer uses Affirm. BNPL transaction fees can be up to three times the typical fee retailers pay to accept credit cards [9]. Affirm does not charge any late fees, likely increasing its reliance on fees from retailers. BNPL providers that do charge late fees typically earn 25% of their revenue from late fees and 75% from retailer fees [33].

Affirm performs a soft credit check when a consumer applies for a loan [38]. A soft credit check provides information about a consumer's credit score, lines of credit and payment history, but it does not appear on credit reports or affect the consumer's credit score. The consumer's payment history with Affirm is also considered. If the consumer is approved to receive a loan from Affirm, installments are due biweekly or monthly. The shortest-term loans consist of four biweekly installments and charge zero interest. As the price of an item increases, Affirm becomes more likely to offer a longer-term payment plan that may charge interest [38]. Affirm does not provide information about the exact formula they use to determine loan eligibility or how they decide which repayment plans to offer to each customer. Affirm has reported that descriptively, average order values at its partner retailers are 85% higher among customers who use its BNPL service [33].

The Affirm website provides details on how partnering with a retailer makes BNPL visible and easy to use during in-store checkout. The partnership allows the retailer to integrate the BNPL option directly into their in-store checkout system. The cashier can select "Affirm" as the payment option from the retailer's point of sale system, which triggers a request to Affirm to send a checkout link to the customer directly via email or text. Upon receiving the link, the consumer fills out a short application on her smart phone and is then able to pay with Affirm if she is approved.

Affirm recommends specific actions its partner retailers can take to increase the salience of BNPL for in-store shoppers. In addition to the co-branded advertising discussed above, Affirm recommends training sales associates to tell customers about BNPL, answer questions about it, and assist with any problems that arise in the application process. Affirm's website says that a common pathway to a customer using Affirm in-store is when "A customer learns about Affirm from a sales associate." Affirm provides an employee training program to its partner retailers that teaches their sales associates how to increase use of Affirm in-store.

[39] provides descriptive data on Walmart customers in 2019. Walmart has more customers aged 25-34 than in any other age group; over 20% of Walmart customers are in this category. More women than men shop at Walmart. The average Walmart shopper has an annual household income of $76,000, above the U.S. average of $67,521.

1.3 Model

I consider optimal consumption, lending and repayment decisions for an individual borrower in a world with BNPL. I model the borrower's decision, taking lenders' behavior as given. First, I describe relevant features of the lenders. Second, I introduce the borrower's problem and describe characteristics of her optimal consumption, lending and repayment choices. Although preferences and choices are specific to the individual, I omit i subscripts to simplify notation.

1.3.1 Lenders

There are two lenders: a "traditional" (non-BNPL) lender and a BNPL lender. The traditional lender offers a loan to individual i of up to the traditional credit limit B^T at traditional interest rate r^T . The BNPL lender offers individual i a loan of up to B^B at interest rate r^B where $r^B \leq r^T$ and r^B may be zero. $B^B \leq \bar{B}^B$, where $\bar{B}^B$ is the maximum possible BNPL credit limit. For Affirm, for example, $\bar{B}^B$ is $17,500. The maximum possible traditional credit limit is $\bar{B}^T$ where $\bar{B}^T > \bar{B}^B$. An individual has a credit score S, where a higher score indicates that the individual is more creditworthy based on past borrowing and repayment behavior. Both traditional and BNPL borrowing limits are higher for individuals with higher credit scores, so that $\frac{\partial B^T}{\partial S} \geq 0$ and $\frac{\partial B^B}{\partial S} \geq 0$. The traditional lender charges a lower interest rate as the credit score increases: $\frac{\partial r^T}{\partial S} \leq 0$. If $r^B > 0$, the BNPL lender also lowers the interest rate as the credit score increases: $\frac{\partial r^B}{\partial S} \leq 0$. The traditional lender will only lend to individuals with $S > \underline{S}_T$ and the BNPL provider will lend to individuals with $S > \underline{S}_B$, where $\underline{S}_B < \underline{S}_T$. Below $\underline{S}_T$, $B^T = 0$ and below $\underline{S}_B$, $B^B = 0$.

Traditional and BNPL lenders use different business models, which lead to different interest rates and minimum credit scores for the two types of lenders. Traditional lenders earn money from interest payments on their loans, while BNPL lenders primarily earn money from retailer fees and are therefore able to charge lower interest rates. Therefore $r^B \leq r^T$. The difference in how the two types of lenders earn money provides one explanation for why $\underline{S}_B < \underline{S}_T$. The presence of retailer fees, which the BNPL lender receives regardless of whether the BNPL loan is repaid, limits expected net losses for a BNPL lender and makes the BNPL lender more willing to lend to "riskier" individuals with lower S . Affirm does not charge any late fees. Therefore their reason for lending to individuals with lower S is not that they profit from late fees when a "riskier" individual fails to repay a loan.

1.3.2 Borrowers

In each period t , individual i has traditional debt D_t^T and BNPL debt D_t^B. She chooses consumption C_t , traditional borrowing L_t^T , repayment of traditional loans R_t^T , BNPL borrowing L_t^B and repayment of BNPL loans R_t^B . She can fund consumption with borrowing and income Y_t . The individual receives an income shock θ_t in each period that determines whether she receives Y_t^l or Y_t^h , where $Y_t^l < Y_t^h$:

$$Y_t = \begin{cases} Y_t^l & \text{if } \theta_t = 0 \\ Y_t^h & \text{if } \theta_t = 1 \end{cases}$$

Traditional lending and borrowing affect the individual's credit score in the next period: $\frac{\partial S_{t+1}}{\partial L_t^T} \leq 0$ and $\frac{\partial S_{t+1}}{\partial R_t^T} \geq 0$. BNPL lending and repayment do not affect credit scores because BNPL is not reported to credit bureaus. However, the BNPL lender will extend more credit if the individual has a history of repaying BNPL loans. Therefore the individual must consider the following relationships between this period's lending and repayment decisions and next period's credit limits: $\frac{\partial B_{t+1}^T}{\partial R_t^T} \geq 0$, $\frac{\partial B_{t+1}^T}{\partial L_t^T} \leq 0$, $\frac{\partial B_{t+1}^B}{\partial R_t^T} \geq 0$, $\frac{\partial B_{t+1}^B}{\partial L_t^T} \leq 0$, $\frac{\partial B_{t+1}^B}{\partial L_t^B} \leq 0$, and $\frac{\partial B_{t+1}^B}{\partial R_t^B} \geq 0$. While the individual knows the direction of these relationships, the exact formula used to calculate each period's credit limit is proprietary lender information and is therefore unknown to the individual.

Although $r^B \leq r^T$, an individual may choose to borrow from the traditional lender even when she has not reached her BNPL credit limit, so that $L_t^B < B_t^B$ and $L_t^T > 0$. One reason for this behavior could be that the individual would rather borrow and then repay traditional credit in order to improve her credit score, which will increase both her traditional and BNPL credit limits. While an increase in L_t^T decreases S , this decrease could be outweighed in the long term by the increase in S that results from paying off the loan. Another reason could be that BNPL is not available at all retailers, and therefore her options are limited when she funds consumption through BNPL. Credit bureau evidence has empirically shown that BNPL users also utilize traditional credit [7].

The individual's goal in a world with BNPL is to maximize lifetime utility given by:

$$u(C_t) + \beta E_t[\Sigma_{\tau=1}^{T-t} \delta^{\tau} u(C_{t+\tau})] \text{ where } \frac{du}{dC_t} \geq 0 \text{ and } \frac{d^2u}{dC_t^2} \leq 0 \tag{1.1}$$

subject to:

$$C_t \leq Y_t + L_t^T + L_t^B - R_t^T - R_t^B \tag{1.2}$$

$$L_t^T \leq B_t^T \tag{1.3}$$

$$L_t^B \leq B_t^B \tag{1.4}$$

$$R_t^T \leq D_t^T \tag{1.5}$$

$$R_t^B \leq D_t^B \tag{1.6}$$

$$B_{t+1}^T = f(L_t^T, R_t^T, B_t^T) \tag{1.7}$$

$$B_{t+1}^B = g(L_t^T, R_t^T, L_t^B, R_t^B, B_t^B) \tag{1.8}$$

$$D_{t+1}^T = (D_t^T - R_t^T + L_t^T)(1 + r_t^T) \tag{1.9}$$

$$D_{t+1}^B = (D_t^B - R_t^B + L_t^B)(1 + r_t^B) \tag{1.10}$$

The individual is present-biased and discounts the future at rate δ . Preferences are time-inconsistent if $\beta \neq 1$, as established in [40]. β and δ determine the extent to which the individual values future utility. An individual who values the future highly has an incentive to borrow less and repay more of her debt today to obtain a higher credit limit in the future, while the reverse is true for an individual who does not value the future. A higher future credit limit enables higher future consumption. With time-inconsistent preferences, an increased opportunity to borrow could lower the individual's lifetime utility.

Who has the most to gain from access to BNPL? One benefit of BNPL is that it increases the total credit limit, $B_t^T + B_t^B$. The increased credit limit from BNPL allows the individual to increase consumption, which will have the largest utility gains for individuals with low values of Y_t . Individuals with low incomes and credit scores close to $\underline{S}_B$ benefit most from the increased credit limit, as these individuals have low values of B_t^T . A second benefit is that BNPL offers a lower interest rate than traditional loans. As S_t increases, $r_t^T \to 0$ and the benefit of $r_t^B \leq r_t^T$ shrinks towards zero. Therefore as S_t decreases towards $\underline{S}_B$ the individual has more to gain from the low interest rate offered by BNPL. A third benefit is that BNPL enables consumption smoothing. Increased ability to smooth consumption is particularly valuable for individuals who experience a negative income shock in period t , so that $Y_t = Y_t^l$, especially if they have low values of S_t and therefore B_t^T is low and r_t^T is high.

The increase in access to BNPL that I study in this paper is an increase in B_t^B . The increase in B_t^B enables the individual to increase overall consumption, smooth consumption, and increase the amount she borrows at the lower interest rate r_t^B . I examine the effects of the increase in B_t^B on use and repayment of traditional forms of credit, and therefore I am interested in the signs of $\frac{\partial R_t^{T*}}{\partial B_t^B}$ and $\frac{\partial L_t^{T*}}{\partial B_t^B}$.

First, I consider the potential direction of $\frac{\partial L_t^{T*}}{\partial B_t^B}$. BNPL and traditional forms of credit are complements if $\frac{\partial L_t^{T*}}{\partial B_t^B} > 0$, and substitutes if $\frac{\partial L_t^{T*}}{\partial B_t^B} < 0$. Because $r_t^B \leq r_t^T$, if C_t^* is unchanged by the increase in B_t^B the individual will substitute BNPL for traditional credit without changing overall borrowing so that $\frac{\partial L_t^{T*}}{\partial B_t^B} < 0$. However, if $\frac{\partial C_t^*}{\partial B_t^B} > 0$ the individual may increase traditional borrowing in addition to increasing BNPL borrowing so that $\frac{\partial L_t^{T*}}{\partial B_t^B} > 0$. One reason for $\frac{\partial C_t^*}{\partial B_t^B} > 0$ could be that $Y_t = Y_t^l$ and i prefers to smooth consumption, but the extent of consumption smoothing is limited by the amount of available credit. $\frac{\partial C_t^*}{\partial B_t^B} > 0$ is more likely at low values of Y_t because lower-income individuals gain more utility from an increase in consumption.

Second, I consider the potential direction of $\frac{\partial R_t^{T*}}{\partial B_t^B}$. Repaying traditional loans in period t reaps future benefits in terms of higher B_{t+1}^B and B_{t+1}^T but lowers C_t . The individual prefers lower R_t^T if the utility gain from increasing C_t exceeds the gain from increasing future credit limits. When $Y_t = Y_t^l$, $\frac{\partial R_t^{T*}}{\partial B_t^B}$ could be positive: an individual faced with a negative income shock is more likely to pay her debts when BNPL is available to help her smooth consumption. $\frac{\partial R_t^{T*}}{\partial B_t^B} < 0$ could occur if $\frac{\partial L_t^{B*}}{\partial B_t^B}$ is large and positive, leading to a value of D_t^B that is large relative to Y_t . This outcome is more likely at low values of β and δ and could lead to lower repayment of traditional debt.

What do the signs of $\frac{\partial R_t^{T*}}{\partial B_t^B}$ and $\frac{\partial L_t^{T*}}{\partial B_t^B}$ tell us about financial well-being? $\frac{\partial L_t^{T*}}{\partial B_t^B} > 0$ accompanied by $\frac{\partial R_t^{T*}}{\partial B_t^B} \geq 0$ could signify improved financial well-being. Higher borrowing can indicate higher consumption and consumption smoothing, both of which increase utility. With $\frac{\partial R_t^{T*}}{\partial B_t^B} \geq 0$, the increase in utility from consumption does not come at the cost of a failure to repay debt. $\frac{\partial L_t^{T*}}{\partial B_t^B} > 0$ accompanied by $\frac{\partial R_t^{T*}}{\partial B_t^B} < 0$ could be a cause for concern. This combination would indicate that individuals with access to BNPL increase consumption by accumulating debts that they do not repay. Failure to repay traditional debts today leads to higher long-term repayment costs as the debt accumulates interest at rate r_t, as well as lower credit limits and higher interest rates in the future. In the remainder of the paper, I empirically examine the signs of $\frac{\partial R_t^{T*}}{\partial B_t^B}$ and $\frac{\partial L_t^{T*}}{\partial B_t^B}$ by estimating the causal effect of an increase in access to BNPL on use and repayment of non-BNPL forms of credit.

1.4 Data

1.4.1 Walmart Location and BNPL Availability Data

I obtained the date when Affirm became available at Walmart from an announcement on Affirm's Twitter page.[1] Data on the locations of all Walmart stores in the United States is available on the Walmart website. I used shapefiles from the U.S. Census to determine the distance from the center of each census block to the nearest Walmart. My data on the distance from the center of each census block to the nearest Walmart shows that 9.13% of census blocks in the United States are within two miles of Walmart, and 29.48% are within five miles.

1.4.2 Outcome Data

For data on financial well-being, I use the California Policy Lab's University of California Consumer Credit Panel (UC-CCP). This longitudinal panel provides quarterly data on a 1% random sample of California residents from 2004 until the first quarter of 2022. The data includes each consumer's demographic attributes, census block and ZIP code, as well as detailed financial information on the amount consumers borrow; the amount they repay; and consequences of failing to repay, such as accounts going into delinquency or a consumer declaring bankruptcy. BNPL loans are not included in the data, and therefore I measure the effects of BNPL access on use and repayment of forms of credit that are reported to credit bureaus.

The data includes individual-level demographic characteristics, such as gender, education and age. Certain characteristics are not available, including race and income. BNPL is a significant source of consumer lending in California: a report from the Califor-

[1] `https://web.archive.org/web/20230118161951/https://twitter.com/affirm/status/1100737321345794048?lang=en`

nia Department of Financial Protection and Innovation found that 91% of all consumer loans in California in 2020 were from the top six BNPL providers [41].

I drop observations prior to June 2014 because 2010 census block identifiers first appear in the data in June 2014. Beginning the sample in June 2014 also excludes the Great Recession from the pre-treatment period. I also exclude deceased individuals and individuals for whom geographic identifiers are missing, leaving a sample of 9,929,100 observations of 310,284 individuals. Access to BNPL is measured as the distance from the center of the census block to the nearest Walmart.

Figure 1.4 and Figure 1.5 show trends over time in the raw data for all seven outcomes. These graphs provide descriptive evidence that outcomes for consumers living near and far from Walmart follow similar trends over time, suggesting that consumers who do not live close to a Walmart store provide a useful counterfactual for those who do. I use event studies, described in more detail in the Methodology section, to formally test whether the two groups followed parallel trends before BNPL became available.

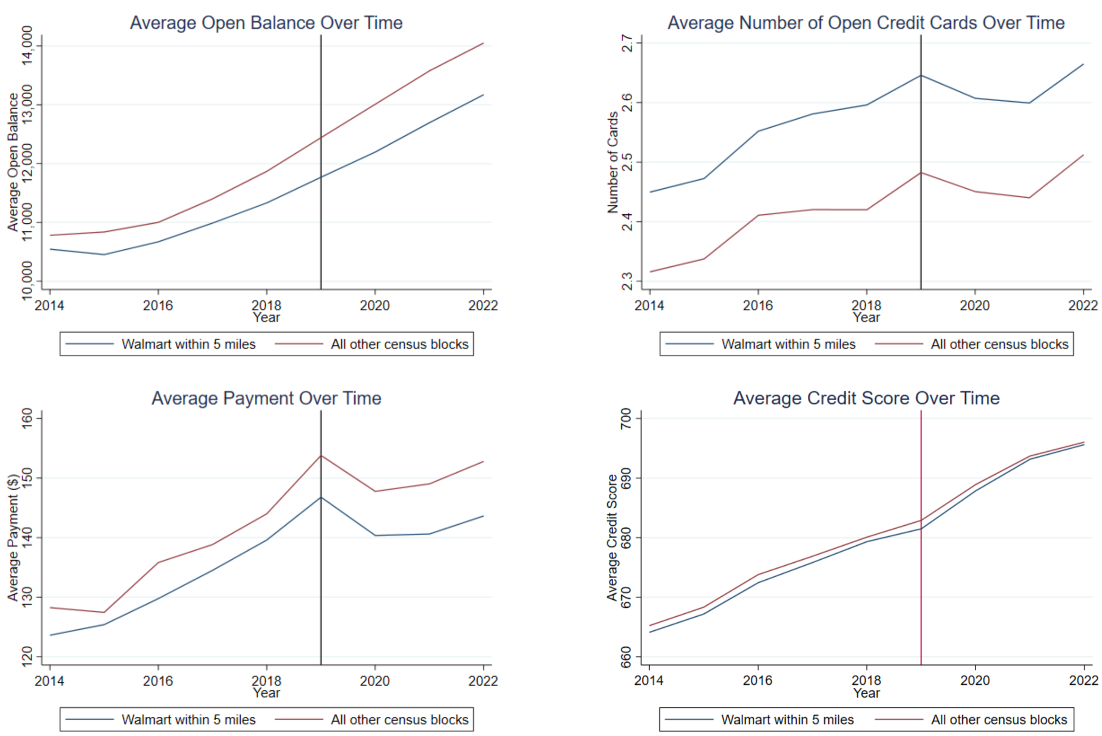

Figure 1.4: Trends in Raw Data on Borrowing Outcomes and Credit Score

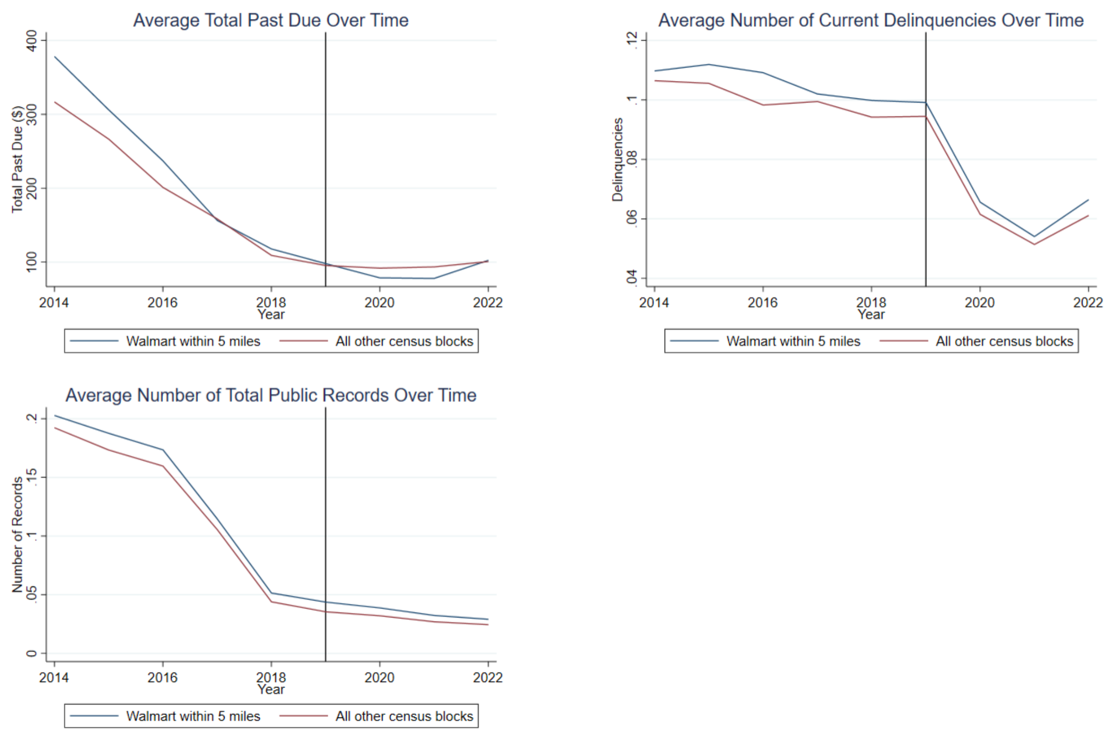

Figure 1.5: Trends in Raw Data on Financial Distress Outcomes

1.4.3 Additional Data

I use demographic data from the American Community Survey (ACS) five-year estimates to obtain census block group-level demographic data on median income, education, race, median age and total population. Table 1.6 compares 2018 demographics of census blocks within five miles of Walmart to all other census blocks. Census blocks near Walmart have a slightly lower median age, a higher percentage of college-educated residents, a higher percentage of black residents and a lower percentage of white residents than census blocks farther from Walmart. This demographic data is used to create a propensity score-trimmed sample, which is used for robustness checks. The demographic variables are not included as controls, as they provide little variation after the inclusion of census block, individual and time fixed effects.

Characteristic	**Average: Walmart farther than 5 Miles**	**Average: Walmart within 5 Miles**	**T-Test P-Value**
% High School	44.8	41.4	0.142
% College	25.4	28.7	0.024
% White	81.8	68.9	0.000
% Black	1.9	5	0.000
% Asian	3.5	10.1	0.000
% Pacific Islander	0.2	0.4	0.161
% Native American	2.7	1.5	0.001
Median Income	$84,449	$97,524	0.001
Median Age	46.9	42.2	0.000
Block Group Population	1,379	1,751	0.032

Data are from American Community Survey five-year estimates at the census block group level for all block groups in California.

Figure 1.6: 2018 Summary Statistics

I construct a novel data set to control for whether each census block is part of a banking desert. A banking desert is defined as a census tract for which the nearest bank branch is more than 10 miles from the center of the census tract. While bank desert demographics have been examined using a single year of data, to my knowledge this is the first data set that includes the banking desert status of all census tracts in the United States in more than one year. Bank desert status provides a measure of credit availability through traditional financial institutions. Controlling for bank desert status helps alleviate the concern that changes in financial outcomes when BNPL becomes available could be related to changes in the availability of other, more traditional sources of credit. While census block fixed effects largely control for the financial environment, the bank desert control is included as an added precaution.

1.5 Methodology

Variation in access to BNPL is necessary to determine a causal relationship between BNPL availability and consumers' financial outcomes. I use two sources of variation in BNPL access: geographic variation in distance to the nearest Walmart and variation over time in whether Walmart offers BNPL. This identification strategy requires three main assumptions. The first assumption is that outcomes in the treatment and control groups would have followed parallel trends in the post-treatment period if BNPL had not become available. I use event studies to support this assumption.

The second assumption is that living close to a Walmart store increases the likelihood of shopping at Walmart. Many individuals use Walmart as their grocery store, and consumers are more likely to grocery shop at a store close to their home: 76% of Americans travel less than two miles to shop for groceries, and 84% travel less than five miles [42]. Data from SafeGraph on the median distance traveled to each Walmart store provides further support for this assumption, as is described in more detail in the Appendix.

Distance to Walmart is a meaningful measure of access to BNPL despite the fact that Walmart's partnership with Affirm also made BNPL available on its website. Distance affects the likelihood of shopping in a Walmart store, and the vast majority of shopping still happens in stores: e-commerce sales as a percent of total commerce peaked at 15% in 2020 and have declined since [43]. Online and in-store shopping tend to be complements rather than substitutes [44], raising the likelihood that Walmart's in-store customers use its website. Furthermore, Walmart's in-store pickup option and same-day delivery for nearby customers both increase the likelihood that customers who live near Walmart will shop on its website more than customers who live far away. Any increase in online access experienced by both the treated and control groups will bias the results towards zero.

The third assumption is that the locations of households and Walmart stores do not depend on BNPL availability at Walmart. The vast majority of Walmart stores chose their locations years before BNPL companies formed, ensuring that Walmart did not choose neighborhoods with more financially distressed residents when BNPL became available. Figure 1.7 supports this assumption by showing that almost all Walmart stores opened before 2010, nearly a decade before BNPL became available. Similarly, consumers who would like to use BNPL at Walmart most likely did not move closer to a Walmart store when BNPL became available. The cost of moving would greatly exceed the savings from using BNPL. To address the possibility that movement between the treatment and control groups could affect the main results, I repeat the main analysis with a sample restricted to exclude individuals whose treatment status switches in the post-treatment period. Results using this restricted sample are nearly identical to the main results, as shown in the Appendix.

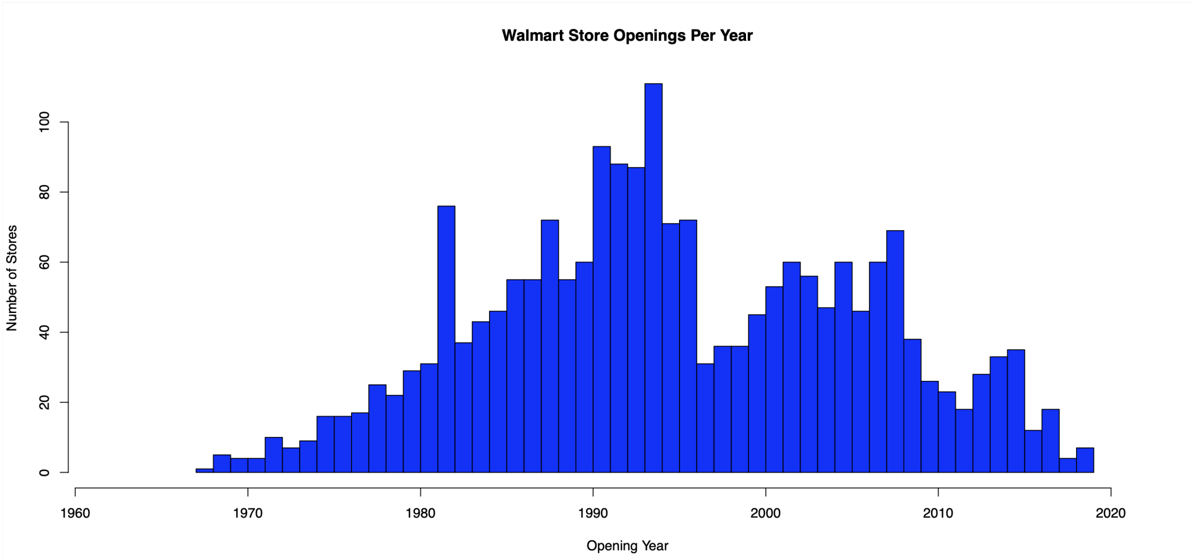

Figure 1.7: Walmart Opening Dates

I use the following regression specification to obtain the main results:

$$Outcome_{ict} = \alpha + \beta Access_{ct} + \gamma X_{it} + \delta Z_{ct} + \psi_i + \phi_c + \rho_t + \epsilon_{ict} \tag{1.11}$$

$Outcome_{ict}$ is an outcome for individual i in census block c at time t. $Access_{ct}$ is an indicator equal to one if the center of census block c is within five miles of a Walmart offering BNPL at time t, and zero otherwise. A two-mile cutoff and the natural log of distance to Walmart are used as alternative measures of treatment. X_{it} is a vector of individual-level controls that vary over time, including age and education. Z_{ct} is a control for whether the census block is in a bank desert at time t. ψ_i is an individual fixed effect, ϕ_c is a census block fixed effect and ρ_t is a year-quarter fixed effect. Standard errors are clustered at the census block level.

I estimate the effect of BNPL access on seven outcomes. Three measure financial distress: total past due, number of current delinquencies, and total public records. Delinquencies are accounts that have been past due for at least 30 days, and total public records is the number of debt-related court proceedings, such as bankruptcy filings. The fourth outcome is credit score, which incorporates information on both financial distress and use of credit. The final three outcomes measure use of non-BNPL credit: number of open credit cards, total open balance, and average monthly payment. Total open balance is the total amount of debt the consumer holds on all loans. Average payment is the average monthly payment on each loan. For all outcomes except credit score, I use the transformation $Outcome_{ict} = ln(Y_{ict} + 1)$ due to skewed distributions of the raw outcomes. I use two robustness checks to address concerns raised about this transformation in [45], described in more detail in the Robustness Checks section. For credit scores, I use the raw credit score as the outcome.

I use event studies with the following specification to assess the plausibility of the parallel trends assumption:

$$Outcome_{ict} = \alpha + \sum_{\substack{j=-5 \\ j\neq -1}}^{3} \beta_j \mathbf{1}\{t = j\} * WithinFiveMiles_c + \gamma X_{it} + \delta Z_{ct} + \psi_i + \phi_c + \rho_t + \epsilon_{ict} \tag{1.12}$$

$WithinFiveMiles_c$ is an indicator for whether census block c is within five miles of a Walmart. As for the main results, standard errors are clustered at the census block level. The number of years since BNPL became available at Walmart is j. Coefficients that are close to, and statistically indistinguishable from, zero in the pre-treatment period provide evidence of parallel pre-trends.

1.6 Results

1.6.1 Main Results

I find that BNPL access decreases financial distress. Table 1.8 displays the main results. The least extreme measure of financial distress is total past due, which is the total amount that is past the due date. I find a 2.4% decrease in total past due, which is significant at the 1% level. This estimate equates to a decrease of \$7.30 from the pre-treatment average of \$304. For the next-most extreme outcome, number of current delinquencies, I find a much smaller effect: a 0.2% decrease, significant at the 1% level. The estimated effect on the number of current delinquencies represents a 0.0016 decrease from the pre-treatment average of 0.08. For the most extreme outcome, total public records, I estimate a statistically insignificant 0.08% decrease. In terms of the model, these results indicate that $\frac{\partial R_t^{T*}}{\partial B_t^B} > 0$: individuals repay more of their traditional debts when their access to BNPL increases.

Main Results						
Measure of Treatment	**Number of Open Credit Cards**	**Total Open Balance**	**Average Payment**	**Number of Current Delinquencies**	**Total Past Due**	**Total Public Records**
5 Mile Cutoff	0.0210***	0.0427***	0.00888	-0.00299***	-0.0243***	-0.000838
	(0.00205)	(0.0104)	(0.00556)	(0.000545)	(0.00409)	(0.000731)
2 Mile Cutoff	0.0183***	0.0267*	0.00214	-0.00285***	-0.0238***	-0.00183*
	(0.00275)	(0.0141)	(0.00752)	(0.000756)	(0.00572)	(0.000999)
Log Distance	-0.0124***	-0.0250***	-0.00497	0.00158***	0.0125***	0.000722*
	(0.00112)	(0.00575)	(0.00307)	(0.000306)	(0.00234)	(0.000408)
Observations	9,929,100	9,929,100	9,929,100	9,929,100	9,929,100	9,929,100

*: p<0.1, **: p<0.05, ***: p<0.01. Clustered standard errors in parentheses. All regressions include quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and census block-level demographic controls. Note that when a continuous measure of distance is used, the coefficients are the estimated effect of being farther away from BNPL, whereas the estimates for five- and two-mile cutoffs show the effect of being closer to BNPL.

Figure 1.8: Main Results

Figure 1.9, Figure 1.10 and Figure 1.11 provide evidence of parallel pre-trends for the three repayment outcomes. F-tests show that pre-treatment estimates in these event studies are also jointly insignificant, as shown in Table 1.12.

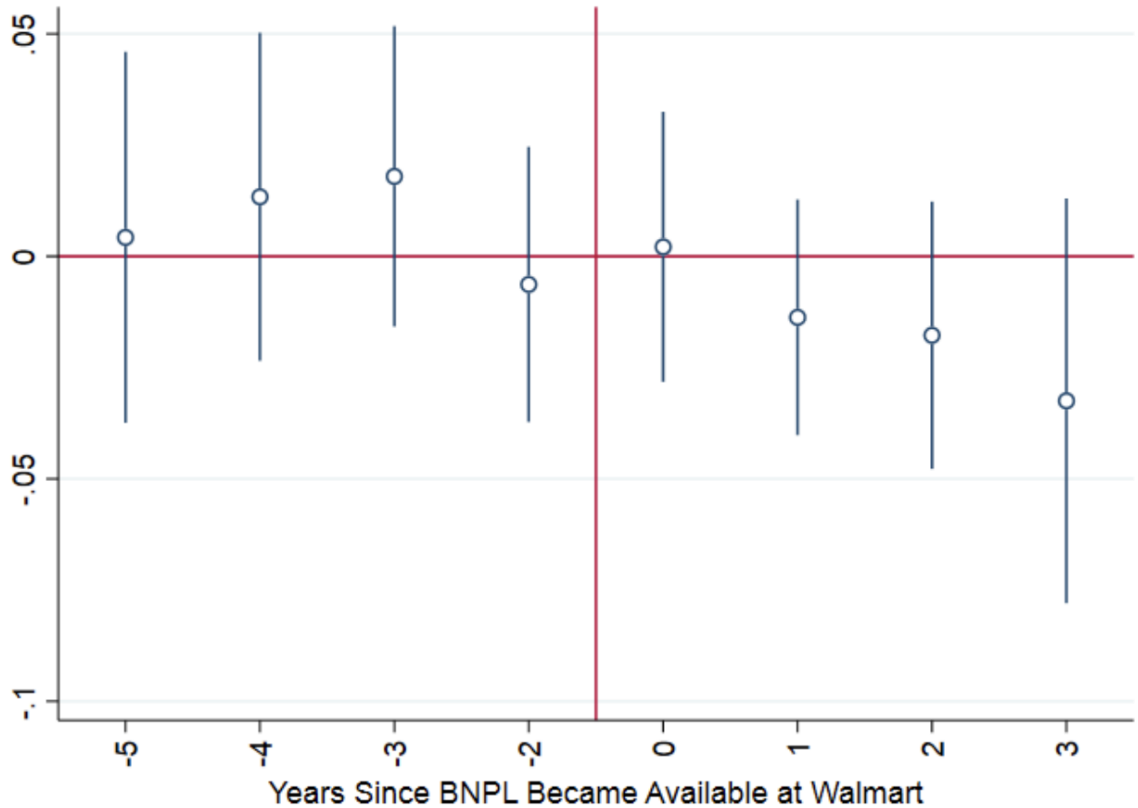

Logged total amount past due is the outcome. Includes quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and a census block-level bank desert control. Standard errors are clustered at the census block level. Five miles from Walmart is used as the cutoff for treatment.

Figure 1.9: Event Study: Total Past Due

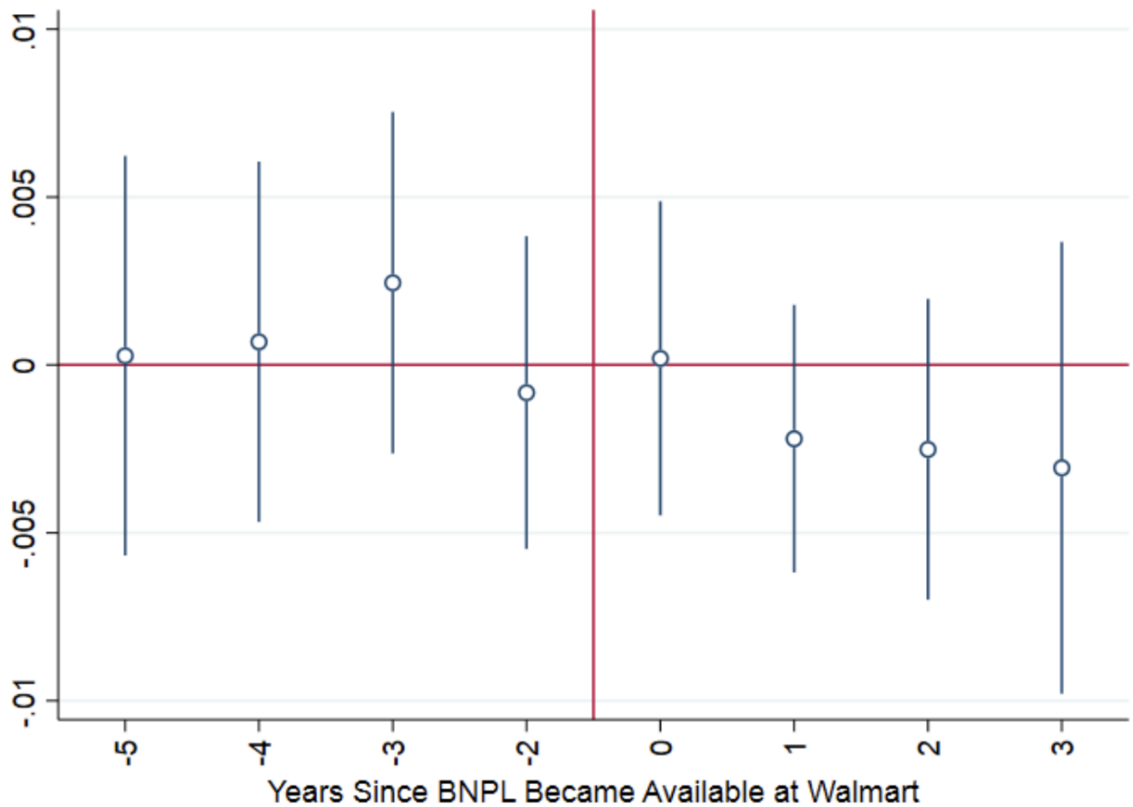

Logged number of current delinquencies is the outcome. Includes quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and census block-level demographic controls. Standard errors are clustered at the census block level. Five miles from Walmart is used as the cutoff for treatment.

Figure 1.10: Event Study: Number of Current Delinquencies

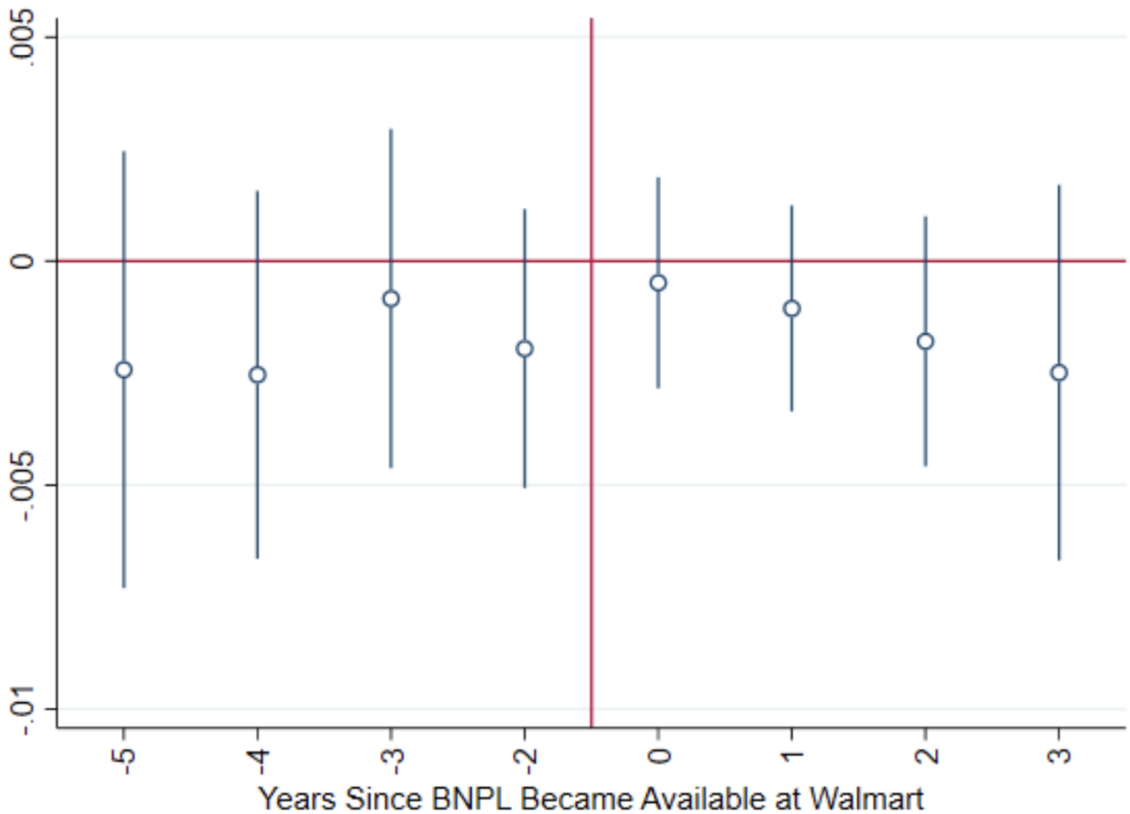

Logged total public records is the outcome. Includes quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and a census block-level bank desert control. Standard errors are clustered at the census block level. Five miles from Walmart is used as the cutoff for treatment.

Figure 1.11: Event Study: Total Public Records

Event Study F-Test Results	
Outcome	**P-Value**
Total Open Balances	0.4688
Number of Open Credit Cards	0.0000
Average Payment	0.6480
Total Past Due	0.6476
Number of Current Delinquencies	0.7705
Total Public Records	0.6340
Credit Score	0.2182

P-values from F-tests for joint significance in the pre-treatment period of event studies.

Figure 1.12: Event Study F-Tests

There are two main takeaways from the repayment results. First, despite concerns raised in the literature, the press and among regulators, BNPL does not appear to lead to over-borrowing. Far from causing a "debt spiral", BNPL reduces debt-related financial distress, likely due to the increased ability to smooth consumption while paying low or zero interest. Volatility in expenses or income can lead to a missed payment on a credit card or car loan, especially for financially vulnerable individuals with low savings and limited access to credit. Even a small low-interest loan such as BNPL can help these consumers avoid missing a loan payment in a month in which they incur an unexpected expense or experience a negative income shock.

Second, while BNPL does appear to improve financial well-being, its largest effect is on the least extreme measure of financial distress. BNPL availability does not reduce the likelihood of bankruptcy, but it does help consumers pay monthly bills on time. The small size of BNPL loans may explain why they only affect less extreme measures of financial distress, and why the effects are economically small. While [28] and [29] find that BNPL leads to negative effects on financial well-being in the form of higher overdraft fees and lower bank account balances, I find no evidence of negative effects on more extreme outcomes. Worsening of the financial distress outcomes I study could have longer-lasting impacts than incurring an overdraft or low balance fee from a bank because these outcomes are reported to credit bureaus and can damage credit scores.

Next, I examine the effect of BNPL access on credit scores. Changes in a consumer's credit score matter for two reasons: they provide a measure of her creditworthiness today, and they determine the credit limit and interest rate that lenders will offer her in the future. Table 1.13 shows that BNPL access increases the average credit score by 1.6 points. This estimate is statistically significant at the 1% level. While the effect is economically small, as credit scores range from 300 to 850, it is a positive effect and it shows that BNPL access does not damage credit scores. The small improvement in credit scores likely results from reduced financial distress. Figure 1.14 displays the event study for credit score, which largely supports a causal interpretation of this result. While zero is outside the confidence interval for one pre-treatment estimate, an F-test shows that the pre-treatment estimates are jointly insignificant.

Credit Score	
Measure of Treatment	**Credit Score**
5 Mile Cutoff	1.569***
	(0.294)
2 Mile Cutoff	1.130***
	(0.399)
Log Distance	-1.014***
	(0.164)
Observations	9,929,100

*: p<0.1, **: p<0.05, ***: p<0.01. Robust standard errors in parentheses. Regression includes quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and census block-level demographic controls.

Figure 1.13: Credit Score Results

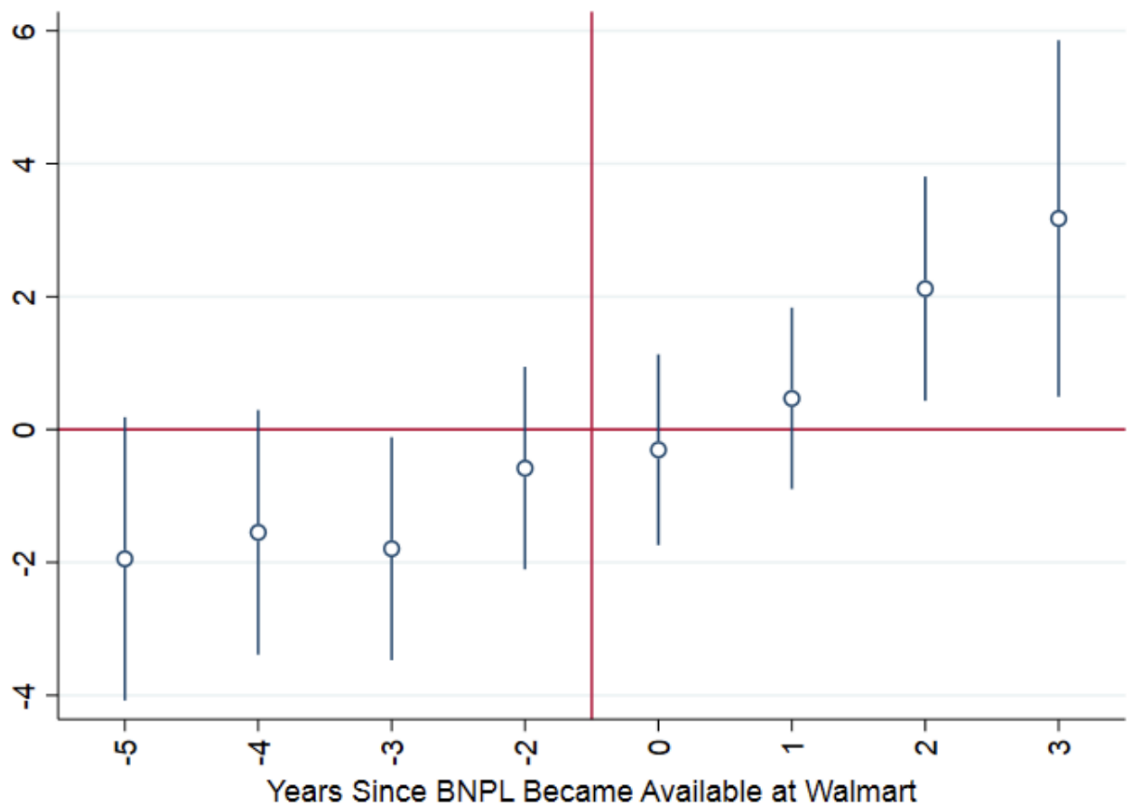

Credit score is the outcome. Includes quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and a census block-level bank desert control. Standard errors are clustered at the census block level. Five miles from Walmart is used as the cutoff for treatment.

Figure 1.14: Event Study: Credit Score

Finally, I examine the effect of BNPL access on use of traditional credit. Two measures of traditional credit use increase in response to BNPL availability: number of open credit cards and total open balance. For my main specification, which uses five miles from the nearest Walmart as the cutoff for treatment, I estimate a 2.1% increase in the number of open credit cards and a 4.3% increase in the total open balance. Both estimates are significant at the 1% level. The estimated increase in the total open balance is $141 for the median consumer, with a total open balance of $3,298 prior to treatment. The estimated effect on number of open credit cards equates to an increase of 0.049 open credit cards, compared with the pre-treatment average of 2.34. The estimated effect on average monthly payment is a statistically insignificant increase of 0.1%.

The estimated effects on total open balances and average monthly payment can be interpreted as causal, while the estimated effect on number of open credit cards should be interpreted as suggestive evidence of an increase. Figure 1.15 and Figure 1.16 provide evidence of parallel trends for total open balance and average monthly payment prior to treatment. Figure 1.17 shows evidence of a pre-trend for number of open credit cards, particularly in the early part of the sample period. F-tests confirm that the pre-treatment estimates for total open balances and average monthly payment are jointly insignificant, while pre-treatment estimates for number of open credit cards are jointly significant, as shown in Table 1.12.

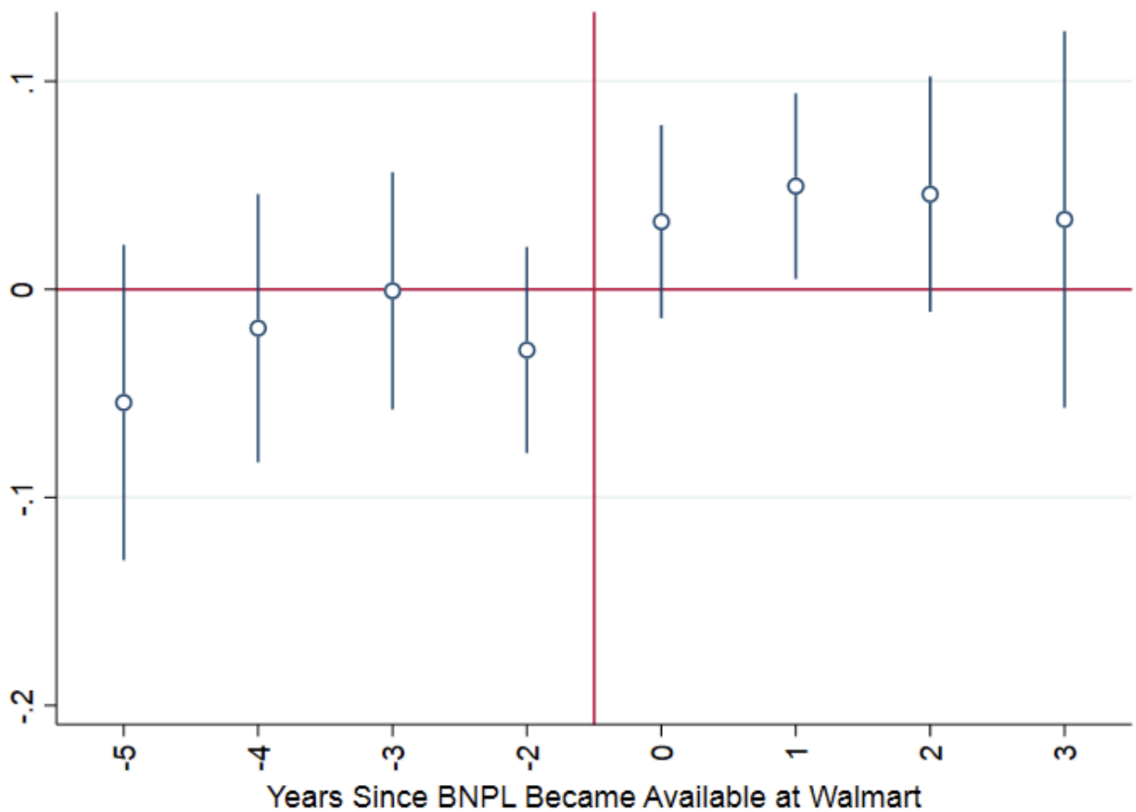

Logged total open balance is the outcome. Includes quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and a census block-level bank desert control. Standard errors are clustered at the census block level. Five miles from Walmart is used as the cutoff for treatment.

Figure 1.15: Event Study: Total Open Balances

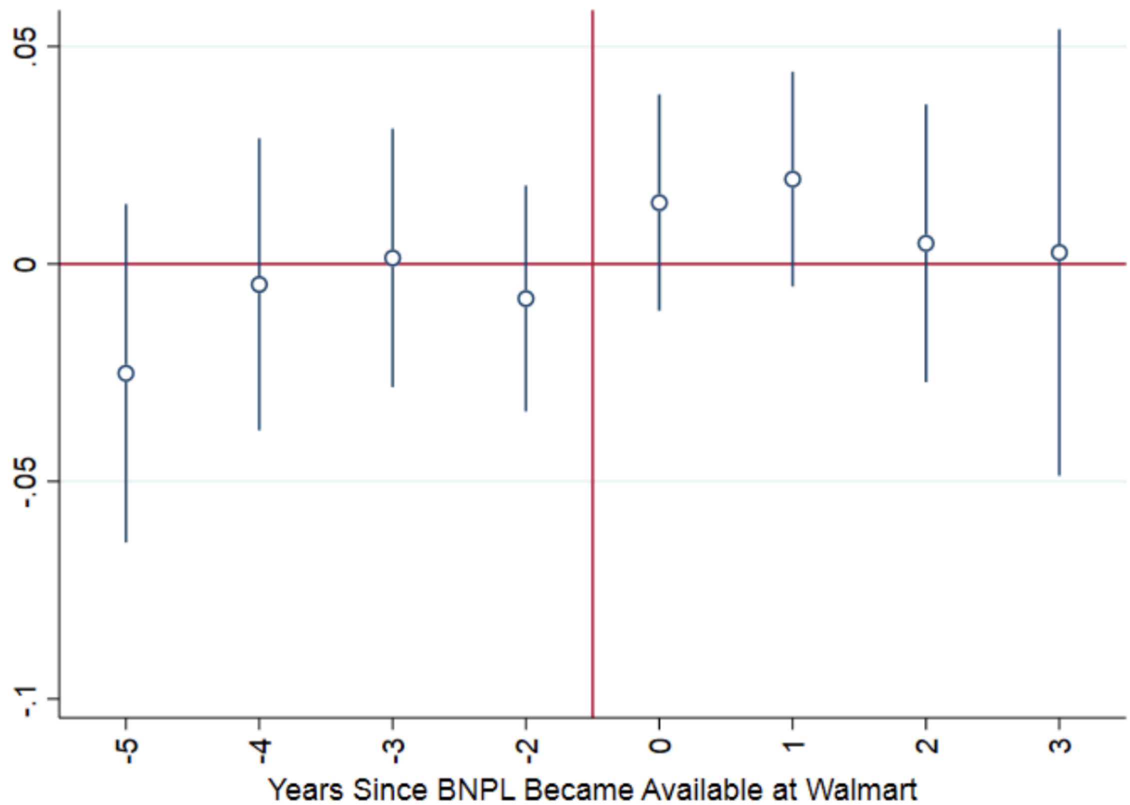

Logged average payment is the outcome. Includes quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and a census block-level bank desert control. Standard errors are clustered at the census block level. Five miles from Walmart is used as the cutoff for treatment.

Figure 1.16: Event Study: Average Monthly Payment

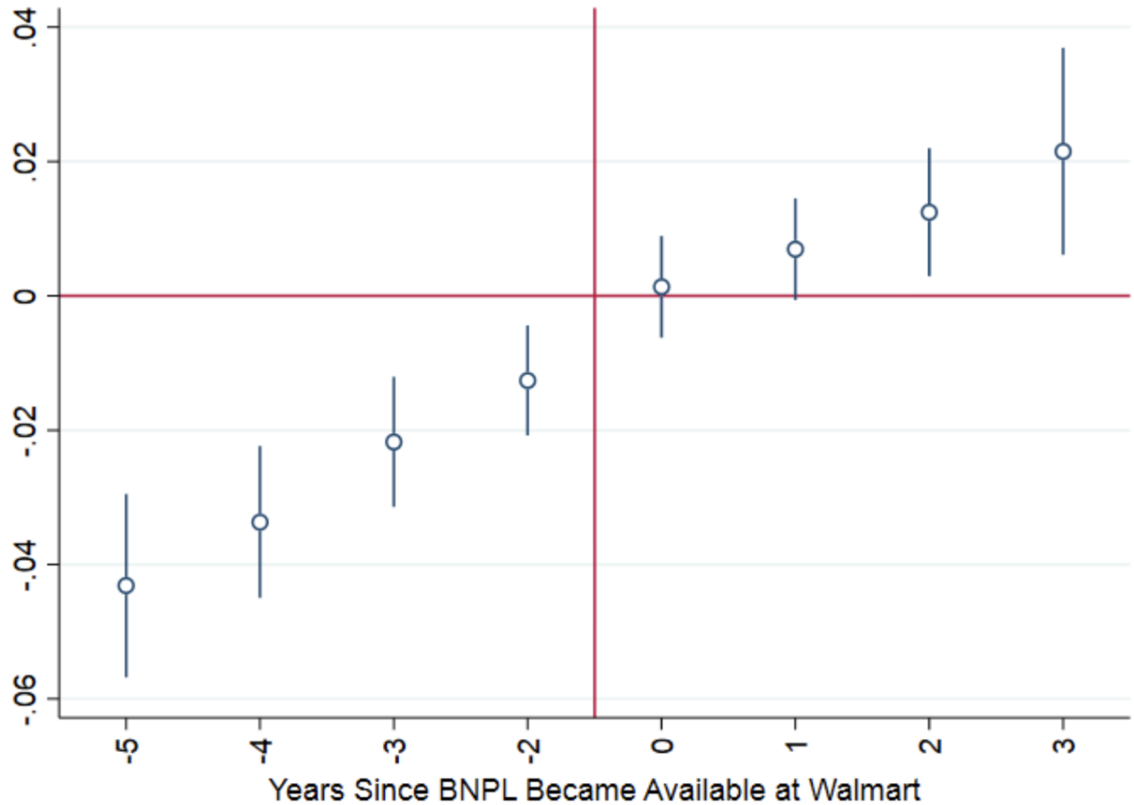

Logged number of open credit cards is the outcome. Includes quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and a census block-level bank desert control. Standard errors are clustered at the census block level. Five miles from Walmart is used as the cutoff for treatment.

Figure 1.17: Event Study: Number of Open Credit Cards

In terms of the model, these results indicate that $\frac{\partial L_t^{T*}}{\partial B_t^B} > 0$. This increase in use of traditional credit is likely driven by two mechanisms: consumers using credit cards to repay BNPL loans and consumers gaining better access to traditional credit as their credit scores improve. Descriptive studies provide evidence that consumers frequently use credit cards to repay BNPL loans [30] and that BNPL users are typically active users of traditional credit [7]. When BNPL is available, consumers may put BNPL payments on a credit card instead of making a purchase with a debit card or cash, which could increase use of traditional credit without changing total consumption.

Increased utilization of credit can lower a credit score, especially if the consumer already uses a high percentage of her available credit. Consumers are advised to use no more than 30% of their available credit to avoid lowering their credit score [46]. However, the finding of an overall increase in credit scores shows that the positive effect of reduced financial distress dominates the negative effect on credit scores of increasing use of traditional credit. Because credit scores increase as a result of BNPL access, consumers may be more likely to be approved for credit cards and may see their credit limits increase, enabling them to increase their use of traditional credit.

1.6.2 Heterogeneity Analysis

While the main results show that BNPL access reduces financial distress on average, it is still possible that groups within the population are worse-off when they have access to BNPL. Groups of particular concern could be consumers with low credit scores, which can result from a history of failing to repay debt, and young consumers. I use heterogeneity analysis to investigate whether the effects of BNPL access vary by consumers' pre-treatment credit score categories or demographic characteristics.

First, I examine heterogeneity by credit score category prior to treatment. I group credit scores into the standard categories defined by FICO. Table 1.18 shows how each category is defined, with the highest credit scores in the "exceptional" category, followed by "very good", "good", "fair" and then "poor." Table 1.19 displays the results. I find that directions of the effects do not vary by group and effects are strongest among individuals who had "fair" credit scores before BNPL was available: decreases in financial distress and increases in use of credit in this group are similar to the main results in terms of both magnitude and significance.

FICO Score Ranges	
Category	**Credit Score Range**
Poor	300-579
Fair	580-669
Good	670-739
Very Good	740-799
Exceptional	800-850

Note:
Source: FICO

Figure 1.18: FICO Credit Score Ranges

Heterogeneity by Pre-Treatment Credit Score Range						
Credit Score Range	**Number of Open Credit Cards**	**Total Open Balance**	**Average Payment**	**Number of Current Delinquencies**	**Total Past Due**	**Total Public Record Amount**
Poor n=1,708,091	0.00784 (0.00596)	0.0590* (0.0341)	0.0284 (0.0183)	-0.00306 (0.00238)	-0.0134 (0.0177)	0.0252 (0.0193)
Fair n=2,100,604	0.0254*** (0.00496)	0.0420* (0.0246)	0.00515 (0.0136)	-0.00321** (0.00153)	-0.0289*** (0.0111)	0.0185 (0.0151)
Good n=1,887,689	0.0256*** (0.00481)	0.0491** (0.0239)	0.00533 (0.0129)	-0.00112 (0.000964)	-0.00976 (0.00785)	0.0188 (0.0123)
Very Good n=1,920,715	0.0185*** (0.00401)	0.0181 (0.0212)	-0.00220 (0.0110)	-0.000558 (0.000448)	-0.00726* (0.00386)	0.00443 (0.00778)
Exceptional n=1,780,735	0.00134 (0.00379)	-0.00686 (0.0166)	0.00227 (0.00893)	-0.000284 (0.000269)	-0.00478* (0.00257)	0.00461 (0.00513)

*: p<0.1, **: p<0.05, ***: p<0.01. Robust standard errors in parentheses. All regressions include quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and census block-level demographic controls. Five miles from Walmart is used as the cutoff for treatment. Credit score ranges are determined by credit score in the last quarter before BNPL became available at Walmart.

Figure 1.19: Heterogeneity by Pre-Treatment Credit Score

The next-most affected category is “good”: effects on all three financial distress measures for this group are negative, but statistically insignificant and close to zero, while estimated effects on credit use are positive and slightly larger in magnitude than for people with “fair’ scores. People with”very good” credit scores increased their number of open credit cards by 1.9% as a result of BNPL, an estimate which is smaller in magnitude to the main results and is significant at the 1% level. No other estimates are significant at the 5% level for individuals in this group. BNPL did not affect people at either extreme of the credit score range: no estimates are significant a the 5% level for individuals with “poor” or “exceptional” credit.

These results do not show evidence of increased financial distress for consumers in any credit score category. The effect sizes vary due to differences across these groups in access to BNPL and traditional credit. People with “poor” credit are unlikely to qualify for BNPL and are therefore unlikely to be affected by BNPL availability. At the other extreme, people with “exceptional” credit are also unlikely to be affected: these people already have access to very low-interest forms of credit, and they are more likely to have a preference for paying bills off immediately rather than using a service such as BNPL. Those with “fair” credit scores have the most to gain from BNPL access, as they do not have access to other low-interest forms of credit but their credit scores are high enough to qualify for BNPL. 22% of the sample is in the “fair” category. People with “good” credit scores have slightly better non-BNPL credit options than people with “fair” credit scores, but still may not qualify for many credit cards and are likely charged high interest rates. Therefore, the strongest effects of BNPL on both borrowing and repayment outcomes are among individuals with “fair” credit scores, followed by those with “good” credit scores. These findings also provide support for the identification strategy, as I find the strongest effects among the consumers who are most likely to use BNPL.

Next, I repeat the main analysis on a sample restricted to people under age 40.

Determining whether BNPL access increases financial distress among young consumers is particularly important because young consumers under 40 are the most likely to use BNPL [5, 47, 48, 9, 7]. Table 1.20 displays the results. The directions and statistical significance of all six estimates are the same as the main results, and magnitudes are larger for individuals under 40 for five of the six outcomes. These findings show that BNPL access reduces financial distress among young consumers. The results also provide support for the identification strategy by showing larger effects in the population most likely to use BNPL. Additional heterogeneity analysis by gender, education, marital status and home ownership appears in Table 1.20 Table 1.21.

Heterogeneity by Age and Gender						
Sample	**Number of Open Credit Cards**	**Total Open Balance**	**Average Payment**	**Number of Current Delinquencies**	**Total Past Due**	**Total Public Record Amount**
Under 40 n=3,661,994	0.0180*** (0.00371)	0.0591*** (0.0197)	0.00694 (0.0104)	-0.00385*** (0.00118)	-0.0299*** (0.00763)	-0.0189*** (0.00534)
Women n=4,433,699	0.0273*** (0.00317)	0.0366** (0.0151)	-0.00139 (0.00800)	-0.00355*** (0.000839)	-0.0248*** (0.00578)	0.00582 (0.00741)
Men n=4,702,460	0.0125*** (0.00288)	0.0400*** (0.0154)	0.0110 (0.00832)	-0.00276*** (0.000798)	-0.0250*** (0.00641)	-0.00240 (0.00850)

*: p<0.1, **: p<0.05, ***: p<0.01. Robust standard errors in parentheses. All regressions include quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and census block-level demographic controls. Five miles from Walmart is used as the cutoff for treatment.

Figure 1.20: Heterogeneity by Age and Gender

Heterogeneity by Education, Marital Status, Homeownership and Pre-Covid						
Sample	**Number of Open Credit Cards**	**Total Open Balance**	**Average Payment**	**Number of Current Delinquencies**	**Total Past Due**	**Total Public Record Amount**
High School Educated	0.0150***	0.0254	0.00610	-0.000784	-0.00733	0.00561
n=2,059,304	(0.00461)	(0.0234)	(0.0123)	(0.00136)	(0.00986)	(0.0117)
College Educated	0.0106***	0.0333**	0.0111	-0.00184**	-0.0133**	0.0140
n=2,711,615	(0.00359)	(0.0169)	(0.00936)	(0.000799)	(0.00609)	(0.00984)
Single	0.0236***	-0.00349	-0.0118	-0.00357**	-0.0410***	-0.0178
n=1,881,291	(0.00491)	(0.0274)	(0.0148)	(0.00172)	(0.0119)	(0.0132)
Married	0.0162***	0.0328**	0.00187	-0.00287***	-0.0227***	-0.00267
n=5,396,689	(0.00277)	(0.0136)	(0.00736)	(0.000729)	(0.00565)	(0.00750)
Homeowner	0.0206***	0.0521***	0.00595	-0.00278***	-0.0216***	-0.00361
n=4,308,891	(0.00301)	(0.0143)	(0.00764)	(0.000725)	(0.00575)	(0.00725)
Before Covid-19	0.0148***	0.0181**	-0.000231	-0.000486	-0.00867*	-3.81e-05
5,557,524	(0.00172)	(0.00910)	(0.00468)	(0.000655)	(0.00443)	(0.00510)

*: p<0.1, **: p<0.05, ***: p<0.01. Robust standard errors in parentheses. All regressions include quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and census block-level demographic controls. Five miles from Walmart is used as the cutoff for treatment.

Figure 1.21: Additional Heterogeneity

A striking takeaway from the heterogeneity analysis is that the directions of the effects of access to BNPL do not vary systematically across demographic groups. While magnitudes and significance do vary by group, showing which individuals are most affected by BNPL access, I do not identify any demographic groups for whom BNPL access worsens financial distress. This finding may help alleviate concerns of systematic misuse of BNPL among groups that are considered vulnerable, such as young adults who are still developing financial literacy [49]. Individuals with fair credit, who have either struggled to repay debts in the past or have yet to build a credit history, also benefit from access to BNPL loans.

1.6.3 Placebo Tests

Event studies support the argument that the main results are causal by showing parallel trends in the treated and control groups before BNPL became available. However, even if Walmart had not introduced BNPL, it is still possible that financial distress would have decreased and credit use would have increased more among individuals living near Walmart than among those who live far away. I use two placebo tests to assess whether the main results capture a trend over time among consumers living near Walmart that would have occurred in the absence of BNPL.

For the first placebo test, I repeat the main analysis using only data from Iowa and West Virginia. In February 2019, Walmart made BNPL available in all its American stores except in these two states. This placebo test provides insight into whether individuals living close to Walmart would have improved debt repayment and increased use of credit more than individuals living farther from Walmart even in the absence of BNPL. For this placebo test, I use a 0.02% national sample of credit bureau data, which I restrict to only include individuals who live in Iowa and West Virginia. This data is identical to the data used to obtain the main results, except that the sample size is smaller and individual-level demographic controls are unavailable. Therefore, I also repeat the main analysis of California data without controls to ensure that the absence of controls does not drive any differences between the main results and the placebo test.

Table 1.22 shows the results. None of the placebo estimates are statistically significant at the 5% level, while five of the six estimates for the main results are statistically significant at this level. The placebo estimates for number of open credit cards, total open balance, number of current delinquencies, and total past due have smaller magnitudes than the main results. Furthermore, the placebo estimates for number of current delinquencies, total past due, and average payment have the opposite sign of the main results. The results of this placebo test support the argument that the main results capture the effect of BNPL, rather than an effect that would have happened even in the absence of BNPL. A limitation of this placebo test is that the sample, while not small at 56,812, is much smaller than the sample used to obtain the main results. One possibility is that the placebo test does not show statistically significant effects due to the smaller sample size. I use a second placebo test with a larger sample size to address this concern.

Placebo Test: Iowa and West Virginia						
	Number of Open Credit Cards	**Total Open Balance**	**Average Payment**	**Number of Current Delinquencies**	**Total Past Due**	**Total Public Records**
Iowa and West Virginia Placebo Test						
BNPL Access	0.0154	0.00942	-0.0257	0.00211	0.00973	-0.0311*
	(0.0244)	(0.109)	(0.0589)	(0.00859)	(0.0471)	(0.0160)
Observations	56,812	56,812	56,812	56,812	56,812	56,812
California Results						
BNPL Access	0.0202***	0.0469***	0.0119**	-0.00287***	-0.0235***	-0.00113*
	(0.00191)	(0.00993)	(0.00538)	(0.000502)	(0.00377)	(0.000672)
Observations	11,016,285	11,016,285	11,016,285	11,016,285	11,016,285	11,016,285

*: p<0.1, **: p<0.05, ***: p<0.01. Clustered standard errors in parentheses. All regressions include quarter-year, census block and individual fixed effects. The placebo regressions only include data from West Virginia and Iowa.

Figure 1.22: Placebo Test: Iowa and West Virginia

For the second placebo test, I repeat the main analysis using an indicator for whether an individual lives within five miles of Target to define the treated group. Target and Walmart sell similar products, and both offer low-cost options for items such as clothing and household goods. This placebo test would determine whether the main results capture effects that are common to all consumers living close to a low-cost, big-box retailer such as Target or Walmart, rather than the effect of BNPL becoming available at Walmart. I restrict the sample for this placebo test to before June 2021, when Target introduced its own BNPL partnership. I also restrict the sample to exclude individuals who live within five miles of Walmart to ensure that these results do not capture the effect of Walmart introducing BNPL. The benefit of this placebo test is that I can conduct it with California data, and therefore the sample size is larger than for the first placebo test. A limitation of this test is that Target and Walmart are different retailers, and individuals who live close to Target may differ from individuals who live close to Walmart. I repeat the main analysis with the same time restriction to facilitate comparison with the placebo results.

Table 1.23 shows the results of the second placebo test. Placebo estimates of effects on the number of open credit cards, total open balance, average payment, and number of current delinquencies are statistically insignificant, with magnitudes that are smaller than the main results and very close to zero. The estimated effect on the total amount past due is statistically significant at the 10% level, and the effect on total public records is significant at the 1% level. However, both of these effects have the opposite direction of the main results, providing further evidence that the main results do not capture a general trend among all consumers living close to a large retailer. This placebo test has a large sample size of 4.96 million observations and is therefore more comparable to the main results in terms of power. Table 1.23 also shows estimated effects of BNPL at Walmart with the sample restricted to before June 2021. Estimated effects of BNPL at Walmart in this time frame are nearly identical to the main results.

	Placebo Test: Within 5 Miles of Target					
	Number of Open Credit Cards	**Total Open Balance**	**Average Payment**	**Number of Current Delinquencies**	**Total Past Due**	**Total Public Records**
Target Placebo Test						
Near Target	-0.00503	0.00179	0.00724	0.000513	0.0127*	0.00412***
	(0.00350)	(0.0183)	(0.00984)	(0.000958)	(0.00742)	(0.00141)
Observations	4,963,731	4,963,731	4,963,731	4,963,731	4,963,731	4,438,380
Comparable Walmart Results						
Near Walmart	0.0193***	0.0396***	0.00761	-0.00287***	-0.0235***	-0.000529
	(0.00203)	(0.0104)	(0.00552)	(0.000573)	(0.00425)	(0.000726)
Observations	8,891,286	8,891,286	8,891,286	8,891,286	8,891,286	8,891,286

*: p<0.1, **: p<0.05, ***: p<0.01. Clustered standard errors in parentheses. All regressions include quarter-year, census block and individual fixed effects. The Target placebo test excludes observations within five miles of Walmart. Both regressions exclude observations after Target introduced BNPL in June 2021.

Figure 1.23: Placebo Test: Target

1.7 Robustness Checks

1.7.1 Alternative Distance Cutoffs

The first robustness check uses two alternative measures of treatment to ensure that the main results do not depend on using a cutoff of exactly five miles to define treatment. Results from two alternative measures of treatment are very similar to the main results. For the first alternative specification, I use a two mile distance from the nearest Walmart as the cutoff for treatment. I estimate that access to BNPL increases the number of open credit cards by 1.83% and increases total open balance by 2.67%. These estimates are significant at the 1% and 10% levels, respectively. As with the main results, the estimated effect on average payment is statistically insignificant and close to zero at 0.2%. For the financial distress outcomes, I estimate a 2.38% decrease in total past due that is significant at the 1% level; a 0.3% decrease in number of current delinquencies that is significant at the 1% level; and a 0.2% decrease in total public records that is significant at the 10% level. The estimated effect on credit score is an increase of 1.13 points, which is nearly identical to the main result and is statistically significant at the 1% level.

For the second alternative specification, I use the log of distance from the nearest Walmart as a continuous measure of distance. I estimate that a 1% increase in distance away from Walmart when BNPL is available decreases number of open credit cards by 1.24%, decreases total open balance by 2.5%, increases number of current delinquencies by 1.58%, and increases total amount past due by 1.25%. All four estimates are significant at the 1% level. As with the main results, the estimated effects on average payment and total public records are statistically insignificant at the 5% level and close to zero. For credit score, I estimate that a 1% increase in distance from Walmart decreases credit score by 0.01 points. This estimate has the same direction as the main result and is also significant at the 1% level.

1.7.2 Propensity Score Trimming

As a second robustness check, I restrict the sample to only include census blocks with propensities between 0.1 and 0.9 of being within five miles of Walmart, using the standard cutoffs established in [50]. Census blocks within five miles of Walmart have statistically significant demographic differences from census blocks more than five miles from Walmart, raising the question of whether census blocks farther than five miles away are an informative control group, even when event studies provide support for the parallel trends assumption. Using propensity scores to trim the sample helps to address this concern. The propensity scores are predicted values from a regression of an indicator for being within five miles of Walmart on demographic variables from the American Community Survey in 2018, the last year prior to treatment. Despite differences in the averages of demographic variables shown in Table 1.6, propensity score trimming barely reduces the sample size, from 9,929,100 to 9,782,360. Results from the propensity score-matched sample are nearly identical to the main results, as shown in Table 1.1.

Sample	Number of Open Credit Cards	Total Open Balance	Average Payment	Number of Current Delinquencies	Total Past Due	Total Public Records
Propensity Score Trimmed	0.0207***	0.0408***	0.00842	-0.00295***	-0.0241***	-0.000871
	(0.00207)	(0.0104)	(0.00561)	(0.000549)	(0.00412)	(0.000736)
Observations	9,782,360	9,782,360	9,782,360	9,782,360	9,782,360	9,792,230
Before COVID-19	0.0148***	0.0181**	-0.000231	-0.000486	-0.00867*	8.51e-05
	(0.00172)	(0.00910)	(0.00468)	(0.000655)	(0.00443)	(0.000644)
Observations	5,557,524	5,557,524	5,557,524	5,557,524	5,557,524	5,557,524

*: $p<0.1$, **: $p<0.05$, ***: $p<0.01$. Clustered standard errors in parentheses. All regressions include quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and a census block-level bank desert control. A five mile cutoff is used to define treatment.

Table 1.1: Robustness Checks

1.7.3 Restriction to Before COVID-19

Next, because the COVID-19 pandemic started during the treatment period, I restrict the sample to before 2020 to determine whether the effect of BNPL was different before the pandemic. With this restriction, the data includes only four treated quarters. Table 1.1 shows that the number of open credit cards increased by an average of 1.48% and total open balance increased by 1.81%. These estimates are significant at the 1% and 5% levels, respectively. Estimates for the three financial distress outcomes have the same direction as the main results, but the magnitudes are smaller and only the change in number of current delinquencies is statistically significant at the 10% level. Magnitudes may have increased over time as buyers became more aware of BNPL availability at Walmart.

1.7.4 Extensive Margin Results and Normalization

For all outcomes Y except credit score, I use the transformed outcome $ln(Y + 1)$ to obtain the main results. I use the standard interpretation of the results as percent changes. While a percent change should be invariant to scaling of the outcome variable, [45] shows that the percent change obtained using the transformed $ln(Y + 1)$ variable does depend on scaling of the outcome variable and can therefore lead to misleading interpretations of the magnitudes of coefficients. I perform two additional analyses to address this concern. First, I use a linear probability model to estimate results at the extensive margin. The outcomes in this analysis are indicator variables and are therefore not subject to the interpretation issue identified in [45]. Second, I use normalized outcomes of the form $\frac{Y}{X}$ as suggested in [45], where X is the population average of Y prior to treatment. These normalized outcomes are already in the form of a percent, and results can therefore be interpreted as percentage point changes at the intensive margin.

The extensive margin results are very similar to the main results. I estimate increases in the probability of using credit: the probability of having any open credit cards increases by 0.8 percentage points, the probability of any open balance increases by 0.1 percentage points, and the probability of having any monthly payment increases by 0.3 percentage points. The estimates for any credit cards and any monthly payment are significant at the 1% level. I also estimate decreases in the probability of financial distress: the probability of having any money past due decreases by 0.361 percentage points and the probability of any current delinquencies decreases by 0.359 percentage points. Both estimates are significant at the 1% level. The estimated effect on the probability of having any public records is small, negative and insignificant.

Results using the normalized outcomes appear in the Appendix. The normalized results are also consistent with the main results: I estimate increases in the number of open credit cards and total open balance, and decreases in the total amount past due and number of current delinquencies. The magnitudes of the normalized estimates are similar to the main results, suggesting that in this setting the use of the $ln(Y+1)$ transformation does not lead to a misleading interpretation of magnitudes.

1.8 Conclusion

The rapid growth of BNPL has raised questions about whether consumers will be able to reap the benefits of these inexpensive, easily accessible loans without incurring unsustainable levels of debt. Because this new form of credit is not subject to many of the regulations designed to protect consumers in more traditional lending markets, BNPL provides a fairly rare modern example of consumers and lenders operating in a credit market with few consumer protection laws.

I provide the first causal evidence of how access to BNPL affects severe measures of financial distress and credit scores. I find that BNPL access reduces financial distress, likely due to the increased ability to smooth consumption when BNPL is available, and that the reduction in financial distress is most pronounced among consumers who had "fair" credit scores prior to treatment. I also find that increased access to BNPL leads to a small improvement in credit scores. These results suggest that consumers benefit from access to BNPL rather than using it to accumulate debts they cannot afford to repay. Policymakers should consider these results as they decide whether new laws are needed to regulate BNPL.

Chapter 2

Do Democracy Vouchers Help Democracy?

2.1 Introduction

Policies that provide public financing for political campaigns have gained popularity in the United States. Between 2010 and 2018, 14 states and 24 municipalities began providing public financing with the goal of limiting candidates' reliance on large, private donations from wealthy citizens [51]. A 2018 poll found that 72% of Americans believe big donors have a disproportionate influence on election outcomes [52]. The perception that elections are decided by wealthy elites rather than ordinary Americans could contribute to low voter turnout: two thirds of people who did not vote in the 2020 election agreed that "voting has little to do with the way that real decisions are made in this country", and when asked about the best way to encourage more people to vote, these non-voters gave "cleaning up government" as the most common answer [53]. With 80 million Americans choosing not to vote in the 2020 election [53], programs that limit private financing and provide public funding for campaigns have the potential to increase political participation by changing the perception that candidates are dependent on wealthy donors and are not responsive to ordinary people.

The city of Seattle, Washington implemented its Democracy Vouchers program in 2017 in the hopes of reducing candidates' reliance on large donations [51]. This program gives every registered voter in Seattle $100 worth of publicly funded Democracy Vouchers to donate to candidates for municipal office. Candidates who accept vouchers agree to limits on non-voucher contributions. In this paper, I study the Democracy Vouchers' effects on political participation, measured through voter registration and turnout. I also examine how the program affected candidates' reliance on large and small contributions, as well as the effect on the total cost of elections. To my knowledge, this paper is the first to provide evidence of the effects of public financing for political campaigns on political participation.

I use data on voter registration, voter turnout and campaign donations in King County, where Seattle is located, from 2009 to 2021. Because Seattle is the only treated group in my sample, I use the method developed by [54] for inference. The first main result is that Democracy Vouchers increase political participation, raising voter turnout by 4.9 percentage points. This estimate equates to a 9% increase from the pre-treatment average in Seattle. I estimate that voter registration increases by 23 voters per precinct, which is equivalent to a 6% increase from the pre-treatment average in Seattle, although the estimated effect on voter registration is statistically insignificant. The increase in voter turnout shows that the Democracy Vouchers program improves political participation. One reason for increased participation could be the unique design of the program, which engages voters by allowing them to decide how public financing should be allocated. The process of researching candidates to decide where to donate vouchers could lead individuals to feel more invested in election outcomes, raising their likelihood of voting. Another factor could be increased faith in the political system due to a reduced influence of large donations and wealthy donors on election outcomes. This improved view of the political system could increase political participation even among voters who do not use their vouchers. I next examine whether large donations became less influential.

The second main result is that Democracy Vouchers reduce candidates' reliance on large contributions by increasing small donations while decreasing large donations. This shift makes candidates more reliant on voucher donors and small cash donors for funding, which may make candidates more responsive to the desires of less wealthy voters. [55] show that small donors are more likely than large donors to be female or belong to an ethnic minority, suggesting that Democracy Vouchers may make politicians more responsive to these groups. Among candidates for city council, who face the lowest contribution limits when they accept vouchers and were the only eligible voucher recipients until 2021, I find that dollars from small contributions of \$100 and under increased by 156% and dollars from contributions of \$250 and under increased by 95%. Dollars from large contributions of more than \$100 decreased by 85%, dollars from contributions of more than \$250 decreased by 93%, and dollars from contributions between \$100 and \$250 decreased by 74%.

Third, I examine the program's effects on total campaign contributions and expenditures. I estimate increases in contributions and expenditures among all candidates, but decreases for city council candidates. Although the magnitudes of these estimates are fairly large, all estimated effects on total contributions and expenditures are statistically insignificant.

This paper contributes to the literature on determinants of political participation. My paper is the first to study the effects of public financing for campaigns on voter registration and turnout. Voter turnout has been found to increase in response to increased newspaper coverage [56], experiencing negative effects of a local policy [57], increased competitiveness of the election [58, 59], online voter registration [60], increased campaign spending in state-level elections [61] and compulsory schooling, which also increases voter registration [62]. Voter turnout decreases in response to broadband Internet access due to a decline in viewership of televised news [63] and is unaffected, both in the general population and among Black voters, by voter identification laws [64].

Furthermore, this paper contributes to the emerging literature that examines various outcomes of the Democracy Vouchers program. My paper is the first to study the program's effects on political participation. [65] examine demographic changes in the donor pool resulting from the program and find that when vouchers were available, the donor pool moved closer to being demographically representative of the electorate. By contrast, [66] finds that voucher availability made the donor pool less representative of the electorate. The most closely related paper is [67], which uses difference-in-differences to study Democracy Vouchers' effects on campaign finance outcomes. I perform my campaign finance analysis using the [54] method and my control group consists of other cities in King County, while their control group is composed of several large cities in Washington and California. I also include data from Seattle's 2021 election, while their sample ends in 2020. One advantage of my approach is that I am able to provide evidence of parallel pre-trends, while [67] do not have parallel pre-trends for their campaign finance outcomes. Both this paper and [67] find increases in small donations, but including 2021 data reduces the magnitude substantially: I estimate that 95-156% more dollars were given in small donations to city council candidates, compared with their estimate of 350%. While [67] estimate a 53% increase in total contributions per city council race, I estimate a 46% decrease in total contributions to city council candidates. Including data from an additional election is useful in determining whether patterns from the first two elections with Democracy Vouchers are persistent over time.

2.2 Background

Starting with the 2017 election, every registered voter in Seattle received four Democracy Vouchers in the mail nine months before election day. In each of the following months, vouchers were mailed to all voters who registered in the past month, with the last batch of vouchers mailed on October 1. A voter can donate a voucher by mailing it to an eligible campaign. Voters have discretion over whether to send all four vouchers to the same campaign and whether to use fewer than four vouchers. Vouchers can be returned at any point before the general election, and candidates can spend them either in the August primary election or in the November general election (Seattle Municipal Code 2021). Vouchers are only available for Seattle's municipal elections,[1] which are held every odd-numbered year.

Seattle's voters approved the Democracy Vouchers program in a 2015 ballot initiative, which includes 10 years' worth of funding from a property tax increase of $3 million per year. The tax affects both commercial and residential properties, and the average Seattle homeowner pays about $8.00 per year to fund the program (City of Seattle 2021a). The tax increase is not large enough to pay for voucher donations by all of Seattle's residents, likely due to the expectation that many voters will not use their vouchers. When a candidate reaches the expenditure limit, he or she must stop redeeming vouchers. These expenditure limits further reduce the probability that all vouchers will be redeemed. The value of unused vouchers remains in the program budget.

Seattle's municipal campaigns are non-partisan, with the two best-performing candidates from the primary election competing in the general election. Candidates are required to collect at least 150 signatures and at least 150 donations of $10 or more to qualify to receive vouchers. Voucher recipients must accept a lower individual contribu-

[1]Elections for city-level positions

tion limit, which is a limit on the amount of non-voucher money they will accept from any individual. Table 2.1 shows Seattle's individual contribution limits by year. Publicly available data on campaign contributions in Seattle[2] shows that before the Democracy Vouchers program, 54% of total campaign funding came from donations larger than $250; after the program was implemented, only 16% of campaign funding came from these large donations. Democracy Voucher recipients agree to campaign expenditure limits that range from $150,000 to $800,000, depending on the office (Seattle Municipal Code 2021). The same campaign contribution data shows that before Democracy Vouchers were available, 31% of Seattle's municipal campaigns exceeded the lowest expenditure limit of $150,000 and 1.8% spent more than the highest limit of $800,000. After vouchers became available, 26% of campaigns spent more than $150,000 and 1.5% spent more than $800,000.

Year	City Council & City Attorney Candidate Voucher Recipients	Mayoral Candidate Voucher Recipients	Candidates Not Accepting Vouchers
Before 2017	NA	NA	$700
2017	$250	NA	$500
2019	$250	NA	$500
2021	$300	$550	$600

Source: City of Seattle. Individual contribution limits are NA for mayoral candidates in 2017 and 2019 because they were not eligible to receive vouchers until 2021. The voucher program started in 2017.

Table 2.1: Contribution Limits

Data from the City of Seattle on Democracy Voucher contributions shows that participation in the voucher program has steadily increased over time, as is shown in Table 2.2. Only 3.4% of voters participated in the program in 2017, equating to $1,390,650 in total campaign contributions from Democracy Vouchers. Participation was higher in the 2019 election, with 7.7% of voters participating and candidates receiving $3,511,825. In the 2021 election, voucher donations reached $4,476,000, with 9% of registered voters

[2] https://www.pdc.wa.gov/political-disclosure-reporting-data/browse-search-data/contributions.

participating. [68] found anecdotal evidence after the 2017 election that many voters were aware of the program, but were not interested enough in municipal elections to prioritize researching candidates and mailing in vouchers before the deadline. Others assumed the vouchers were junk mail and recycled them.

Year	Number of Voucher Donors	Percent of Registered Voters	Total Voucher Amount
2017	15,353	3.36	$1,390,650
2019	36,694	7.71	$3,511,825
2021	42,775	8.96	$4,476,000

Source: Voucher use data provided by the City of Seattle. Percent of registered voters calculated using precinct-level data on the number of registered voters in King County, Washington.

Table 2.2: Voucher Use Statistics

Although voter participation has increased continuously, candidate participation declined in the 2021 election. The candidate participation rate in the voucher program increased from 44% in 2017 to 91% in 2019, but decreased back to 56% in 2021 [69]. Data on voucher contributions from the City of Seattle shows that candidates received an average of $114,260 in vouchers per campaign, which is equivalent to 93% of the average campaign's total funding. Democracy Vouchers therefore became a significant source of funding for municipal campaigns despite the low participation rate among voters. The remainder of the paper explores the program's effects on political participation and campaign finance.

2.3 Data

2.3.1 Sample and Time Frame

I use data from King County, the county in Washington where Seattle is located, for all outcomes. The other cities in King County are used as a counterfactual for Seattle.

King County is Washington's most populous county, containing one third of the state's population. Table 2.3 shows the populations of the control cities, as well as the city council size and term limits for each city. While Seattle has the largest population, the control group does not exclusively consist of small towns: five of Washington's ten largest cities are in King County [70]. All of the control cities have a mayor and a city council consisting of either five or seven members. City council and mayoral terms are typically four years, with a two-year city council term in one city and two-year mayoral terms in 11 cities.

Both Seattle and the control cities have unusually high levels of political engagement: before Democracy Vouchers were available, Seattle's average voter turnout was 52.9% and the average in King County's other cities was 46.7%, both greatly exceeding the national average of 27% turnout in local elections [71]. Local economic shocks to Seattle affect the entire county, as Seattle is the county's largest city and economic center [72]. Economic shocks affect the amount of disposable income available for cash donations, making the cities surrounding Seattle a useful control group.

Name	Population	Council Seats	Council Term	Mayor Term
Algona	3,107	5	4	2
Auburn	74,527	7	4	4
Beaux Arts Village	333	5	4	4
Bellevue	134,630	7	4	4
Black Diamond	4,291	7	4	4
Bothell	41,207	7	4	4
Burien	49,785	7	4	2
Carnation	1,733	5	2	2
Clyde Hill	3,150	5	4	4
Covington	18,689	7	4	4
Des Moines	30,715	7	4	4
Duvall	7,373	7	4	4
Enumclaw	11,494	7	4	4
Federal Way	92,859	7	4	4
Hunts Point	508	5	4	4
Issaquah	33,682	7	4	4
Kenmore	21,575	7	4	2
Kent	124,292	7	4	4
Kirkland	84,721	7	4	2
Lake Forest Park	13,059	7	4	4
Maple Valley	24,682	7	4	2
Medina	3,120	7	4	2
Mercer Island	24,120	7	4	2
Milton	7,219	7	4	4
Newcastle	10,994	7	4	2
Normandy Park	6,568	7	4	2
North Bend	6,305	7	4	4
Pacific	6,954	7	4	4
Redmond	57,959	7	4	4
Renton	97,234	7	4	4
Sammamish	50,163	7	4	4
SeaTac	27,859	7	4	4
Seattle	653,017	9	4	4
Shoreline	54,774	7	4	4
Skykomish	130	5	4	2
Snoqualmie	12,115	7	4	4
Tukwila	19,757	7	4	4
Woodinville	11,373	7	4	4
Yarrow Point	1,172	5	4	4

This table lists the cities in King County, Washington. Populations are from 2015 ACS 5-year estimates. Information about local government composition and term duration is from local government websites.

Table 2.3: Control Cities

I use data from the odd-numbered years from 2009 to 2021. I only use data from odd-numbered years because Seattle holds municipal elections in these years, and they are therefore the only years when Democracy Vouchers can be used. Federal elections and Washington's statewide elections are held in even-numbered years, and I exclude these elections because they garner much more attention, higher voter participation and more donations than local elections. King County holds county-wide elections in odd-numbered years, so voters in every city in my sample have an election in which they can participate in all years in my sample. However, seven of the 39 cities in King County do not hold municipal elections in odd-numbered years. Those seven cities are included in voter registration and turnout data, but not in campaign finance data.

2.3.2 Political Participation

To measure voter registration and turnout, I use precinct-level data from King County for the odd-numbered years from 2009 to 2021. King County election officials provided these data upon my request. For each precinct, I am able to observe the number of registered voters and the percent of registered voters who cast a ballot in each municipal election. King County has about 2,550 precincts in each election in the data, with the exact number ranging from a minimum of 2,514 in 2013 to a maximum of 2,611 in 2019. Each precinct has an average of 473 registered voters. The political participation data contains 17,890 observations, one for each precinct in each year. All registered voters in all precincts in Seattle received Democracy Vouchers in 2017, 2019 and 2021.

Elections in King County are coordinated at the county level. Voters in each city in the county elect county-level officials as well as city-specific officials. Each city is divided into precincts, which are the smallest level at which elections are organized. Precinct-level data provides the finest available geographic variation in political participation.

Analysis of voter registration and turnout is therefore performed at the precinct level. Precinct boundaries are re-drawn between elections, typically resulting in minor changes to a few precincts. I can always observe the city in which a precinct is located, and precinct boundaries are drawn so that they never span multiple cities. Therefore, the changes in precincts over time do not interfere with my analysis, but they do mean that I am unable to use precinct fixed effects when I estimate the vouchers' effects on voter registration and turnout.

2.3.3 Campaign Finance

I use publicly available data on campaign contributions in King County to measure the Democracy Vouchers' impact on campaign finance. These data are available online from the Washington Public Disclosure Commission.[3] The data include all donations to campaigns in King County beginning in 2007, but I exclude 2007 from my analysis because census tract-level demographic controls are unavailable for that year. I adjust the contributions for inflation by converting all contributions to 2019 dollars. I combine these data with voucher contribution data, which is publicly available from the Seattle Ethics and Election Commission.[4] When the outcome is a dollar amount, I always use 2019 dollars. However, in parts of my analysis I use restrictions such as "donations under $100", and for these restrictions I use the original amount rather than adjusting for inflation. I use the original amount for these restrictions because the value of each voter's Democracy Vouchers was not adjusted for inflation from 2017 to 2021, and because the distribution of donations shows that voters are always most likely to donate a nominal dollar amount ending in zero or five.

[3] Data are available here: https://www.pdc.wa.gov/political-disclosure-reporting-data/browse-search-data/contributions.

[4] Data are available here: https://www.seattle.gov/democracyvoucher/program-data/distributed-voucher-funds

I began with data on 864,453 campaign contributions. Restricting my sample to candidates for mayor, city council and city attorney in all cities removed 130,079 observations. Accounting for corrections removed 2,733 more observations.[5] I also removed data on the 27,315 Democracy Vouchers that were returned to the city blank or assigned to an invalid campaign. After these restrictions, I have a sample of 704,326 contributions. Each observation in these data is a contribution from an individual to a campaign. I also have campaign-level data on 944 campaigns from 2009-2021, in which each observation is one candidate's campaign in one election. The campaign-level data show the total amounts of contributions and expenditures for each campaign.

I aggregate all campaign finance data to the city-year level for my main analysis. I chose the city level over the campaign level because the Democracy Vouchers program could affect the number of candidates running for office (discussed in more detail in the Appendix). Campaign-level analysis could show a decrease in the average cost of a campaign even if the overall election became much more expensive because more candidates were running for office. Another option would be to perform the analysis at the race level. Race-level analysis was the focus of [67], and I complement their analysis by studying outcomes at the city level.

The campaign finance outcomes are total contributions, total expenditures, small contributions, and large contributions. I use two cutoffs to separate "small" and "large" contributions: \$100 and \$250. \$100 is chosen because it is the total amount that each individual receives in Democracy Vouchers, and \$250 is chosen because it was the most common individual contribution limit for voucher recipient candidates during the 2017-2021 period. A contribution of exactly the cutoff amount is counted as a small contribution. Previous literature has raised the question of whether to aggregate donations

[5] After reporting a contribution, campaigns can issue a correction, saying the amount initially reported was incorrect and providing the correct amount. I adjusted the contributions to reflect the corrected amount, rather than the original amount. I removed contributions that were corrected to zero.

at the donor-campaign-year level before assigning them to a size category [73]. I do not perform this aggregation for two reasons. First, multiple small contributions are a different choice by the donor than one large contribution. Second, 89% of non-voucher donors only make one contribution per campaign per election year. The choice of whether to aggregate in this case only affects the size category assignment of non-voucher donations, as individuals can only donate up to $100 in Democracy Vouchers in each election.

The small and large contribution outcomes are a city-year level measure of the total dollar amount of campaign contributions that came from donations under or over the cutoff, respectively. For example, if only four donations were made in Bellevue in 2009 in the amounts of $100, $100, $300 and $400, then Bellevue would have $900 in total contributions, $200 in small contributions and $700 in large contributions in 2009 regardless of whether the $100 or the $250 cutoff was used.

The control cities vary in size, as shown in Table 2.3, and certain small cities either did not report any contributions in some elections or did not have a contribution in every size category. To obtain a balanced panel at the city-year level for total contributions and expenditures, I used a value of zero for total contributions and expenditures if one of the 32 cities that has municipal elections in odd-numbered years did not report any contributions. For the remaining outcomes, which are the amount of contributions in a certain dollar amount category, I included each of the 32 cities in the control group if it ever reported a contribution in that category from 2009-2021. If the city is included in the control group for the outcome, I used a value of zero for any years when the city did not report contributions in the category in question. The rationale for allowing the control group to vary slightly by outcome is that if a city never any has donations over $250, it does not provide a useful counterfactual for donations over $250 in Seattle. The sample size in the campaign finance analysis varies slightly by outcome due to this variation in the control group.

Because the individual-level campaign contribution data includes the donor's name, I can use names to identify whether a voucher donor also gave a cash donation and whether the same individual donated in multiple years. Matching by name reveals that both cash and voucher donors were very unlikely to donate in more than one election from 2009-2021. The average cash donor gave a donation in 1.29 elections, and the average voucher donor gave a donation in 1.27 elections. I do not perform analysis of individual-level changes in donation behavior when Democracy Vouchers become available, because this matching exercise demonstrates that donating in multiple elections is a very rare behavior even among the 9.2% of the King County population that is politically engaged enough to donate in at least one of the seven elections I study. More detailed information about this matching exercise can be found in the Appendix.

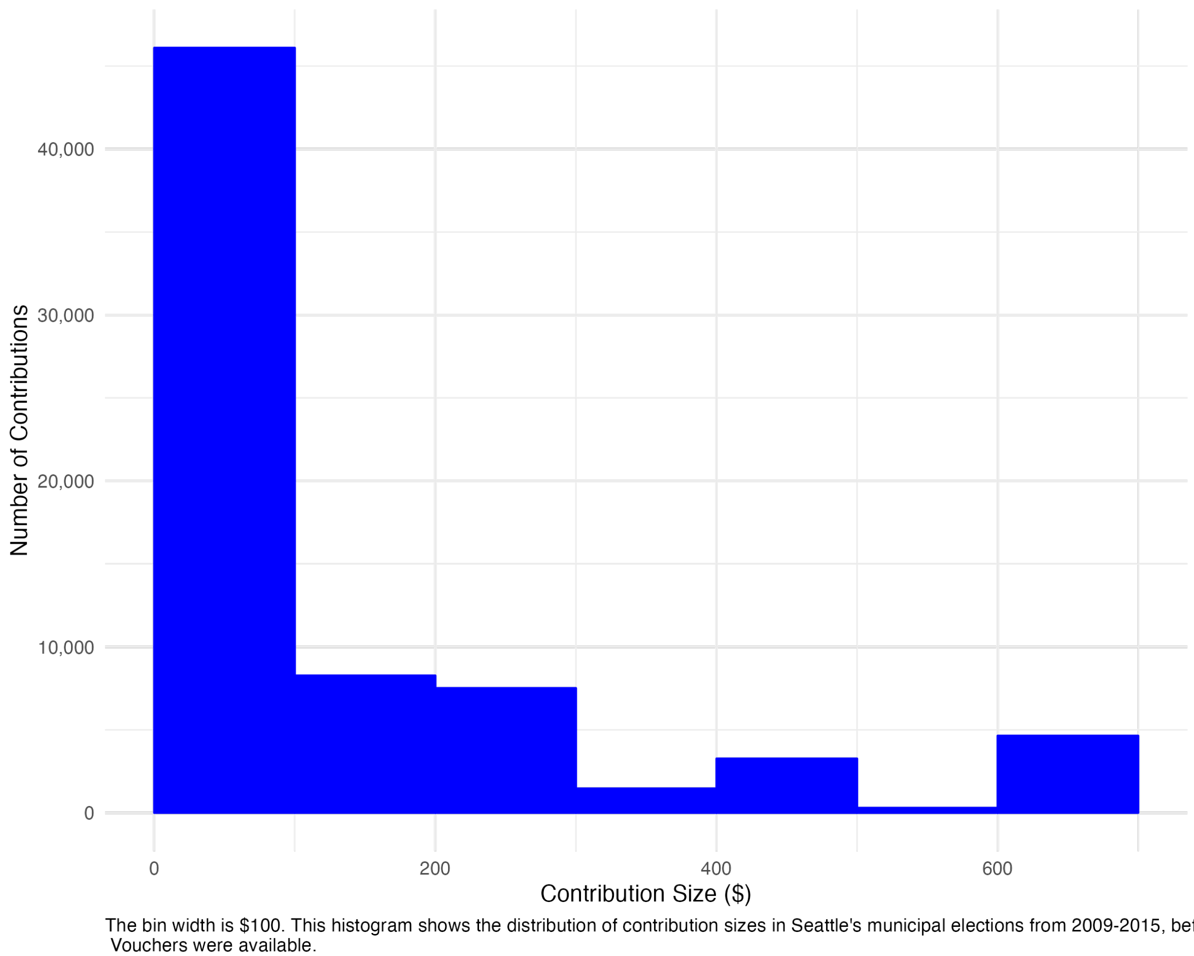

Figure 2.1: Histogram of Contributions Before Vouchers

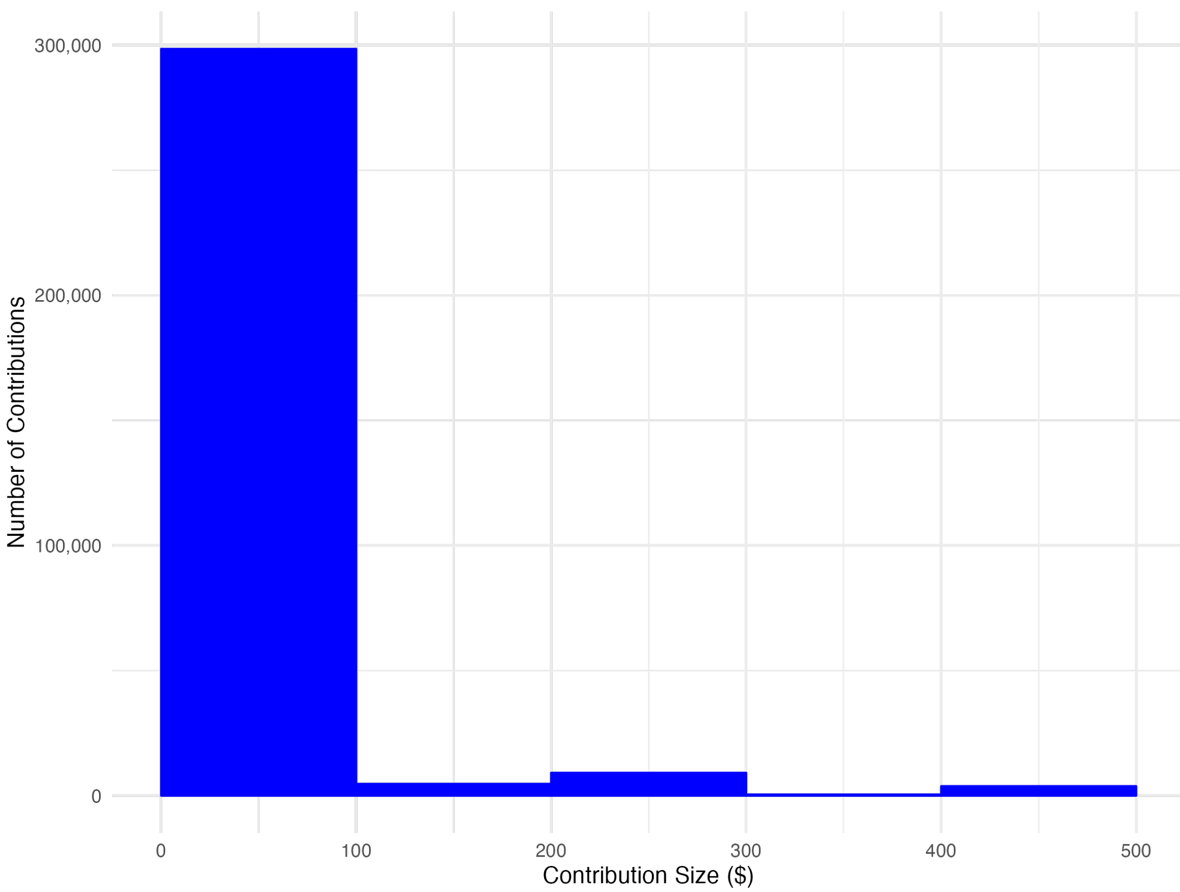

The bin width is $100. This histogram shows the distribution of contribution sizes in Seattle's municipal elections from 2017-2021, wh
Vouchers were available. For Democracy Voucher donations only, donations are aggregated at the donor-campaign-year level so th
gave all four of her vouchers to one campaign, that appears as one contribution in the histogram.

Figure 2.2: Histogram of Contributions After Vouchers

Figure 2.1 and Figure 2.2 are histograms showing the distribution of contribution sizes in Seattle before and after the Democracy Vouchers program. Each contribution is one contribution from one donor. In each histogram, the leftmost bar shows the number of contributions of $100 or less. The distribution of contributions became more left-skewed when Democracy Vouchers became available, with contributions in the smallest category of $100 or less increasing from 64% to 94% of all contributions. This shift in the distribution of contributions aligns with the program's goal.

2.3.4 Demographic Controls

I use census tract-level data from American Community Survey five-year estimates to obtain demographic characteristics of each precinct, including age, gender, race, income, education and total population. Census tract boundaries are different than precinct boundaries, allowing one tract to span multiple precincts. The demographic characteristics I use for each precinct are therefore weighted averages of the census tracts that overlap with that precinct, where the weights are determined by the percent of area that overlaps. For example, if 20% of Precinct 1 is in Census Tract A, 50% is in Census Tract B and 30% is in Census Tract C, then Precinct 1's demographics are a weighted average of Tracts A, B and C with a weight of 0.2 given to A, 0.5 given to B and 0.3 given to C. I use geographic area because I cannot observe the percent of the precinct's population that resides in each census tract. Table 2.4 compares city-level averages of demographic variables, as well as the number of registered voters and voter turnout, between Seattle and the other cities in King County. This table uses data from 2009 until 2015, the year of the last election before Democracy Vouchers became available. In Section 4, I describe the fixed effects and controls I include to ensure that my results are not driven by demographic differences between Seattle and the control cities.

	Control Cities		Seattle			
	Mean	SD	Mean	SD	Difference	SE
Average Contribution ($)	192.0	132.6	213.8	46.8	21.8	27.9
Median Age	37.8	3.7	36.0	0.1	-1.8	0.4
Percent Black	5.7	5.7	7.5	0.2	1.8	0.7
Percent White	69.6	11.3	70.6	0.9	1.0	1.4
Percent Asian	15.3	6.8	13.8	0.7	-1.5	0.8
Percent Pacific Islander	0.8	1.0	0.4	0.0	-0.4	0.1
Percent Native American	0.7	0.6	0.8	0.1	0.1	0.1
Percent Grade 8 or Less	3.3	2.8	3.1	0.2	-0.2	0.3
Percent Grade 9-12	5.4	3.2	4.1	0.5	-1.4	0.4
Percent High School Degree	19.2	8.4	12.6	1.1	-6.6	1.1
Percent Some College	22.4	4.7	21.1	0.5	-1.3	0.6
Percent Associate's Degree	8.1	1.7	6.8	0.1	-1.3	0.2
Percent Bachelor's Degree	26.7	10.3	32.5	1.1	5.8	1.3
Percent Graduate Degree	14.8	8.8	19.8	1.1	5.0	1.1
Average Income ($)	40,192	15,167.9	41,811.4	2,301.1	1,619.2	2,076.5

Control cities are all other cities in King County, WA besides Seattle. All summary statistics displayed are averages from 2009-2015. All variables are measured at the city-year level. Demographics are from the American Community Survey. Average contribution amount is measured using data on all contributions in King County.

Table 2.4: Summary Statistics

2.4 Methodology

I use the method that [54] developed for difference-in-differences with one treated group, recognizing that the standard cluster-robust variance estimator does not perform well in the case of only one treated group [54, 74, 75]. The [54] method is robust to heteroskedasticity, while the two other methods developed for difference-in-differences with one treated group [74, 75] rely on homoskedasticity assumptions that would cause them to under-reject the null hypothesis in my setting [54]. The wild bootstrap, often used when the number of clusters is small, can perform well with few treated groups except in the case of only one treated group [76, 77].

With one treated group, the standard difference-in-differences estimator is equal to the difference-in-differences estimand plus a second difference in differences: the difference in the pre-post difference in average errors between the treated and control groups [74]. As the number of control groups grows large, the pre-post difference in average errors for the control groups shrinks to zero. However, with only one treated group, the pre-post difference in average errors for the treated group does not disappear, making the difference-in-differences estimator inconsistent. Obtaining a consistent estimator of the difference-in-differences estimand in this setting requires using information from the control group to estimate the pre-post difference in average errors for the treatment group [74]. [54] allow for heteroskedasticity when they estimate the average errors of the treatment group. They use estimated heteroskedasticity to rescale the pre-post difference in the control groups' average residuals, making the control groups' average residuals informative about the treated group's pre-post difference in average errors [54].

2.4.1 Political Participation

For voter registration and turnout, my specification is:

$$Y_{pct} = \beta_0 + \beta_1 Seattle_{pc} * Voucher_t + \gamma_c + \psi_t + \theta X_{pct} + \epsilon_{pct} \tag{2.1}$$

Y_{pct} is voter registration or voter turnout in precinct p in city c in election year t . $Seattle_{pc}$ is an indicator for whether Seattle is the city in which precinct p is located. $Voucher_t$ is an indicator for whether Democracy Vouchers are available in year t . X_{pct} is a vector of precinct-level demographic controls, including race, gender, age, education, income, and total precinct population. β_1 is the coefficient of interest. I use city and time fixed effects, and standard errors are clustered at the city level.

2.4.2 Campaign Finance

For all campaign finance outcomes, the regression specification is:

$$ln(Y_{ct} + 1) = \beta_0 + \beta_1 Seattle_c * Voucher_t + \gamma_c + \psi_t + \theta X_{ct} + \epsilon_{ct} \quad (2.2)$$

Y_{ct} is the outcome in city c in election year t . $Seattle_c$ is an indicator for whether city c is Seattle. $Voucher_t$ is an indicator for whether Democracy Vouchers were available in Seattle in election year t . X_{ct} is a vector of city demographic controls, including race, gender, age, education, income, and total city population. β_1 is the coefficient of interest. I include city and year fixed effects, and I cluster at the city level.

To measure the vouchers' effects on campaign finance, I use four outcomes: total campaign contributions, total campaign expenditures, amount of small contributions, and amount of large contributions. Each of these outcomes is a dollar amount measured at the city-year level. I perform the campaign finance analyses for my entire sample, as well as for a sample restricted to candidates for city council. The vast majority of campaigns in the data (88%) are for city council, and city council candidates were the only candidates to receive vouchers in the 2017 and 2019 elections. Although mayoral and city attorney candidates did not receive vouchers in those elections, I include them in the full-sample analysis because mayoral candidates did receive vouchers starting in 2021 and because lower individual contribution limits for city council candidates could potentially have spillover effects on contributions to mayoral and city attorney candidates in 2017 and 2019. 11% of campaigns in my sample are for mayor and 1% are for city attorney. Figure 2.3 shows that 64% of total campaign funding went to city council candidates, 31% went to mayoral candidates and 5% went to city attorney candidates.

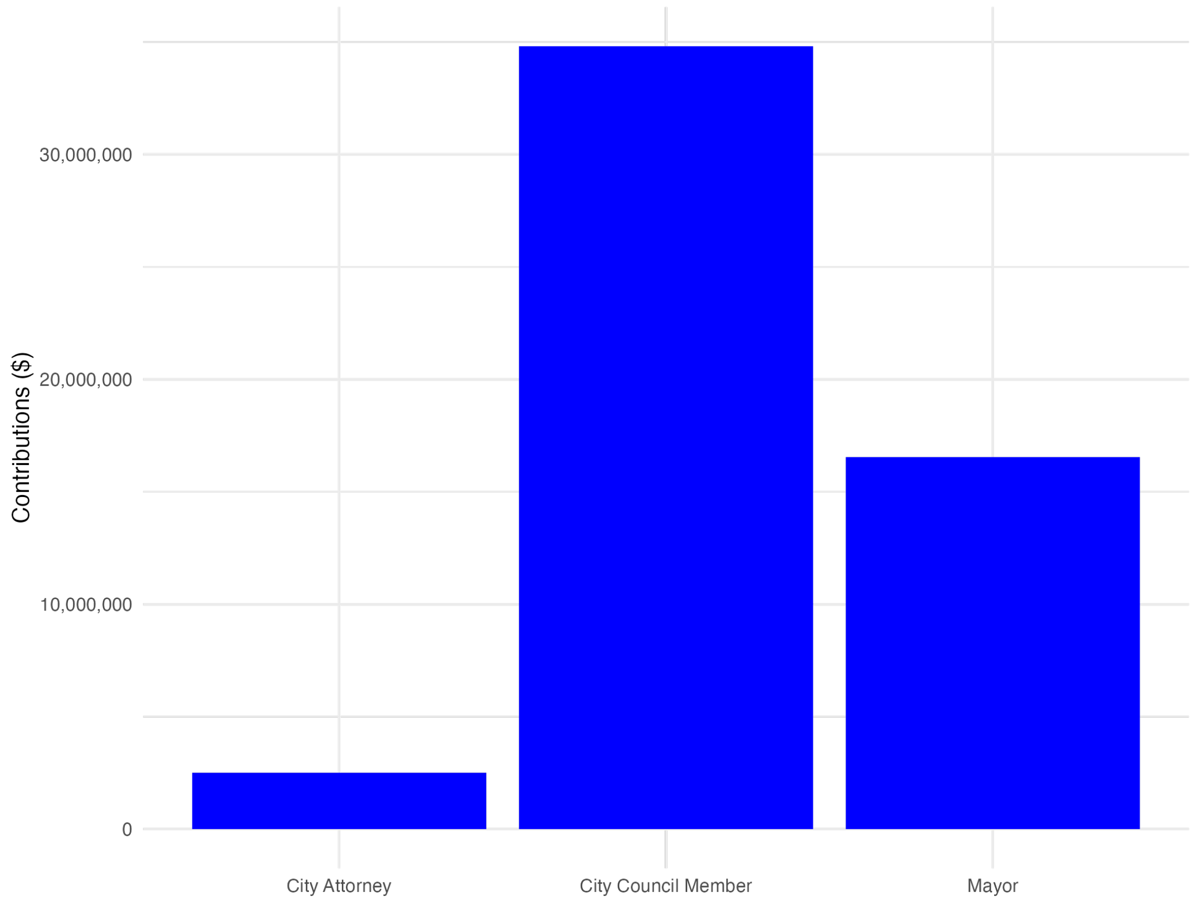

This graph shows the total amount of money donated to candidates for each municipal office from 2009-2021.

Figure 2.3: Contributions by Political Office

I log the financial dependent variables due to a skewed distribution of residuals when the dependent variables are not logged. I do not perform the campaign finance analysis at the precinct level because donations to any campaign can come from anywhere in the city.

Because individuals can donate to a candidate outside their home city, spillover effects pose a potential threat to identification for campaign finance outcomes. This problem would arise if the contribution limits placed on Democracy Voucher recipient candidates motivated Seattle residents to donate to campaigns outside their city. An investigation of the extent of spillover effects shows that the percent of all dollars contributed that were outflows from Seattle has always been small and remained fairly stable from 2009-2021, only slightly decreasing from 3.1% to 3% when vouchers became available. The rarity and stability of these spillovers suggests that they are unlikely to be a substantial driver of the main results.

2.4.3 Event Studies

I use event studies to assess the plausibility of the parallel trends assumption for difference-in-differences. The event study regressions have the following form:

$$Y_{ct} = \beta_0 + \sum_{\substack{j=2009 \\ j \neq 2015}}^{2021} \beta_j Seattle_c \mathbf{1}(t = j) + \psi_c + \gamma_t + \theta X_{ct} + \epsilon_{ct} \tag{2.3}$$

Y_{ct} is the outcome in city c in year t . γ_t is a year fixed effect, and ψ_c is a city fixed effect. Logs and precinct-level analysis are used for an outcome if they were used in the difference-in-differences regression, so that each event study matches the corresponding difference-in-differences regression. 2015 is the omitted year because the last election without Democracy Vouchers was held in that year. I use the [54] method to obtain p-values for the event study, and I calculate confidence intervals using those p-values.

2.5 Results

I find that the Democracy Vouchers program increases political participation. Estimated effects on both voter turnout and voter registration are positive, although only the estimated effect on voter turnout is statistically significant. Figure 2.4 provides evidence of similar pre-trends for voter registration and turnout in Seattle and the control cities. Analysis of campaign finance outcomes shows that small contributions, defined as either contributions $100 and under or $250 and under, increased while contributions over those cutoffs decreased. I do not find evidence of statistically significant changes in total contributions and expenditures.[6] Figure 2.5 and Figure 2.6 display event studies for the campaign finance outcomes, which show evidence of similar pre-trends in Seattle and the control cities.

The event studies suggest that the parallel trends assumption necessary for identification is plausible. Event studies for the political participation outcomes are better-powered to detect violations of the parallel trends assumption than event studies for campaign finance outcomes due to larger sample sizes and a larger proportion of treated units for the political participation outcomes. For the main results, these same features make statistical significance a more demanding test for the campaign finance outcomes than for the political participation outcomes.

[6]The synthetic control approaches from [78] and [79] produce estimates that are consistent with the main results discussed in this section. Synthetic control results are not included in the main text because placebo tests suggest that they are unreliable for causal inference, likely due to the short pre-treatment period. These results are available in the Online Appendix.

Outcome	DiD Coefficient	Ferman & Pinto P-Value	Sample Size	DiD Coefficient (Percent)	Seattle Pre-Treatment Average
Political Participation					
Voter Turnout	4.904	.085	17,890	9%	53
Voter Registration	23.393	.2822	17,890	6%	412
Small Contributions					
Under $100	1.411	<0.0001	217	310%	$786,059
Under $100 (City Council)	0.942	.0703	203	156%	$556,453
Under $250	1.269	<0.0001	217	256%	$1,634,864
Under $250 (City Council)	0.667	.1655	203	95%	$1,100,594
Large Contributions					
Over $100	-0.176	.7339	224	-16%	$2,825,903
Over $100 (City Council)	-1.906	<0.0001	210	-85%	$1,764,798
Over $250	-0.453	.4305	224	-36%	$1,977,098
Over $250 (City Council)	-2.677	<0.0001	210	-93%	$1,220,657
Between $100 and $250	-0.047	.8574	217	-5%	$407,718
Between $100 and $250 (City Council)	-1.357	.0354	203	-74%	$266,474
Total Contributions and Expenditures					
Total Contributions	0.338	.4133	224	40%	$3,637,067
Total Contributions (City Council)	-0.624	.247	210	-46%	$2,348,940
Total Expenditures	0.451	.3218	224	57%	$3,691,085
Total Expenditures (City Council)	-0.783	.1731	210	-54%	$2,395,427

All dependent variables are logged except voter registration and voter turnout. Baseline averages are in original units. P-values are calculated using the method in Ferman and Pinto (2019). Demographic controls are included in all regressions. The unit of analysis is the city for all outcomes except voter registration and turnout, for which the unit of analysis is the precinct. Clustering is at the city level.

Table 2.5: Main Results

2.5.1 Political Participation

I use two measures of political participation: voter registration and voter turnout. Only registered voters receive Democracy Vouchers, creating an incentive to register to vote for individuals who would like to donate vouchers. Voter turnout could increase if either the decrease in large donations or the process of donating vouchers raises an individual's enthusiasm for her preferred candidate, increasing her likelihood of voting. Democracy Vouchers could improve political participation by reducing cynicism about the political system and its responsiveness to only the desires of wealthy donors.

Table 2.5 displays the estimated effects on political participation. Voter turnout, measured as the percent of registered voters who vote in an election, increases by 4.9 percentage points. This estimate is statistically significant at the 10% level and represents a 9% increase from the pre-treatment mean of 53%. This result suggests that public financing for political campaigns can increase political engagement. I estimate that the Democracy Vouchers program increases voter registration by 23 voters per precinct, a 6% increase from the pre-treatment mean of 412. However, this estimate is statistically insignificant.

While increasing political participation was not an official goal of the voucher program, it is a beneficial unintended consequence of the program. Democracies are only responsive to the desires of people who vote, and an increase in voter turnout leads to political outcomes that reflect the preferences of a larger subset of the population. The particular form of public financing may be an important determinant of its effect on political participation: because Democracy Vouchers are mailed to all registered voters and donating them is an interactive experience, they are salient. The process of researching candidates and choosing where to send vouchers could lead to individuals becoming more invested in the outcome of the election than they would be in an alternative public financing system that directly provides government grants to candidates.

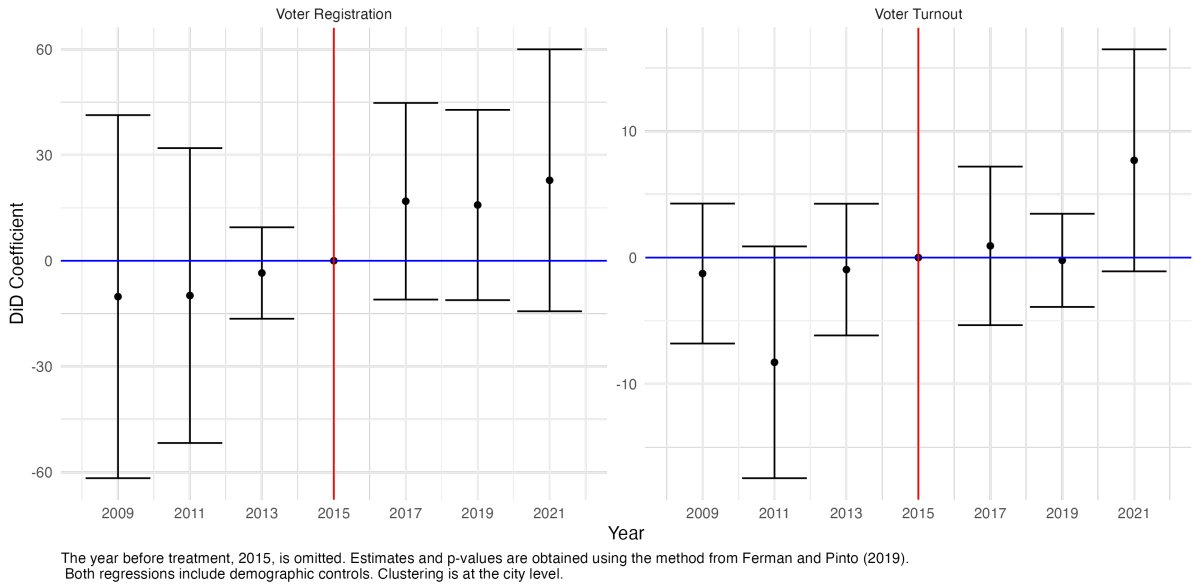

Figure 2.4: Event Studies: Political Participation

2.5.2 Small Contributions

Because the Democracy Vouchers program restricts the size of contributions for program participants and provides a new source of small contributions, it may affect the amounts of campaign contributions that come in the form of small contributions and large contributions. I first define small contributions as contributions less than or equal to $100, the value of each voter's Democracy Vouchers. This category includes all Democracy Voucher donations as well as non-voucher donations of $100 or less. I replicate my analysis with $250 as the cutoff for a small donation. I chose $250 as the cutoff because $250 was the maximum individual donation to a city council candidate who received Democracy Vouchers in the first two years of the program. This cutoff increased in 2021. Results using the 2021 cutoff are similar to the main results and are discussed in detail in the Robustness Checks section.

I find that the Democracy Vouchers program increases the amount of funding from donations of $100 and under, as well as the amount of funding from donations of $250 and under. This finding aligns with the descriptive statistic that the average contribution in Seattle decreased by 77% under the program, from $201.35 before 2017 to $45.95 in the 2017, 2019 and 2021 elections.[7] The increase in small donations is consistent with the program's goal of making candidates more reliant on small rather than large contributions.

Table **??** displays my results for these outcomes. First, I examine the vouchers' effects on contributions of $100 and under and find an increase of 310%. This estimate is significant at the 1% level, and Seattle's pre-treatment average was $786,059. When I restrict my analysis to candidates for city council, I estimate an increase of 156% that is significant at the 10% level. Seattle's pre-treatment average for city council candidates was $556,453.

Second, I estimate the vouchers' effects on contributions of $250 and under. I estimate that Democracy Vouchers increased these contributions by 256%. This estimate is statistically significant at the 1% level, and the pre-treatment average is $1,634,864. This result is informative about the behavior of donors who would have given $250 or less in the absence of vouchers. If these donors substituted vouchers for their cash donation, giving $100 in vouchers and $150 in cash, I would not see a change in the amount of dollars from contributions of $250 or less. Therefore, this result provides evidence that Democracy Vouchers did more than crowd out cash donations. When I restrict this analysis to city council candidates, the pre-treatment average is $1,100,594, and I estimate a statistically insignificant increase of 95%.

Together, the results suggest that most of the increase in small contributions came from an increase in donations of $100 or less. I find the strongest evidence of an increase in small contributions when I use $100 as the cutoff. The new small donations include

[7] A contribution is defined here as one contribution from one donor. Note that it is possible for one donor to make multiple contributions.

Democracy Vouchers as well as any cash donations of $100 or less that were inspired by the program.

Although most voucher recipients were city council candidates, estimated effects are smaller when the sample is restricted to candidates for city council, likely for two reasons. First, mayoral and city attorney candidates became eligible to receive vouchers in 2021, and the resulting increases in voucher and small cash contributions are excluded from the city council-only sample. Second, individuals who received Democracy Vouchers in 2017 and 2019 may have become more likely to make small contributions to all municipal candidates in those years, including those who were ineligible for vouchers. Restricting the sample to city council therefore leads to an underestimate of the program's total effect on small contributions.

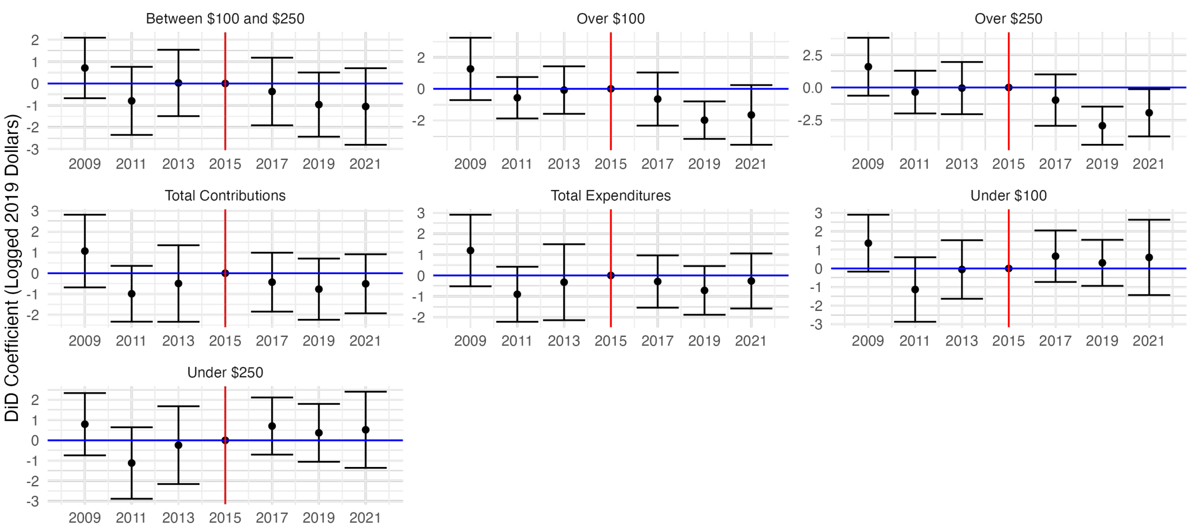

All outcomes are logged. The year before treatment, 2015, is omitted. Estimates and p-values are obtained using the method from Ferman and Pinto (2019).
All regressions include demographic controls. Clustering is at the city level.

Figure 2.5: Event Studies: Campaign Finance

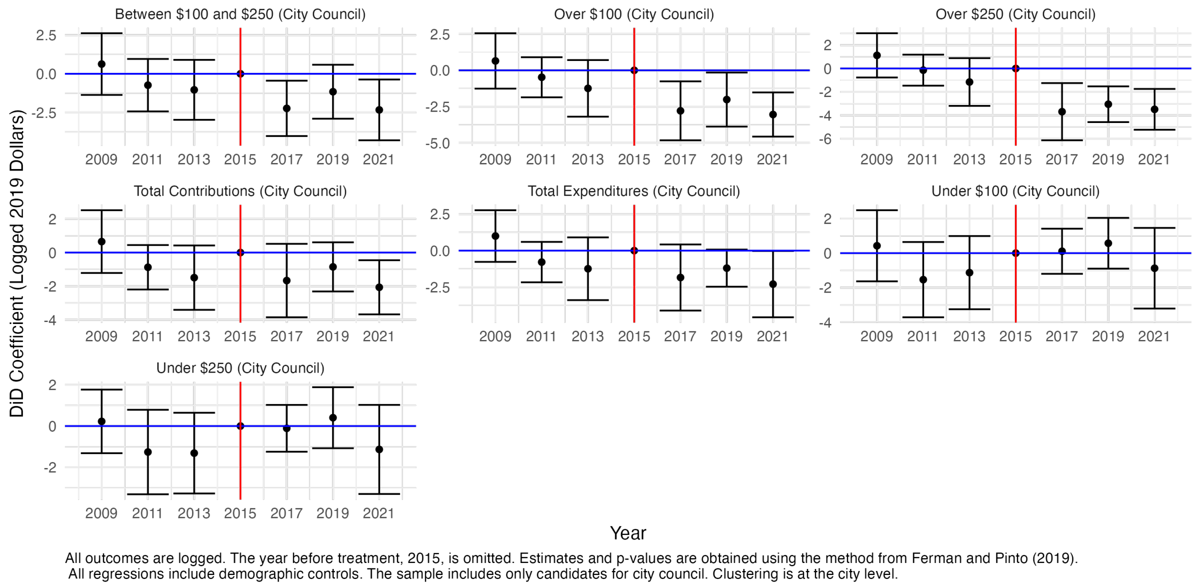

Figure 2.6: Event Studies: Campaign Finance (City Council Only)

2.5.3 Large Contributions

I find evidence that the amount of campaign funding coming from large donations decreases when Democracy Vouchers are available, which is consistent with the program's goals. I perform this analysis with both $250 and $100 as the cutoff for a large donation. I also estimate the change in contributions between $100 and $250. The results appear in Table 2.5.

I estimate a 36% decrease in contributions of more than $250. Seattle's pre-treatment average was $1,977,098. This estimate is statistically insignificant. For contributions of more than $100, I estimate a statistically insignificant decrease of 16%. When I restrict my sample to candidates for city council, I find a 93% decrease in contributions over $250 and an 85% decrease in contributions over $100. Both estimates are statistically significant at the 1% level. Pre-treatment averages were $1,220,657 and $1,764,798 respectively. Although point estimates are negative for the results that include all candidates, the estimates for city council candidates are larger in magnitude and statistically significant.

Stronger evidence of a decrease in large contributions to city council candidates than to all candidates likely has two causes. First, city council candidates were the only candidates eligible to receive vouchers in two of the three elections with vouchers and were therefore the only candidates subject to lower contribution limits in those elections. Second, mayoral candidates had a larger individual contribution limit of $550 if they accepted vouchers when they became eligible in 2021. An analysis of mayoral candidates only shows that contributions over $550 decreased when mayoral candidates became eligible to receive vouchers in 2021, as shown in Table 2.8.

The decrease in contributions over $250 provides insight into the type of city council candidate that chose to accept Democracy Vouchers: if the only candidates who accepted vouchers were those who would not have received any contributions over $250 in the absence of the program, then I would not find a decrease in contributions over $250. Candidates who expected to receive contributions over $250 if they chose not to accept vouchers decided that choosing to accept vouchers would maximize their total funds despite the decrease in the individual contribution limit, showing that they expected a compensating increase in small donations if they chose to accept vouchers and accepted the $250 contribution limit.

I also examine changes in donations between $100 and $250. I estimate a statistically insignificant decrease of 5% for all candidates and a 74% decrease for city council candidates, which is significant at the 5% level. The pre-treatment average for city council candidates in Seattle was $266,474. Contributions under $250 are unaffected by the lower contribution limit for candidates who accept Democracy Vouchers, and therefore a decrease in contributions of this size is not driven by the $250 individual contribution limit binding. The decrease in contributions between $100 and $250 is therefore evidence that for some donors, optimal cash contributions decrease when vouchers become available. This group of donors may experience a “warm glow” from donating vouchers that replaces the “warm glow” they could otherwise gain from donating cash. The total decreases in donations over $100 and donations over $250 are likely due to a combination of donors for whom the optimal cash donation decreases and donors whose optimal cash donation exceeds the $250 limit, but who are restricted by the limit.

One potential concern about these results is that the decreases in contributions over $100 and contributions over $250 may be driven by the decrease in the cap on campaign contributions from individuals from $700 to $500 for non-voucher recipients that occurred when the Democracy Vouchers program began. I use robustness checks to examine this possibility and conclude that the decrease is at least partially driven by the Democracy Vouchers program, and does not solely result from the decrease in the contribution cap.

Another concern is that wealthy donors may have substituted their contributions to Political Action Committees (PACs), which are independent entities that can spend unlimited amounts as long as they do not directly coordinate with a candidate’s campaign.[8] Although King County provides partial data on PAC spending, the data does not include sufficient detail to analyze whether changes in PAC spending resulted from the Democracy Vouchers program.

2.5.4 Total Campaign Contributions and Expenditures

I examine the Democracy Vouchers’ effects on total campaign contributions and expenditures to determine how the program affected the total cost of elections. These two outcomes are closely linked but are not quite equal, as candidates do not always spend the full amount they receive in contributions. Table 2.5 displays the results for total

[8] Traditional PACs are allowed to coordinate with a candidate’s campaign and face strict contribution limits, while the 2010 *Citizens United v. Federal Election Commission* Supreme Court ruling enabled “Super PACs” to receive unlimited contributions as long as they do not coordinate with the candidate’s campaign.

campaign contributions and expenditures.

For all candidates, I estimate a 40% increase in total contributions and a 57% increase in total expenditures. By contrast, in a sample restricted to only city council candidates I estimate a 46% decrease in total contributions and a 54% decrease in total expenditures. Although the estimates are large in magnitude, none are statistically significant. The difference in the signs of the estimates for the full sample and city council candidates is consistent with the evidence above that large contributions decreased more for city council candidates than for the full sample of candidates. The larger decrease in funding from large contributions for city council candidates may result from those candidates being eligible for vouchers in more elections than mayoral or city attorney candidates. It is possible that if all candidates had been eligible for vouchers in all three elections from 2017 to 2021, the decrease in large contributions would have outweighed the increase in small contributions to make elections less expensive for all candidates.

The estimated effect on total contributions can be used to calculate the extent of crowding out. The 46% decrease in total contributions to city council candidates implies crowding out of 60%, which means that each dollar donated in Democracy Vouchers reduced private contributions by $0.60. This estimate is larger than the magnitude found in [67]. However, the estimated effect on total contributions is imprecise, which is demonstrated by the statistical insignificance of estimates with large magnitudes. The 60% estimate should therefore not be taken as a precise estimate of crowding out.

2.6 Robustness Checks

2.6.1 Permutation Tests

An alternative to the [54] approach to difference-in-differences with one treated group is the permutation test method used in [80] and [81]. This approach repeats the difference-in-differences regression for each city in the control group, replacing the Seattle indicator with an indicator for the control city. Then the estimate from the main difference-in-differences regression is compared to the placebo estimates to determine statistical significance. Because the method from [54] uses a different method than standard difference-in-differences techniques to calculate both point estimates and p-values, the permutation test here functions as a robustness check for both the point estimates and levels of statistical significance in the main results.

Both [81] and [80] used fifty control units, meaning that 5% significance from a two-tailed test could be achieved only if their treated unit was ranked at the top or bottom of the distribution of placebo estimates. In my setting, with only 38 control units for political participation outcomes and 31 control units for campaign finance outcomes, 5% significance from a two-tailed test cannot be achieved even if the estimate for Seattle is ranked at the top or bottom of the distribution, and 10% significance can only be achieved if Seattle is ranked at the top or bottom.

Outcome	Coefficient	P-Value
$100 and Under	1.158	0.581
Over $100	-0.100	0.938
$250 and Under	1.102	0.323
Over $250	-0.348	0.875
Between $100 and $250	-0.074	0.903
Total Contributions	0.530	0.812
Total Expenditures	0.491	0.812
Voter Registration	17.456	0.450
Voter Turnout	7.233	0.100

All dependent variables are logged except voter registration and voter turnout. P-values are calculated using permutation tests. Demographic controls are included in all regressions. Clustering is at the city level.

Table 2.6: Permutation Test Results

Table 2.6 displays the results of the permutation tests. While all estimates have the same directions as the main results and similar magnitudes, only the estimated effect on voter turnout achieves statistical significance at the 10% level. The point estimate is a 7.2 percentage point increase, larger than the effect found with the [54]. The lack of statistical significance for other outcomes is unsurprising due to the size of the control group, as this test is considered very demanding even with a larger number of control units [81]. Finding statistical significance at the 10% level for voter turnout with this demanding test provides further evidence that the Democracy Vouchers program increased political participation.

2.6.2 Controls for Number of Races

Outcome	DiD Coefficient	Ferman & Pinto P-Value
Voter Turnout	4.670	0.085
Voter Registration	23.327	0.27
Total Expenditures (City Council)	-0.074	0.77
Total Expenditures	0.005	0.983
Total Contributions (City Council)	0.096	0.706
Total Contributions	-0.089	0.73
Over $250 (City Council)	-2.641	<0.0001
Over $250	-1.683	<0.0001
Over $100 (City Council)	-1.602	0.013
Over $100	-1.198	0.001
Between $100 and $250 (City Council)	-0.799	0.061
Between $100 and $250	-0.728	0.13
Under $250 (City Council)	1.716	0.004
Under $250	0.960	0.042
Under $100 (City Council)	2.120	<0.0001
Under $100	1.225	<0.0001

All dependent variables are logged except voter registration and voter turnout. P-values are calculated using the method in Ferman and Pinto (2019). Demographic controls, a control for number of city council races, and a control for whether a mayoral election was held are included in all regressions. Clustering is at the city level.

Table 2.7: Control for Number of Races

One potential concern is that changes in campaign finance outcomes could result from a change in the number of city council races or whether a mayoral election is being held in a given year. To address this concern, I include controls for the number of city council and mayoral elections being held in each city in each year. Table 2.7 displays the results. Estimates and statistical significance for the political participation outcomes are nearly identical to the main results. The signs of all large and small contribution estimates, such as estimates under $100, remain the same. The directions of the estimated effects on total contributions and total expenditures change when these controls are included, but the estimates remain statistically insignificant. Overall, the results of this robustness check are similar to the main results. Event studies including these controls, which appear nearly identical to the main event studies, can be found in the Appendix.

2.6.3 2021 Changes to the Program

Two significant changes were made to the Democracy Vouchers program in 2021. First, mayoral candidates became eligible to receive Democracy Vouchers for the first time. Mayoral candidates who accepted vouchers were subject to an individual contribution limit of \$550. Second, the individual contribution limit for city council candidates who received vouchers increased from \$250 to \$300.

Because the program in 2021 differed from the program in 2017 and 2019, I exclude the years 2017 and 2019 to analyze the difference between the 2021 election and the elections before Democracy Vouchers were available, using the 2021 cutoffs to define small and large contributions. As with the main results, I find increases in small contributions and decreases in large contributions.

Outcome	DiD Estimate	Ferman & Pinto P-Value	Sample Size
Mayor: Contributions Over \$550	-2.666	0.001	105
Mayor: Contributions \$550 and Under	0.976	0.012	98
City Council: Contributions Over \$300	-1.923	0.050	210
City Council: Contributions \$300 and Under	0.699	0.325	203

All dependent variables are logged. The dependent variable for each regression is the logged total dollar amount over or under the specified cutoff, measured at the city-year level. Only the years 2009-2015 and 2021 are included. These regressions use the cutoffs for Democracy Voucher recipient candidates that were introduced in the 2021 election. Demographic controls are included in all regressions. Clustering is at the city level.

Table 2.8: Before 2017 vs. 2021

Table 2.8 shows the results. For city council candidates, I estimate an 85% decrease in contributions over \$300 and a 101% increase in contributions of \$300 and under. The estimated decrease in large contributions is statistically significant at the 5% level, while the estimated increase in small contributions is not. For mayoral candidates, contributions over $550 are estimated to decrease by 93% and contributions of $550 and under are estimated to increase by 165%. Both estimates are significant at the 1% level.

2.6.4 Maximum Contribution Limit

Because the law that implemented the Democracy Vouchers program also lowered the maximum contribution amount from $700 to $500, I expect that the reduction in the amount of contributions over $250 is partially driven by the decrease in this cap.

To determine whether this cap reduction is the primary mechanism reducing the amount of contributions over $250, I first use the number of contributions over $250 as my outcome, rather than the dollar amount. Table 2.9 displays the result. I use this outcome because individuals switching from a donation of more than $500 to a donation of $500 in response to the new contribution cap would not lower the number of contributions over $250. A decrease in the number of contributions over $250 should be driven by the Democracy Vouchers. I estimate a 24% decrease, which is statistically significant at the 1% level. Evidence of parallel trends is weaker for this outcome than for the main results, as can be seen in Figure **??**.

Outcome	DiD Coefficient	Ferman & Pinto P-Value	N	DiD Coefficient (Percent)	Seattle Pre-Treatment Average
Between $250 and $500	-0.657	.1633	210	-48%	$279,099
Number of Campaigns	-1.536	.1898	224	-6%	28
Number of Donors	38,140.483	<0.0001	224	315%	12,094
Number Over $250	-641.099	<0.0001	224	-24%	2,710
Over $250 ($500 Topcoded)	-0.565	.3253	224	-43%	$1,148,263

Note:
Dependent variables are logged unless they start with 'Number'. Baseline averages are in original units.
P-values are calculated using the method in Ferman and Pinto (2019).
Demographic controls are included in all regressions. The unit of analysis is the city. Clustering is at the city level.

Table 2.9: Additional Robustness Checks

Second, I use the dollar amount of contributions between $250 and $500 as an outcome. The result appears in Table 2.9. I expect that individuals who prefer to donate more than $500 donated exactly $500 when the new contribution cap was imposed. Individuals who prefer to make donations between $250 and $500 (not including individuals who donate exactly $500) should not be affected by the overall contribution cap being lowered from $700 to $500. Therefore, if I find a smaller amount of donations between $250 and $500, I can attribute that effect to the Democracy Vouchers. I estimate a 48% decrease. Although the magnitude of this estimate is fairly large and has the expected direction, it is imprecisely estimated and statistically insignificant.

Third, I use the amount of contributions over $250 as my outcome, replacing all contributions over $500 in the data with $500. The result appears in Table 2.9. Individuals who donated more than $500 prior to 2017 would most likely have donated $500 if the contribution limit had not decreased. Coding these donations as $500 allows me to test whether the decrease in the amount of donations over $500 was driven by the decrease in the contribution cap. When I perform this robustness check, I find a statistically insignificant 43% decrease. This result is similar in magnitude to the one I found without top-coding.

Overall, the results of these robustness checks are consistent with the conclusion that the decrease in the amount of contributions over $250 was at least partially driven by the Democracy Vouchers, and was not solely the result of the decrease in the cap on contributions to non-voucher recipients. Finding positive results or precisely estimated null results for these robustness checks would have cast doubt on the conclusion that large donations decreased due to the Democracy Vouchers program.

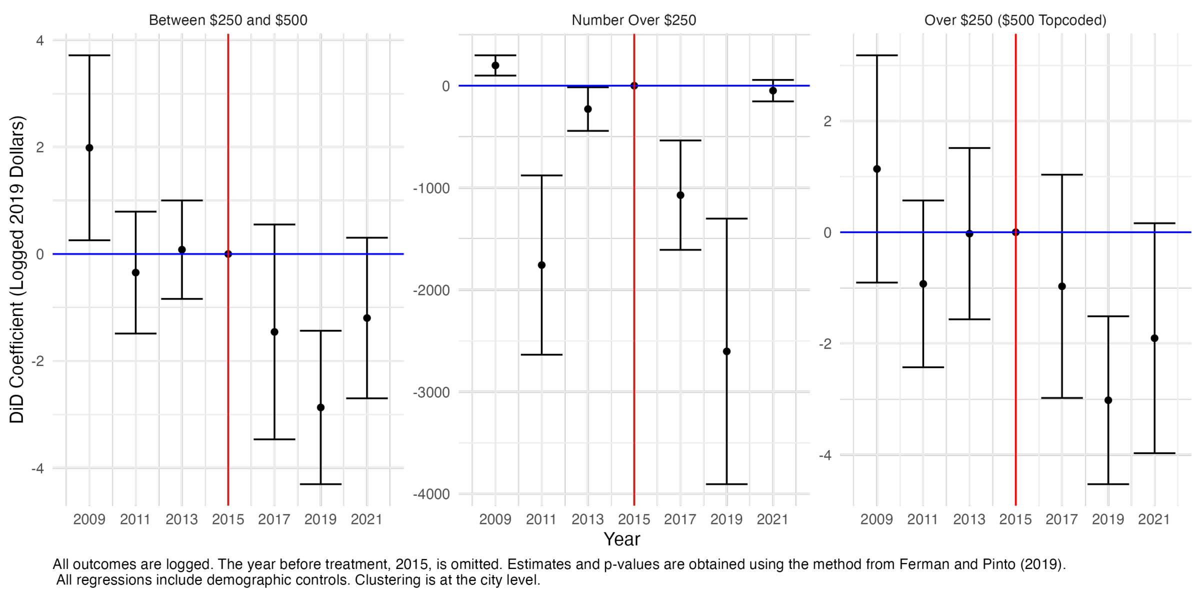

Figure 2.7: Event Studies for Robustness Checks

2.7 Conclusion

Unlike most forms of public financing, the Democracy Vouchers program allows individuals to allocate government funding that can supplement or replace their cash donations. By limiting candidates' reliance on large donations, the program could increase political participation if voters believe candidates are becoming more responsive to the average voter rather than wealthy donors. Political participation could also increase if the process of donating vouchers leads voters to feel more invested in the outcome of an election. I study the program's effects on voter registration, voter turnout and campaign contributions.

My results show that the Democracy Vouchers program increased political participation, raising voter turnout by 4.9 percentage points. This result is a favorable unintended consequence of the policy, as increasing voter turnout was not one of its explicit goals. To determine whether candidates became less reliant on large contributions, which could be a factor driving the increase in political participation, I examine changes in the composition of contributions. I find that the Democracy Vouchers program reduced the amount of funding that candidates received from large contributions and increased the amount they received from small contributions. These are policy-relevant findings as public financing for political campaigns gains popularity and as policymakers decide how that public financing should be allocated.

This paper is the first to examine the effects of public financing for political campaigns on political participation. Seattle's Democracy Voucher's program has three main elements that could affect political participation: the provision of public financing, the contribution limits placed on candidates who accept public financing, and the involvement of voters in allocating the financing. By studying this program, I am not able to decompose which of these three elements played the largest role in driving my results or whether all three elements are necessary to increase political participation. Studying alternative public financing policies to determine whether they have similar effects on political participation is a promising area for future research.

The magnitudes of my effects are remarkable considering both the low rate of voter participation in the program and the fact that vouchers were only available for municipal elections, which tend to interest voters less than statewide or federal elections. My results provide encouraging evidence that public finance programs can improve political participation and shift the composition of political contributions, making candidates more reliant on small rather than large donations.

Chapter 3

Marijuana Legalization and Fertility

3.1 Introduction

State-level legalization of marijuana in the United States has had wide-reaching effects, many of which are not yet understood. Health outcomes affected by marijuana legalization include self-reported health [82], sick days taken from work [83], alcohol consumption [84] and opioid mortality [85]. Marijuana legalization could also affect another health outcome that has been absent from policy discussions: fertility. Any impact on fertility is an unintended consequence of the policy, which should be considered as more states and the federal government evaluate whether to legalize marijuana.

The medical literature provides descriptive evidence that marijuana use could either raise or lower fertility. In this paper, I examine how marijuana legalization affects birth rates to determine which effect is dominant. The physical effects of marijuana use could lower the likelihood of pregnancy, while the psychological effects could raise it. First, marijuana use has physical effects on both men's and women's bodies that could lower the likelihood of pregnancy. Marijuana use is associated with a 29% decrease in sperm counts [86] and ovulation delays of up to 3.5 days [87], both of which make conception less likely. Second, marijuana use could lead to sexual behaviors that raise the likelihood of pregnancy. Using marijuana can heighten enjoyment of sexual activity and diminishes the ability to think about long-term consequences of failing to use contraception. Marijuana use is associated with an increase of 1.1-1.3 sexual encounters per month [88] and a decrease in the likelihood of using contraception, with marijuana users being 5.46 percentage points more likely to use condoms inconsistently than non-users [89]. Medical researchers have been unable to establish whether these relationships are causal or which effect is dominant due to legal limitations on studying marijuana use.

In this paper, I study the effects of marijuana legalization on fertility and provide the first causal evidence of how recreational marijuana legalization affects birth rates. I use a difference-in-differences (DID) design that exploits variation in marijuana legalization across states and over time to examine how recreational marijuana laws (RMLs) affect birth rates, marijuana use, sexual activity and gonorrhea cases. I measure changes in the birth rate using restricted-access birth data from the National Vital Statistics System. To gain insight into the mechanisms driving changes in the birth rate, I use restricted-access data on marijuana use and sexual activity from the National Longitudinal Survey of Youth. I use the estimator from [90] to show that my results are robust to both heterogeneous treatment effects and contamination bias that could arise in estimating the effects of RMLs in states that already had medical marijuana laws (MMLs).

Because marijuana legalization must increase marijuana use in order to affect fertility through the channels established in the medical literature, I begin by estimating effects on marijuana use and find suggestive evidence that days of marijuana use per month increase by 41% in response to RMLs. This result is consistent with past findings that marijuana legalization increases marijuana use [91, 92, 93, 94, 95, 96, 97, 98, 99, 100, 101, 102, 103]. Heterogeneity analysis shows that men increase marijuana use more than women in response to RMLs.

Next, I find the main result of the paper: RMLs lead to a 2.78% decline in the average birth rate. This result provides evidence that marijuana's physical effects, which suppress the likelihood of pregnancy conditional on sexual activity, have the dominant effect on fertility. Age heterogeneity analysis shows that these effects are most pronounced for women close to the end of their child-bearing years: the largest decrease in the birth rate occurs among women aged 30-34, closely followed by women 35-39 and then by women 40-44. The birth rate in all three of these age groups declines by over 6%. This heterogeneity analysis suggests that women are having fewer total children in response to RMLs rather than delaying births. Event studies show that RMLs begin to affect the birth rate in the first two quarters after legalization.

Third, I investigate how RMLs affect sexual activity to determine whether the decrease in the birth rate results from a decline in sexual activity or from the fertility-suppressing physical effects of marijuana use. I find suggestive evidence that RMLs increase the probability of being sexually active by 3.6 percentage points. This finding supports the association between marijuana use and increased sexual activity in the medical literature. I do not find statistically significant effects of RMLs on gonorrhea cases or the probability of having sex with a stranger, a measure of risky sexual behavior. While RMLs increase sexual activity, this change is modest and is less important than marijuana-induced physical changes in determining the overall effect on fertility.

Finally, I examine how MMLs affect fertility. I find suggestive evidence that days of marijuana use increase by 23% in response to MMLs and causal evidence that MMLs lead to a small, statistically insignificant decrease in birth rates. The smaller birth rate effect of MMLs is consistent with the smaller increase in marijuana use resulting from these laws. This result contrasts with a finding in the previous literature that MMLs increase birth rates [93]. I expand on the previous literature by using 21 additional years of data on birth rates and by using the estimator from [90] as a robustness check to show that my results are driven by neither negative weights arising from heterogeneous treatment effects nor contamination bias from multiple treatments.

This paper makes two main contributions to the literature. First, it adds to the growing literature on the health effects of marijuana legalization, which is more developed for MMLs than for RMLs. This paper provides the first causal evidence of how RMLs affect fertility and provides new evidence on how MMLs affect fertility by using new data and recent advances in the difference-in-differences literature. The most closely related paper is [93], which investigates the effects of MMLs on fertility. While the relationship between recreational marijuana *consumption* and fertility has been studied in the medical literature [86, 87, 88, 89], these studies have been associative rather than causal and have not studied the effects of legalizing recreational marijuana. [104] study the effects of both MMLs and RMLs on perinatal health and find that RMLs increase maternal hospitalizations due to marijuana use disorder while lowering tobacco use disorder hospitalizations by about the same amount. Neither type of law affects newborn health. [104] use the multiperiod difference-in-differences (DID_M) estimator from [90] for robustness to heterogeneous treatment effects and multiple treatments, as I also do in this paper. A thorough review of the literature on health effects of marijuana legalization is available in [105].

Second, this paper contributes to the literature on determinants of fertility in developed countries. The literature shows that determinants of fertility include unemployment [106], increased insurance coverage of contraceptives [107] and laws that reduce access to contraception and abortion [108]. This paper is the first to provide causal evidence that RMLs impact fertility.

3.2 Background

3.2.1 Marijuana Legalization

Marijuana use has been criminalized at the federal level in the United States since 1937. While the federal law has not changed, the majority of states have now passed MMLs, and a smaller group of states has also passed RMLs. In 1996, California became the first state to implement a MML and has since been followed by thirty-one other states. The timeline of state-level MMLs from 1996-2019 appears in Table 3.1. When a state passes a MML, a citizen of that state can obtain a physician recommendation for marijuana. Marijuana can only be legally purchased and used by the individual who obtained the recommendation, although researchers have found some evidence of spillovers to recreational use when medical marijuana is legalized [94].

State	MML	RML
Alaska	March 4, 1999	February 24, 2015
Arizona	November 29, 2010	
Arkansas	November 9, 2016	
California	November 6, 1996	November 9, 2016
Colorado	December 28, 2000	December 10, 2012
Connecticut	October 1, 2012	
Delaware	July 1, 2011	
District of Columbia	July 27, 2010	February 26, 2015
Florida	January 3, 2017	
Hawaii	June 14, 2000	
Illinois	January 1, 2014	
Louisiana	May 19, 2016	
Maine	December 23, 1999	January 30, 2017
Maryland	June 1, 2014	
Massachusetts	January 1, 2013	December 15, 2016
Michigan	December 4, 2008	December 6, 2018
Minnesota	May 30, 2014	
Missouri	December 6, 2018	
Montana	November 2, 2004	
Nevada	October 1, 2001	January 1, 2017
New Hampshire	July 23, 2013	
New Jersey	June 1, 2010	
New Mexico	June 1, 2007	
New York	July 5, 2014	
North Dakota	December 8, 2016	
Ohio	September 8, 2016	
Oklahoma	July 26, 2018	
Oregon	December 3, 1998	July 1, 2015
Pennsylvania	May 17, 2016	
Rhode Island	January 3, 2006	
Utah	December 3, 2018	
Vermont	July 1, 2004	July 1, 2018
Washington	December 3, 1998	December 6, 2012
West Virginia	July 1, 2019	

Source: Meinhofer et al. (2021). MML: Medical Marijuana Law. RML: Recreational Marijuana Law. The date in the table is the date when the law became effective. Only RMLs and MMLs that became effective in 2019 or earlier are included.

Table 3.1: Marijuana Legalization Dates

When a state passes a RML, any individual who is at least 21 years old can legally buy and use marijuana. States with RMLs have various restrictions on purchases and use of marijuana. For example, Colorado has a RML but does not allow marijuana use in any public place. Colorado and Washington became the first two states to pass RMLs in 2012. Thirteen more states had RMLs by 2020. Table 3.1 shows the timeline of RMLs from 2012-2019. All states with RMLs already had MMLs. Even after a RML is passed, consumers with a medical need for marijuana have an incentive to obtain a physician recommendation for marijuana. Medical marijuana is exempt from sales tax, and a physician recommendation enables the consumer to purchase stronger marijuana and larger quantities of marijuana than a recreational consumer [109]. RMLs could potentially increase marijuana use among individuals who already used marijuana, because legalization makes marijuana easier to buy. Individuals could also begin using marijuana for the first time if they had previously avoided marijuana only because it was illegal.

Marijuana legalization has become a more popular policy in the past two decades: 67% of Americans believed that recreational marijuana use should be legal in 2019, compared with 31% in 1999 [110]. However, it is still controversial. Proponents argue that legalizing marijuana allows law enforcement to focus on more serious crimes, provides a valuable source of tax revenue, preserves individual freedom, and does not harm users' health. Opponents believe that legalization will increase the number of car accidents, encourage both marijuana use and the use of more addictive drugs, and damage societal morality [111]. The policy debate over the health effects of RMLs has been largely restricted to the questions of how marijuana use affects the lungs, the brain and the likelihood of seeking more dangerous or addictive drugs after trying marijuana. Because marijuana legalization is a recent development, researchers are still developing the understanding of its full effects that will be necessary to inform future policy making. One of the effects that researchers have yet to fully explore is marijuana legalization's impact on fertility.

3.2.2 Fertility

Fertility rates in the United States have been consistently below the replacement rate in recent decades [112], raising the prospect of an aging population with a working class that is too small to support the elderly. Fertility rates vary between states: in 2018, Vermont had the lowest fertility rate at 47.2 births per 1,000 women aged 15-44, while South Dakota had the highest rate at 73.6 [113]. The absence of fertility from policy debates over the effects of state-level marijuana legalization suggests that marijuana legalization has not been endogenous to fertility: states are not legalizing marijuana with the goal of increasing birth rates, nor are they refusing to legalize marijuana to decrease sexual activity or encourage condom use. I empirically test this assumption by using event studies to show the absence of differing pre-trends in fertility between states that did and did not legalize marijuana.

3.2.3 Gonorrhea Cases

Gonorrhea cases can provide insight into changes in sexual activity and contraceptive use. Gonorrhea is a sexually transmitted disease that can be cured in less than two weeks with antibiotics. Condoms, the only birth control method that prevents gonorrhea, are 80% effective at preventing transmission [114]. Condom use is more likely than other types of birth control to be affected by marijuana use because individuals must actively choose to initiate condom use during each sexual encounter. When an individual has used marijuana prior to a sexual encounter and is less able to think through the consequences of his or her decisions, condom use is less likely. An increase in gonorrhea cases can be driven by a decrease in condom use or an increase in overall sexual activity with condom use remaining constant.

Gonorrhea is most common outside of long-term, mutually monogamous relationships [115]. Therefore, an increase in gonorrhea cases can also show an increase in sexual activity outside of long-term, mutually monogamous relationships. The number of gonorrhea cases among men is commonly used as an outcome in this literature due to its short duration, which ensures that the times of contraction and diagnosis are close together [93, 116, 117, 114]. [117] finds that drunk driving laws reduce gonorrhea cases among men by 14%; [93] find that MMLs increase gonorrhea cases among men by 12%; and [116] finds that increased access to alcohol upon turning 21 has a null effect on gonorrhea rates, with a 5.2% increase at the upper bound of the 95% confidence interval.

3.2.4 Mechanisms

Marijuana use has behavioral and physiological effects on fertility that could counteract one another. Marijuana use is associated with an increased likelihood of engaging in sexual activity [88, 118] and lower likelihood of using contraceptives conditional on engaging in sexual activity [89, 119, 120]. Marijuana use can also heighten sensory perception, increasing the hedonic effects of sexual activity [121, 122], and decreasing pain associated with sex [123]. Sixty-five percent of medical marijuana recommendations are for chronic pain [124]. A decrease in the severity of chronic pain could affect fertility by enabling a patient to be sexually active if chronic pain had precluded enjoyment of sexual activity.

On their own, these behavioral effects should increase fertility. However, marijuana use also has physical effects that could decrease fertility. Marijuana use in women can delay ovulation and impair embryo implantation [87]. In men, it can reduce sperm count and motility [86, 87, 125]. Marijuana use can disrupt reproductive hormones in both men and women [87, 125]. These effects could counteract the behavioral response to marijuana use, leading to a plateau or decrease in fertility. Whichever fertility effect dominates will determine the overall effect of marijuana legalization on the birth rate.

RMLs and MMLs could have different effects on fertility for two main reasons. First, RMLs increase access to marijuana in the entire adult population, while MMLs increase access primarily among individuals with physician recommendations to use marijuana. While MMLs produce some spillovers to recreational use [92, 94, 126], they provide less expansive marijuana access than RMLs. Effects of RMLs are therefore likely to be larger in magnitude.

Second, these two types of laws could increase marijuana use in different sub-populations, as individuals who obtain physician recommendations for medical marijuana may differ from those who access marijuana only through a recreational dispensary. MMLs affect a subset of the population that is disproportionately likely to suffer from a health condition such as chronic pain. If these individuals tend to be past their child-bearing years, then RMLs can be expected to have stronger effects on fertility. Similarly, if one type of law increases marijuana use more in men and the other type increases use more among women, then the effects of the laws could vary depending on which gender experiences stronger physical or psychological effects of marijuana use.

Birth rates can be affected by changes in miscarriage rates, and [127] finds an association between a male partner's use of marijuana and an increased probability of miscarriage. However, I find that effects of both RMLs and MMLs on miscarriages and stillbirths are statistically insignificant and close to zero, as discussed in the Appendix. Therefore an increase in miscarriage or stillbirth is unlikely to drive changes in the birth rate when marijuana is legalized. A change in abortion is another potential mechanism. I investigate effects on abortion and find statistically insignificant effects, but cannot draw a strong conclusion due to imprecise estimates. Details of this analysis are also provided in the Appendix.

3.3 Data

3.3.1 Birth Rates

To determine how marijuana legalization affects birth rates, I use restricted-access birth data from the National Vital Statistics System.[1] This data includes every birth in the United States from 1989-2019 and provide the date and location of the birth, as well as demographic information about the mother. I aggregate the data to obtain birth counts at the state-quarter-year level and use the birth count to calculate the birth rate. The birth rate is total number of births per 1,000 population in a quarter, multiplied by four to obtain the annualized birth rate as is standard in calculating birth rates [128]. For heterogeneity analysis, I calculate the birth rate within a specific group. For example, I determine the birth rate among Hispanic mothers by finding the number of births to Hispanic mothers in each state-quarter-year and using the total Hispanic population to calculate the birth rate. I use state-level demographic information from the Surveillance, Epidemiology and End Results Program (SEER) to obtain the total population in each state in each year, as well as the number of people in each group that I use for heterogeneity analysis. I have 6,324 state-quarter-year level observations of birth rates. The average birth rate is 3.44. Figure 3.1 shows the average birth rate by the type of marijuana law in effect at the end of 2019. This graph provides suggestive evidence of similar trends in birth rates in the three groups of states. Figure 3.2 shows the percent of births by mother's age. Women aged 25-29 were responsible for the largest share of births from 1989-2019. Births to women under 25 have declined substantially

[1] National Center for Health Statistics. Natality- All Counties (Years: 1989-2019), as compiled from data provided by the 57 vital statistics jurisdictions through the Vital Statistics Cooperative Program.

since 1989, while births to women aged 25 and older have increased.

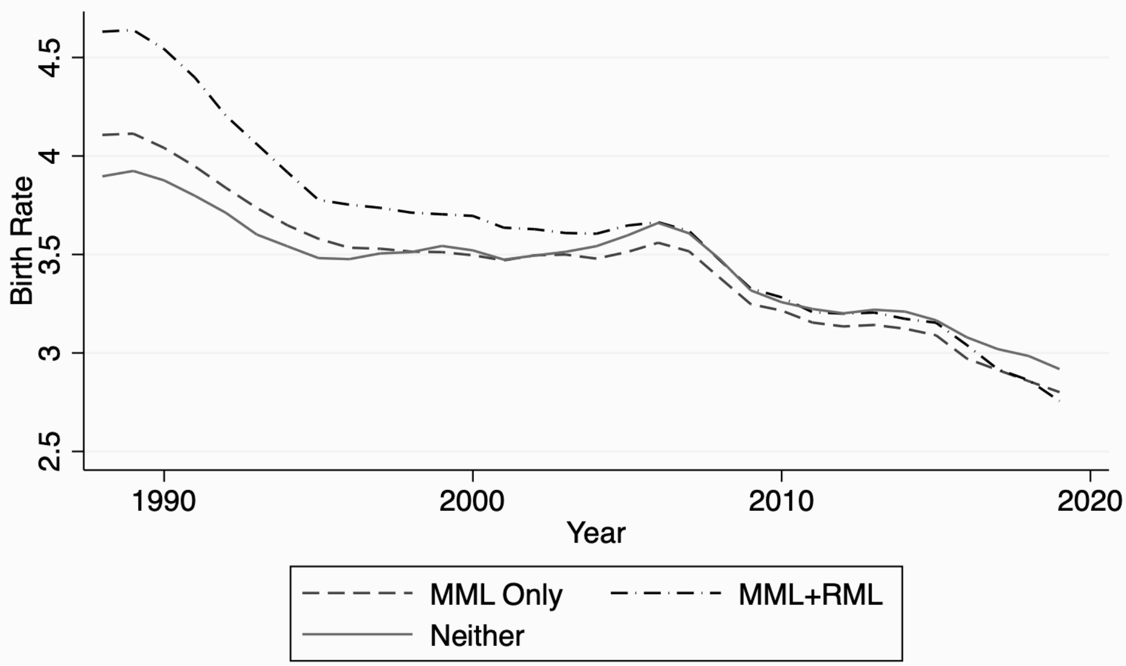

Birth rates calculated using data from the National Vital Statistics System. The birth rate is measured at the state-quarter level and is calculated by dividing the total number of births per quarter by 1,000 population and then multiplying by four. MML only: the state only legalized medical marijuana by 2019. MML + RML: the state legalized both medical and recreational marijuana by 2019. Neither: the state legalized neither medical nor recreational marijuana by 2019.

Figure 3.1: Birth Rate by Treatment Status

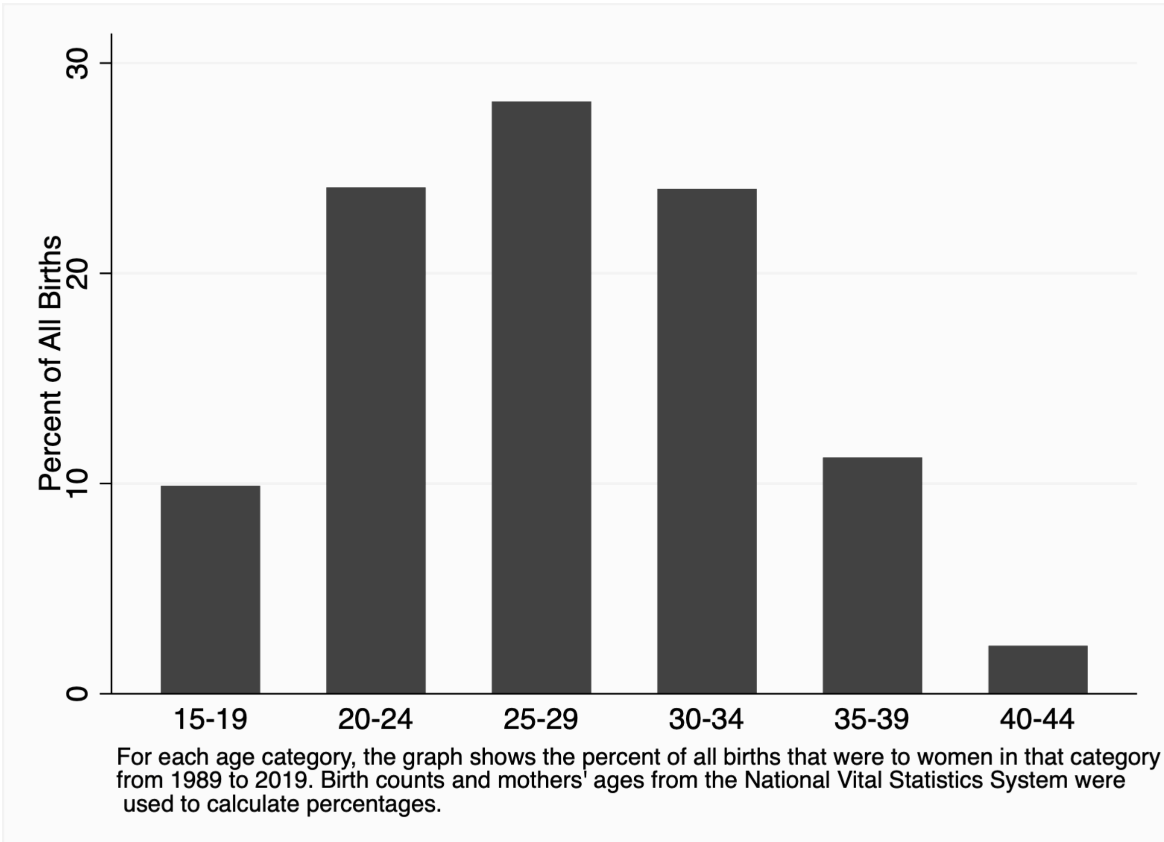

Figure 3.2: Birth Rate by Age Group

3.3.2 Demographic and Policy Controls

I include controls from several sources to account for time-varying factors that affect birth rates. Shares of the total population by race and age group are provided by SEER. Shares of the total population by ethnicity and educational achievement are from U.S. Census data. These controls are included because changes in the demographic composition of a state's population could affect the birth rate: for example, an increase in the share of the population that is elderly will lower the birth rate. I control for changes in abortion policy, which can influence the birth rate by affecting the probability that a pregnancy will end in abortion. I use data from [129] on three categories of laws that restrict access to abortion: ambulatory surgical center laws, admitting privilege laws, and transfer agreement laws. Changes in same-sex marriage laws could potentially impact birth rates by increasing the share of married-couple households that cannot have children accidentally. I therefore control for whether same-sex marriage is legal. The date when same-sex marriage became legal in each state was obtained from the Georgetown Law Library.

Changes in economic conditions can also affect the birth rate, as couples often decide to delay births during recessions [130, 106]. Therefore I include controls for the quarterly unemployment rate and median household income, which are available from the Local Area Unemployment Statistics and Small Area Income and Poverty Estimates, respectively. Access to health care and social safety net programs for low-income women may influence the birth rate by reducing the cost of pregnancy-related healthcare and increasing available household resources. The Kaiser Family Foundation provides data on whether and when states expanded Medicaid under the Affordable Care Act, as well as the Medicaid income eligibility threshold for pregnant women. The United States Department of Agriculture provides data on whether the state's Women Infants and Children (WIC) program has transitioned to using Electronic Benefit Transfers (EBT). WIC provides food to pregnant and postpartum women, as well as children under age five, and is easier to use when it is provided via EBT.

To isolate effects of marijuana legalization, I control for other policies that can affect marijuana consumption. I control for cigarette and beer tax rates, as cigarette and beer price changes may affect marijuana consumption [131, 132]. State cigarette tax rates were obtained from the Campaign for Tobacco-Free Kids.[2] State beer taxes were available from the Tax Policy Center. I also control for decriminalization of marijuana. Decriminalization reduces the penalty for using or possessing marijuana, usually from jail

[2] Cigarette tax rates are available starting in 2000. The 2000 tax rate is used for years prior to 2000.

time to a fine for the first offense. Repeat offenders may still receive jail time. These laws do not allow commercial sale or home cultivation of marijuana, requiring users to purchase marijuana illegally. Decriminalization therefore does not increase access to marijuana in the same way as legalization, and the fine may still deter risk-averse individuals from using marijuana when it is decriminalized.

Table 3.2 shows the differences in the average birth rate and average values of the control variables for states with RMLs compared with other states. These tables show that states with RMLs are demographically similar to states without them.[3]

	All Other States		Recreational Marijuana Legalizers			
	Mean	Std. Dev.	Mean	Std. Dev.	Diff. in Means	Std. Error
Birth Rate	3.39	0.60	3.54	1.02	0.15	0.02
Days of Marijuana Use	1.79	6.21	2.09	6.61	-0.30	0.045
Any Marijuana Use	0.23	0.42	0.26	0.44	-0.03	0.002
Number of Sexual Encounters	6.13	37.0	6.42	39.01	-0.29	0.260
Sexually Active	0.45	0.50	0.45	0.50	0.01	0.003
Sex with a Stranger	0.06	0.23	0.07	0.25	-0.01	0.002
Gonorrhea Cases	3,849.38	4,438.64	4,019.62	8,001.14	-170.25	411.62
Total Population	5,445,718.64	5,255,348.28	6,162,780.08	8,261,225.33	717,061.44	201,898.16
Unemployment Rate	5.34	1.78	5.89	2.00	0.56	0.05
Median Household Income	42,488.57	12,100.37	46,420.61	12,946.37	3,932.03	343.83
Percent White	81.76	13.94	83.25	13.80	1.49	0.38
Percent Black	12.68	10.27	9.36	13.67	-3.32	0.34
Percent Native American	5.64	11.68	7.52	5.38	1.89	0.22
Percent Hispanic	9.25	9.76	12.99	11.12	3.74	0.29
Percent Aged 0-19	27.80	2.57	27.15	2.62	-0.66	0.07
Percent Aged 20-39	28.38	2.52	29.27	3.67	0.89	0.09
Percent Aged 40-64	30.36	3.19	30.73	3.35	0.38	0.09
Percent Aged 65+	13.46	2.04	12.85	2.49	-0.61	0.06
Percent High School	80.23	8.00	83.41	5.84	3.17	0.18
Percent College	22.92	5.87	26.77	7.45	3.85	0.19

Table 3.2: Summary Statistics

3.3.3 Marijuana Use and Sexual Activity

I use restricted-access data from the 1997 National Longitudinal Survey of Youth (NLSY) to measure the effects of marijuana legalization on marijuana use and levels

[3]The Appendix provides equivalent summary statistics for states with and without MMLs, and shows similar evidence of demographic similarities between the two groups of states.

of sexual activity. The NLSY provides individual-level panel data on 8,984 individuals from 1997-2017. From 1997 to 2011, each individual was surveyed once per year; after 2011, the survey was only conducted in 2013, 2015 and 2017. The restricted-access data include geographic identifiers and individual demographic information, such as race, age, sex, marital status and education level. NLSY respondents are of child-bearing age or slightly younger throughout the survey: respondents' ages ranged from 12 to 18 when the survey began in 1997 and from 32 to 38 in the final year, 2017. The NLSY provides weights to make the survey representative at the national level.

Table 3.2 shows that states that passed RMLs had higher average days of marijuana use per month and a higher percentage of respondents who had used any marijuana in the past month than states without RMLs: 2.09 compared to 1.79 and 26% compared to 23%, respectively.

3.3.4 Gonorrhea Cases

To estimate the effects of marijuana legalization on gonorrhea cases, I use data from the Centers for Disease Control and Prevention. Data are publicly available and provide gonorrhea case counts at the state-year level from 2000-2019. I study gonorrhea cases among men rather than women because gonorrhea case numbers are measured more accurately for men. Up to 80% of women with gonorrhea are asymptomatic, while only 10-20% of men are asymptomatic [133]. Men's symptoms are typically painful, alarming and appear within 14 days of exposure, creating a high likelihood that a man's gonorrhea case will be reported to a medical professional within two weeks of exposure to gonorrhea [116], while women without symptoms are unaware that they need to seek treatment. Due to the gender difference in the likelihood of developing symptoms, studying only gonorrhea cases among men is standard in the literature [116, 93].

3.4 Methodology

3.4.1 Birth Rates

I use a difference-in-differences design to analyze the effects of RMLs and MMLs on fertility, taking advantage of variation in when states adopted these laws. I use event studies with the following specification:

$$Y_{st} = \alpha + \sum_{\substack{j=-7 \\ j\neq -1}}^{4} \beta_j \mathbf{1}(t=j) + \phi X_{st} + \gamma_s + \rho_t + \theta_s t + \epsilon_{st} \tag{3.1}$$

Y_{st} is the birth rate in state s from births conceived in semester t. The semester of conception is backed out from the semester of birth under the assumption that each pregnancy lasts nine months. I use the semester of conception because an increase in marijuana use can have an immediate effect on conception, but an effect on births does not materialize until the end of a pregnancy.

The number of semesters since marijuana legalization is j .[4] The first treated semester for each state is the semester when the RML or MML became effective. Precise dates when treatment began appear in Table 3.1. Although some states did not yet have operational dispensaries when their laws became effective, the laws could begin to affect marijuana use as soon as they become effective by changing the penalty associated with marijuana use and the behavior of law enforcement. I conduct separate event studies for RMLs and MMLs, as is standard in the literature [104]. The semester immediately prior to treatment is omitted from the event studies.

X_{st} is a vector of controls: shares of the total population by race, ethnicity, age, and

[4] A semester is six months, either January-June or July-December.

education; unemployment rate; median household income; state cigarette tax; state beer tax; an indicator for whether the state has expanded Medicaid; indicators for abortion restrictions in the form of ambulatory surgical center laws, admitting privilege laws, and transfer agreement laws; an indicator for whether same-sex marriage is legal; the Medicaid eligibility threshold for pregnant women as a percentage of the federal poverty level; a WIC EBT indicator; and an indicator for whether marijuana has been decriminalized. The regressions include state fixed effects, year-quarter fixed effects and state-specific linear time trends. The state fixed effects control for the time-invariant characteristics of each state, the year-quarter fixed effects control for changes in the birth rate over time that are common to all states, and the state-specific linear time trends control for linear changes in the birth rate over time within each state. Standard errors are clustered at the state level. Regressions are weighted by the total state population, as is standard in the literature when birth rates are used as an outcome [134, 135, 136].

Event studies are used to test the plausibility of the parallel trends assumption and examine dynamic effects of marijuana legalization on birth rates. In addition to event studies, I estimate two-way fixed effects regressions to obtain estimates of the average effect of marijuana legalization on birth rates in the post-treatment period. The two-way fixed effects estimates can be interpreted as causal when the event study provides evidence of small, statistically insignificant differences in trends between the treated and control groups prior to treatment. I use the following regression specification:

$$Y_{st} = \beta_0 + \beta_1 Recreational_{st} + \beta_2 Medical_{st} + \beta_3 X_{st} + \gamma_s + \rho_t + \theta_s t + \epsilon_{st} \tag{3.2}$$

Y_{st} is the birth rate in state s at time t, where time is measured at the year-quarter level and t is the estimated quarter of conception. The quarter of conception is backed out from the quarter of birth under the assumption that each pregnancy lasts nine months. $Recreational_{st}$ is a dummy variable indicating whether a RML is in effect in state s in conception quarter t. $Medical_{st}$ is an indicator for whether a MML is in effect in state s in conception quarter t. The first treated quarter for each state is the quarter when the RML or MML became effective. The coefficients of interest are β_1 and β_2. The regressions include state fixed effects, year-quarter fixed effects and state-specific linear time trends. Controls, weighting and clustering are the same as in the event studies.

For the main results, I study effects on general birth rates for women 14-44, as well as birth rates for women 21-30. While the birth rate to women 14-44 is the standard measure of birth rates, I also study effects on women 21-30 because births are most common among women in this age group, as Figure 3.2 shows. I perform heterogeneity analysis by race, education level and age.

I use the same methodology to study gonorrhea cases with year rather than year-quarter fixed effects because the gonorrhea data is annual rather than quarterly. A state is considered treated in the first year in which a RML or MML becomes effective. I follow the CDC in using the number of gonorrhea cases per 100,000 as the gonorrhea case rate.

3.4.2 Marijuana Use and Sexual Activity

I use individual-level data from the NLSY to measure the effects of marijuana legalization on marijuana use and sexual activity. I therefore use a slightly different specification for event studies for these outcomes:

$$Y_{ist} = \alpha + \sum_{\substack{j=-5 \\ j\neq -1}}^{4} \beta_j \mathbf{1}(t=j) + \phi X_{st} + \kappa Z_{ist} + \gamma_s + \rho_t + \theta_s t + \epsilon_{st} \tag{3.3}$$

Y_{ist} is the outcome for person i in state s in semester t. Z_{ist} is a vector of person-level demographic controls, including age, marital status, race and education. The two treatments, state-level controls, state and year-quarter fixed effects, and state-specific linear time trend are the same as in the birth rate analysis. Regressions are weighted using NLSY sample weights. Standard errors are clustered at the state level.

I use the following specification for two-way fixed effects regressions for these outcomes:

$$Y_{ist} = \alpha + \beta_1 Recreational_{st} + \beta_2 Medical_{st} + \beta_3 X_{st} + \beta_4 Z_{ist} + \gamma_s + \rho_t + \theta_s t + \epsilon_{ist} \quad (3.4)$$

Y_{ist} is the outcome for respondent i in state s in quarter t. As the NLSY data provides the precise date of the interview, $Recreational_{st}$ is equal to one if a RML was effective during the quarter in which the interview was conducted and $Medical_{st}$ is equal to one if a MML was effective during the quarter in which the interview was conducted.

3.4.3 The DID_M Estimator

Recent advances in the difference-in-differences literature have shown that two-way fixed effects estimates can be biased when treatment timing is staggered and treatment effects are heterogeneous across groups or over time. These advances are relevant to my setting, as the timing of RMLs and MMLs varies across states and effects of each law may also vary by state. Several estimators have been proposed to produce unbiased estimates in two-way fixed effects designs with staggered adoption [137]. I use the DID_M estimator from [90] for two reasons. First, this estimator is unbiased in the presence of treatment effect heterogeneity across both time and states. Second, the estimator eliminates bias that arises from the presence of a second treatment, while other estimators in this literature are designed for a single treatment [137]. The DID_M estimator is therefore ideally suited to my setting, in which the two treatments are RMLs and MMLs. Identifying the effects of RMLs and MMLs is even used as an example in [90] of an ideal setting in which their method can be used. [104] also uses the DID_M estimator as their preferred method of addressing bias from negative weights when studying effects of MMLs and RMLs on perinatal health.

In a two-way fixed effects design with staggered adoption, the coefficient of interest is a weighted sum of estimated average treatment effects on the treated units (ATTs), where the weights may be negative. Negative weights lead to biased estimates: for example, they could lead to a negative estimate even if all the ATTs were positive. The presence of multiple treatments is another potential source of bias. The standard approach of controlling for the presence of the first treatment when estimating the effect of the second treatment is insufficient to address this concern, unless the effects of both treatments are constant over time [90]. To estimate average treatment effects, the DID_M estimator compares the evolution of an outcome from period $t-1$ to t for a state whose treatment status changes from $t-1$ to t to the evolution of the same outcome for a state that was untreated at both $t-1$ and t [90].

Regressions using the DID_M estimator follow the same specifications as the two-way fixed effects regressions above. When the DID_M estimator is used, reported coefficients are DID_M estimates of average treatment effects in the post-treatment period. [90] emphasize that their method should only be used in cases in which two-way fixed effects estimates are substantially biased by negative weights, as their estimator is less efficient than two-way fixed effects. I therefore use their method only as a robustness check, because a comparison between their method and the standard two-way fixed effects estimate shows that the two-way fixed effects estimate is not substantially biased by negative weights. An important note in interpreting these comparisons is that the

DID_M estimand is different from the two-way fixed effects estimand, and therefore the average effect estimated using DID_M should not be expected to be exactly the same as the two-way fixed effects estimate even in the absence of bias from negative weights.

3.5 Results

3.5.1 Marijuana Use

Figure 3.3 shows the event study for effects of RMLs on days of marijuana use in the past month. The event study provides suggestive evidence of an increase in days of marijuana use when a RML becomes effective, with the largest increase occurring in the first year after implementation. Point estimates are positive in the first two semesters after legalization. The second post-treatment estimate is the only statistically significant dynamic effect, and the first and third post-treatment estimates have large confidence intervals.

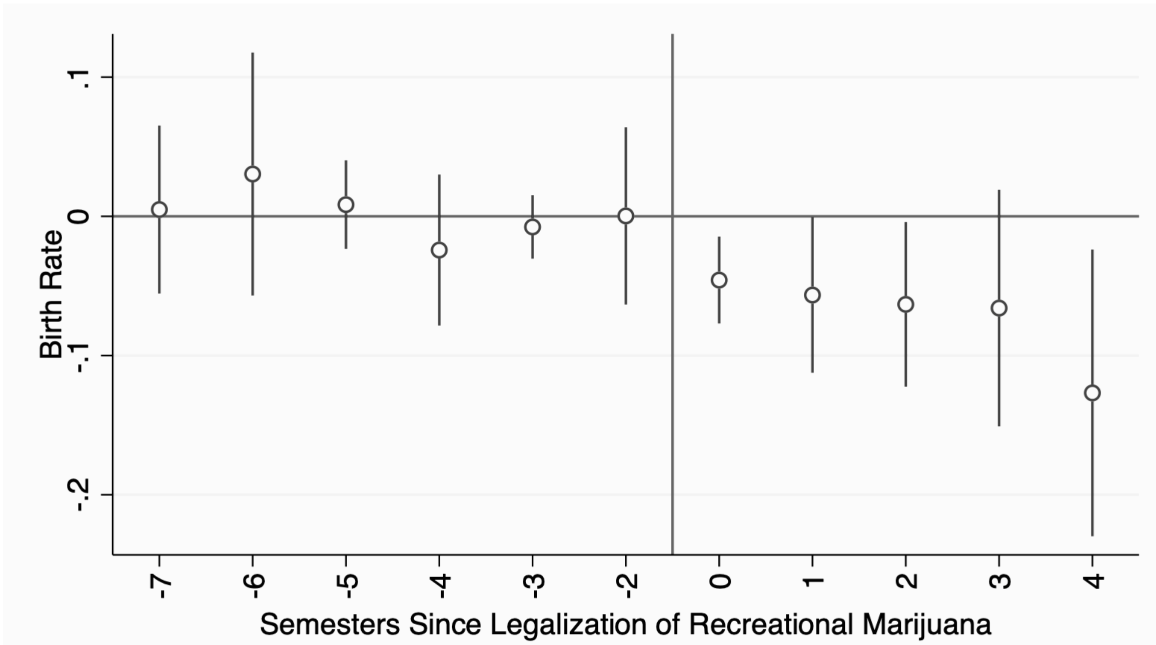

The outcome is the birth rate, measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. Coefficients estimated using two-way fixed effects. Includes state fixed effects, year-quarter fixed effects and state-specific linear time trends as well as controls for share of total state population by race, ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; and decriminalization of marijuana. The regression is weighted by state population and standard errors are clustered at the state level.

Figure 3.3: RML Event Study for Days of Marijuana Use

For the parallel trends assumption to be plausible, the pre-treatment estimates should be close to zero and statistically insignificant. One of the pre-treatment estimates is statistically significant, and the estimates are not close enough to zero to provide strong support for the parallel trends assumption. Estimated effects of RMLs on marijuana use should therefore be interpreted as suggestive evidence, rather than as causal effects.

The large confidence intervals in the post-treatment period of the event study suggest that more power is needed to precisely estimate the effect of RMLs on marijuana use. The two-way fixed effects regression is better-powered than the event study because it estimates fewer parameters. Table 3.3 shows the two-way fixed effects estimate: RMLs increase monthly days of marijuana use by an average of 0.76 days. This estimate is significant at the 5% level and represents a 41% increase from the sample average of 1.85 days of marijuana use per month. I find that RMLs do not have a statistically significant effect on the probability of any marijuana use in the past month, and the point estimate for this outcome is very small.[5]

	Full Sample (TWFE)	Full Sample (DID_M)	Men Only	Women Only
Panel A: Days of Marijuana Use				
Any RML Effective	0.762***	-1.52	1.751**	-0.243
	(0.299)	[-4.50, 1.46]	(0.841)	(0.466)
Any MML Effective	0.425**	-4.03	0.259	0.578**
	(0.198)	[-9.28, 1.22]	(0.238)	(0.249)
Observations	102,940	102,940	50,526	52,414
Panel B: Any Marijuana Use				
Any RML Effective	-0.0157	-0.15	-0.0144	-0.0145
	(0.0125)	[-0.35, 0.046]	(0.0175)	(0.0102)
Any MML Effective	0.00237	-0.10	-0.00565	0.0102
	(0.00895)	[-0.256,0.057]	(0.0115)	(0.0112)
Observations	112,558	112,558	54,975	57,583

*: p<0.1 **: p<0.05 ***: p<0.01. Standard errors clustered at the state level in parentheses. For DID_M Estimates only, 95% confidence interval bounds are in brackets. Outcomes are days of marijuana use last month and an indicator for any marijuana use last month. Outcomes are measured at the individual level using data from the National Longitudinal Survey of Youth from 1997-2017. Coefficients estimated using two-way fixed effects. Regressions include state fixed effects, year-quarter fixed effects and state-specific linear time trends, as well as controls for share of total state population by race, ethnicity, age, and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; decriminalization of marijuana; and the respondent's age, sex, education and marital status.

Table 3.3: Marijuana Use Results

These results provide evidence that RMLs increase marijuana use, with the increase occurring mainly at the intensive margin: marijuana legalization primarily increases the amount of marijuana use among individuals who already used marijuana, rather than

[5] These results are robust to the use of a Poisson regression rather than ordinary least squares for days of marijuana use, use of the wild cluster bootstrap to determine statistical significance, and removal of states that passed MMLs but not RMLs from the control group. Details of these robustness checks are available in the Appendix.

creating new marijuana users. Estimates that include individual fixed effects produce even stronger evidence that the increase in marijuana use occurs only at the intensive margin, as is discussed in more detail in the Robustness Checks section.

An important caveat in interpreting these results is that the oldest respondent in the NLSY data in the final survey year is 38 years old. Past literature has found that individuals who use marijuana for the first time in response to legalization tend to be older adults [138] and that older adults increase marijuana use more than young adults in response to legalization [139, 96, 98]. One possibility is that effect sizes would be larger, and that increases would also occur at the extensive margin, if older adults were included in the sample. However, the age range of NLSY respondents is well-suited to this paper because they are in their child-bearing years and their response to marijuana legalization may induce changes in their fertility.

The results for marijuana use are consistent with findings in the literature that RMLs increase marijuana use [100, 101, 102, 103], as well as evidence that MMLs increase marijuana use [91, 92, 93, 94, 95, 96, 97, 98, 99]. Other studies have found that RMLs are not associated with statistically significant increases in marijuana use [138, 140] and that MMLs are not associated with increases in the probability of marijuana use among adolescents or young adults [141, 142, 143, 144, 145, 96, 98, 140].

Heterogeneity analysis by sex, displayed in Table 3.3, suggests that the increase in days of marijuana use from RMLs is driven by men rather than women. Marijuana use increases by 1.8 days per month among men, with a statistically insignificant effect on women. This finding complements existing evidence that men use marijuana for recreational purposes at higher rates than women [146, 147, 148].

3.5.2 Birth Rates

Figure 3.4 and Figure 3.5 display event studies for the effects of RMLs on birth rates among women aged 14-44 and 21-30, respectively. Pre-treatment estimates in both event studies are statistically insignificant and close to zero, which provides support for the parallel trends assumption. Both event studies show evidence of a decline in the birth rate after a RML is implemented: all post-treatment estimates are negative, and all but one are statistically significant.

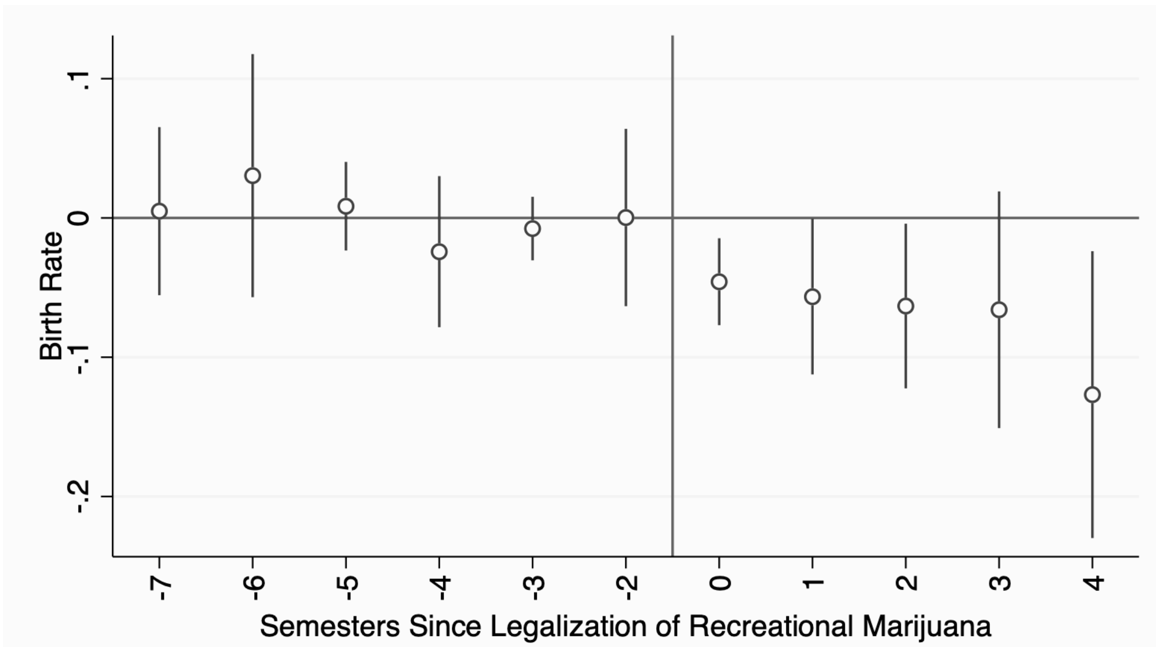

The outcome is the birth rate, measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. Coefficients estimated using two-way fixed effects. Includes state fixed effects, year-quarter fixed effects and state-specific linear time trends as well as controls for share of total state population by race, ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; and decriminalization of marijuana. The regression is weighted by state population and standard errors are clustered at the state level.

Figure 3.4: RML Event Study for Birth Rate (Women 14-44)

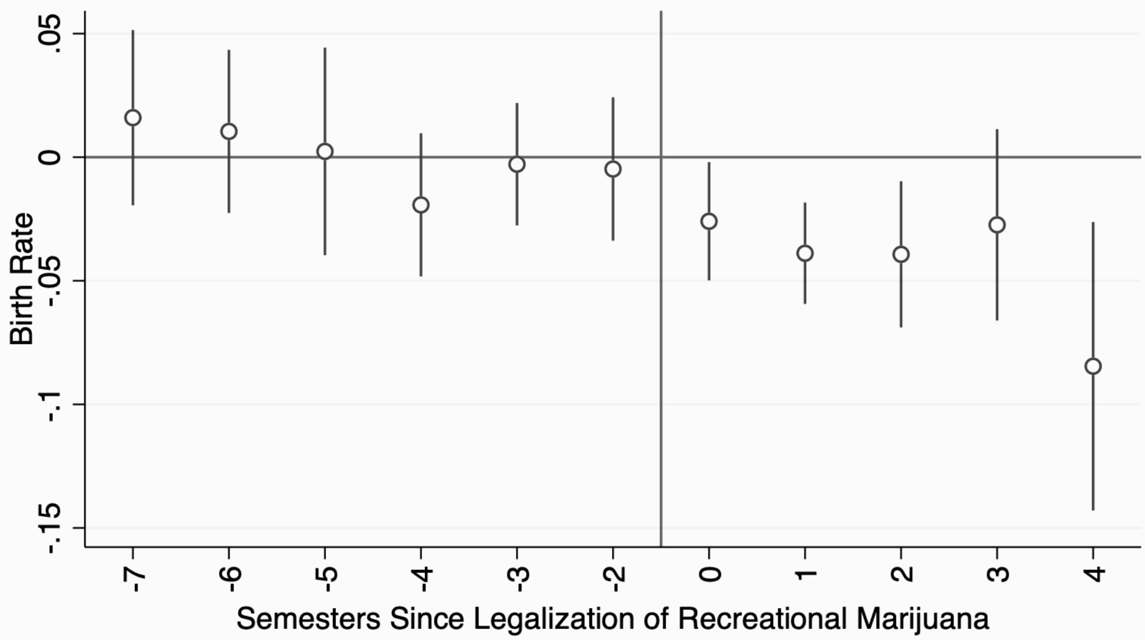

The outcome is the birth rate, measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. Coefficients estimated using two-way fixed effects. Includes state fixed effects, year-quarter fixed effects and state-specific linear time trends as well as controls for share of total state population by race, ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; and decriminalization of marijuana. The regression is weighted by state population and standard errors are clustered at the state level.

Figure 3.5: RML Event Study for Birth Rate (Women 21-30)

Table 3.4 displays two-way fixed effects estimates of the effects of RMLs on birth rates from 1989-2019. A RML is estimated to decrease the birth rate among women ages 14-44 by 0.097 and to decrease the birth rate among women ages 21-30 by 0.057. Both estimates are significant at the 1% level. The average birth rate among women 14-44 is 3.49, and the average birth rate among women 21-30 is 1.91. Therefore, I estimate a 2.78% decrease from the mean birth rate for women 14-44 and a 2.98% decrease from the mean birth rate for women 21-30.[6] These estimates suggest that the physical effects of

[6]These results are robust to the use of a shorter pre-treatment period, use of birth counts instead of birth rates, use of the wild cluster bootstrap to determine statistical significance, removal of states that passed MMLs but not RMLs from the control group, and removal of state-specific linear trends. Details of these robustness checks are available in the Appendix.

marijuana use, which lower fertility, tend to dominate the behavioral effects that could increase fertility.

	Women 14-44	Women 21-30
Panel A: TWFE		
Any RML Effective	-0.097***	-0.057***
	(0.036)	(0.018)
Any MML Effective	-0.027	-0.004
	(0.032)	(0.016)
Panel B: DID_M		
Any RML Effective	-0.034	-0.027
	[-0.102, 0.033]	[-0.059, 0.005]
Any MML Effective	-0.001	-0.007
	[-0.039, 0.037]	[-0.027, 0.013]
Observations	6,324	6,324

*: p<0.1 **: p<0.05 ***: p<0.01. The outcome is the birth rate, which is measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. All regressions include controls for share of total population by race, ethnicity, age and education, as well as unemployment rate, median household income, state cigarette and beer tax rates, Medicaid expansion, abortion laws, same sex marriage laws, Medicaid eligibility threshold, WIC EBT indicator, and decriminalization of marijuana. State fixed effects, year-quarter fixed effects and a state-specific time trend are included, and regressions are weighted by state population. Standard errors are clustered at the state level. TWFE denotes two-way fixed effects, with standard errors in parentheses. DID_M denotes the estimator from Chaisemartin and D'Haultfoeuille (2020), with bounds of the 95% confidence interval in brackets. RML: recreational marijuana law. MML: medical marijuana law.

Table 3.4: Birth Rate Results

Table 3.5 displays the results of the heterogeneity analysis by age group. The effects of RMLs on birth rates follow a roughly u-shaped pattern as age increases, with the largest negative effect occurring among women 30-34. In terms of percent changes from the mean, the birth rate for women 30-34 declines by 6.4%, the birth rate for women 35-39 declines by 6.3% and the birth rate for women 40-44 declines by 6.1%. These results show that the largest decreases in the birth rate occur among women close to the end of their child-bearing years. Birth rates for women 25-29 decrease by 2.1%, while birth rates for women under 25 do not change by a statistically significant amount. This evidence is consistent with findings that RMLs increase marijuana use among individuals aged 30-49 [101] and increase frequent use in individuals aged 26 and older [100]. This age heterogeneity analysis suggests that RMLs are not merely delaying births, as women's fertility declines sharply after the age of 35 [149] and Figure **??** shows that mothers tend to give birth before age 35.

Panel A: Heterogeneity by Race and Education						
	Hispanic	**Black**	**White**	**High School**	**College**	
Any RML Effective	-0.199	-0.244	-0.049	-0.196	-0.281	
	(0.135)	(0.345)	(0.270)	(0.198)	(0.172)	
Any MML Effective	-0.148	-0.273	-0.477**	-0.645***	-0.320*	
	(0.141)	(0.173)	(0.188)	(0.175)	(0.160)	
Observations	6,324	6,324	6,324	6,324	6,324	
Panel B: Heterogeneity by Mother's Age						
	15-19	**20-24**	**25-29**	**30-34**	**35-39**	**40-44**
Any RML Effective	0.036	0.066	-0.304**	-0.723***	-0.322***	-0.063***
	(0.094)	(0.149)	(0.146)	(0.121)	(0.075)	(0.021)
Any MML Effective	-0.099	-0.197	-0.000	-0.122	-0.043	-0.007
	(0.137)	(0.205)	(0.071)	(0.074)	(0.050)	(0.017)
Observations	6,324	6,324	6,324	6,324	6,324	6,324

*: p<0.1 **: p<0.05 ***: p<0.01. The outcome is the birth rate, which is measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. All regressions include controls for share of total population by race, ethnicity, age and education, as well as unemployment rate, median household income, state cigarette and beer tax rates, Medicaid expansion, abortion laws, same sex marriage laws, Medicaid eligibility threshold, WIC EBT indicator, and decriminalization of marijuana. State fixed effects, year-quarter fixed effects and a state-specific time trend are included, and regressions are weighted by state population. Standard errors are clustered at the state level. Standard errors in parentheses.
RML: recreational marijuana law. MML: medical marijuana law.

Table 3.5: Birth Rate Heterogeneity Results

3.5.3 Sexual Activity

I next examine how RMLs affect the probability that an individual has been sexually active in the past month. Figure 3.6 shows the event study. While all pre-treatment estimates are statistically insignificant, the event study does not provide strong support for the parallel trends assumption because the point estimates trend upward during the pre-treatment period. Therefore I interpret estimated effects on this outcome as suggestive evidence, rather than as causal effects. In the post-treatment period, all but one of the point estimates are positive, and three of the four positive point estimates are statistically significant. These post-treatment estimates suggest that RMLs increase the probability of being sexually active. The event study estimates are imprecise, with a particularly large confidence interval for the first post-treatment estimate. The two-way fixed effects regression is better-powered and provides a more precise estimate of the average post-treatment effect.

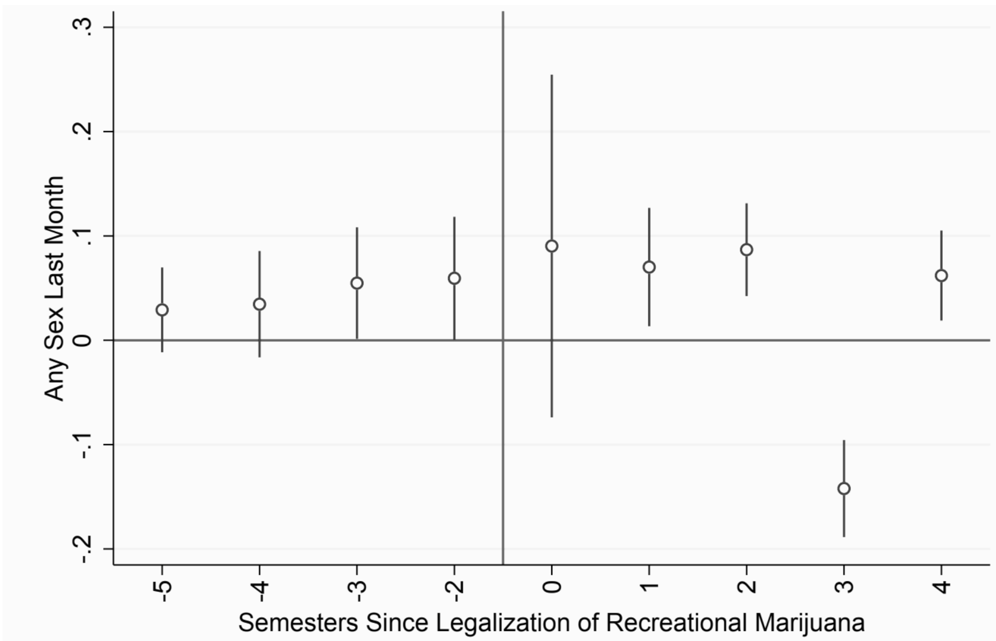

An indicator for any sexual activity in the last month is the outcome, measured annually at the individual level using data from the National Longitudinal Survey of Youth from 1997-2017. Coefficients are estimated using two-way fixed effects. Includes state fixed effects, year-quarter fixed effects and state-specific linear time trends as well as controls for share of total population by race, ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; decriminalization of marijuana; and the respondent's race, age, sex, education and marital status.

Figure 3.6: RML Event Study for Any Sexual Activity Last Month

Table 3.6 shows the two-way fixed effects estimate for this outcome. RMLs are estimated to increase the probability of being sexually active in the past month by 3.6 percentage points, and this effect is statistically significant at the 1% level. The effect remains positive, similar in magnitude and statistically significant at the 1% level when individual fixed effects are included, as shown in the Appendix. Two-way fixed effects estimates of the effect of RMLs on the number of sexual encounters per month and the probability of having sex with a stranger, a measure of risky sexual behavior, are statistically insignificant. The number of sexual encounters is estimated to increase by 0.64, which is equivalent to 9.9% of the mean in states with RMLs. The estimated effect on the probability of having sex with a stranger is close to zero.

	Full Sample	Men Only	Women Only
Panel A: Number of Sexual Encounters			
Any RML Effective	0.642	1.651	-0.216
	(0.819)	(1.420)	(0.628)
Any MML Effective	1.559**	3.353**	-0.516
	(0.773)	(0.505)	(0.0118)
Observations	112,557	54,975	57,582
Panel B: Any Sexual Encounters			
Any RML Effective	0.036***	0.049**	0.017
	(0.011)	(0.020)	(0.021)
Any MML Effective	0.006	0.0253**	-0.0164
	(0.009)	(0.012)	(0.012)
Observations	112,558	54,975	57,583
Panel C: Sex with a Stranger			
Any RML Effective	-0.006	-0.034**	0.0231*
	(0.009)	(0.016)	(0.013)
Any MML Effective	-0.007	-0.006	-0.009**
	(0.006)	(0.009)	(0.004)
Observations	82,217	39,674	42,543
Panel D: Gonorrhea Cases			
Any RML Effective	6.070		
	(7.370)		
Any MML Effective	0.750		
	(7.240)		
Observations	1,020		

*: p<0.1 **: p<0.05 ***: p<0.01. Standard errors clustered at the state level in parentheses. All outcomes except gonorrhea cases are measured at the individual level using data from the National Longitudinal Survey of Youth from 1997-2017. Gonorrhea cases: number of gonorrhea cases among men per 100,000 population, measured at the state-year level using CDC data from 2000-2019. Coefficients estimated using two-way fixed effects. Regressions include state and time fixed effects, state-specific linear time trends, and controls for share of total state population by race, ethnicity, age, and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; decriminalization of marijuana; and the NLSY respondent's age, sex, education and marital status for NLSY outcomes. Regressions using NLSY data are weighted by NLSY sample weights. The regression for gonorrhea cases is weighted by state population.

Table 3.6: Sexual Activity Results

These findings suggest that marijuana legalization increases sexual activity, which could result in a higher birth rate. The decrease in the birth rate that results from RMLs shows the strength of marijuana's negative physical effects on fertility, which dominate the behavioral effects that could increase fertility.

3.5.4 Gonorrhea Cases

Next, I study how RMLs affect gonorrhea cases. Figure 3.7 shows the event study for gonorrhea cases. Estimates in the pre-treatment period are statistically insignificant and close to zero, providing evidence of parallel trends prior to RML implementation. Four of the five post-treatment estimates are positive, but none are statistically significant and estimated effects two, three and four years after RML implementation have large confidence intervals.

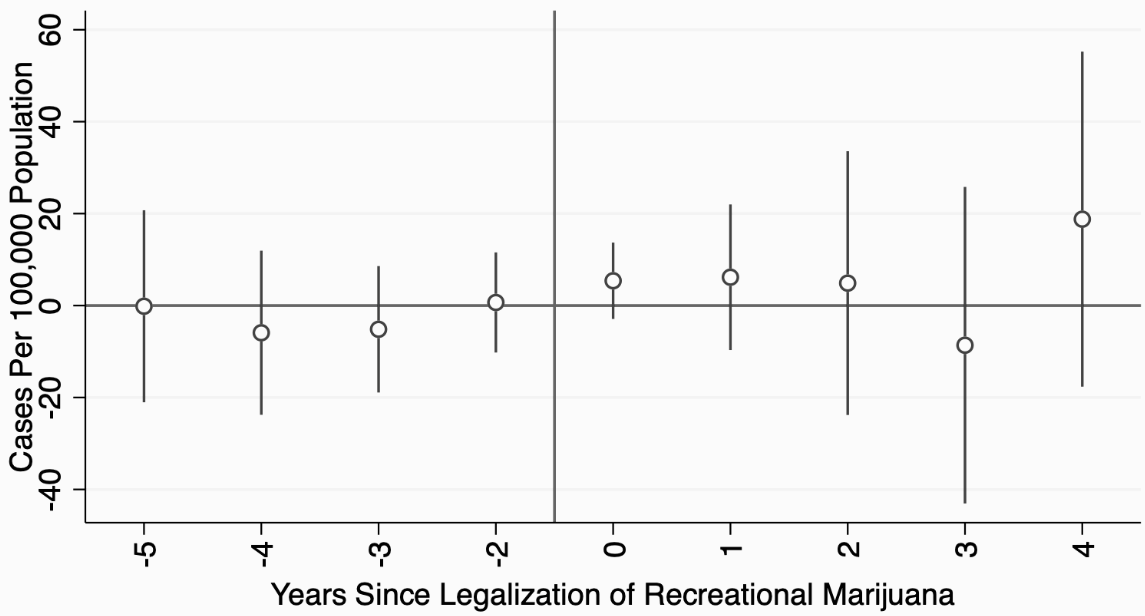

The outcome is the number of gonorrhea cases among men per 100,000 population, measured at the state- year level using data from the Centers for Disease Control and Prevention from 2000-2019. Coefficients estimated using two-way fixed effects. Includes state fixed effects, year fixed effects and state-specific linear time trends as well as controls for share of total state population by race,ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; and decriminalization of marijuana. The regression is weighted by state population and standard errors are clustered at the state level.

Figure 3.7: RML Event Study for Gonorrhea Cases

The two-way fixed effects estimate, shown in Table 3.6, is an increase of 6.1 cases per 100,000, which represents a 5% increase from the mean and is consistent with the positive dynamic effects estimated in the event study. Although this estimate is not close to zero, it is statistically insignificant, with a large standard error. Together, the event study and the two-way fixed effects estimate provide evidence of an increase in gonorrhea cases resulting from RMLs, which supports evidence from the NLSY of an increase in sexual activity. However, these estimates are imprecise and should therefore be interpreted with caution. The lack of precision for these estimates likely results from the smaller sample size for gonorrhea cases than for the other outcomes studied in this paper.

3.6 Effects of Medical Marijuana Laws

In this section, I analyze the effects of MMLs on fertility in order to compare effects of medical and recreational legalization. I begin by examining how MMLs affect marijuana use. Figure 3.8 shows the event study for effects of MMLs on days of marijuana use in the past month. All estimates, both before and after MML implementation, are negative and statistically insignificant. The estimates are imprecise, suggesting that more power is needed to precisely estimate the effect of MMLs on marijuana use. The two-way fixed effects estimate displayed in Table 3.3 is an increase of 0.43 days. This estimate is statistically significant at the 5% level and represents a 23% increase from the mean, which is smaller than the estimated effect of RMLs. The larger magnitude of the effect of RMLs is consistent with the larger increase in access to marijuana under this type of law.

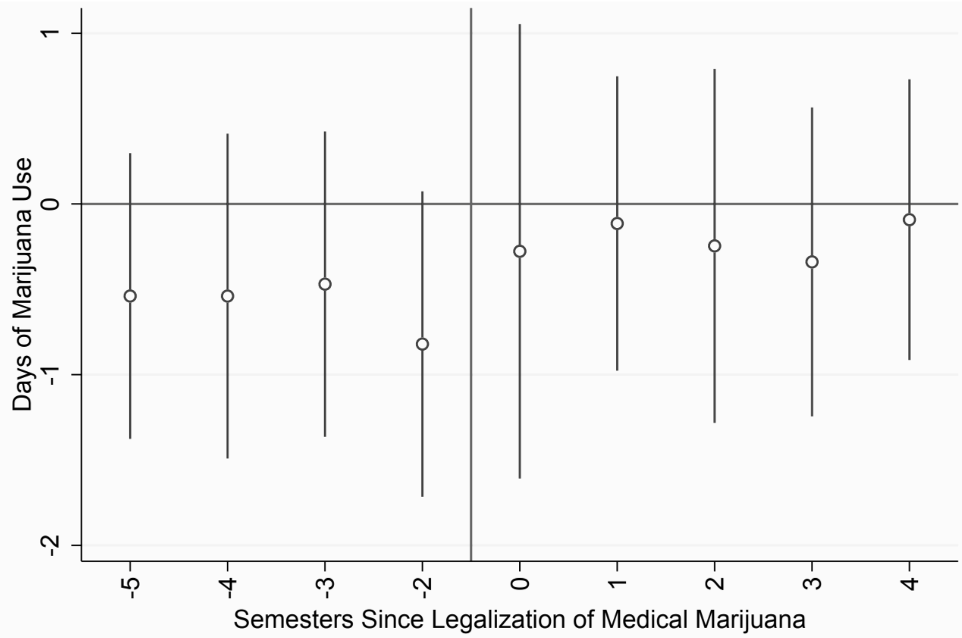

Days of marijuana use per month is the outcome, measured annually at the individual level using data from the National Longitudinal Survey of Youth from 1997-2017. Coefficients are estimated using two-way fixed effects. Includes state fixed effects, year-quarter fixed effects and state-specific linear time trends as well as controls for share of total population by race, ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; decriminalization of marijuana; and the respondent's race, age, sex, education and marital status.

Figure 3.8: MML Event Study for Days of Marijuana Use

Heterogeneity analysis shows that MMLs induce more marijuana use among women than men, while the reverse is true for RMLs. [150] provide similar evidence that women are more likely than men to receive medical cannabis cards, and are more likely than men to increase marijuana use upon acquiring a medical cannabis card. Marijuana's differing physical effects on men and women may contribute to RMLs and MMLs having different effects on fertility. For example, decreased motility of sperm from men's use of marijuana may have a stronger effect on fertility than the suppression of ovulation from women's use of marijuana.

Next, I study how MMLs affect birth rates. Figure 3.9 displays the event study for women aged 21-30. The equivalent event study for women aged 14-44 is available in the Appendix. I find stronger support for the parallel trends assumption for birth rates among women aged 21-30: while none of the pre-treatment estimates in either event study are statistically different from zero, pre-treatment point estimates are closer to zero in the event study for women aged 21-30. I therefore focus on estimated effects of MMLs on birth rates for women aged 21-30, where there is stronger evidence that effects can be interpreted as causal. All dynamic effects in this event study are close to zero and statistically insignificant, suggesting that MMLs have little effect on the birth rate. The two-way fixed effects estimate is a statistically insignificant decrease equivalent to 0.2% of the mean. The estimate has the same sign as the estimated effect of RMLs but a smaller magnitude, which is consistent with the smaller effect of MMLs on marijuana use.

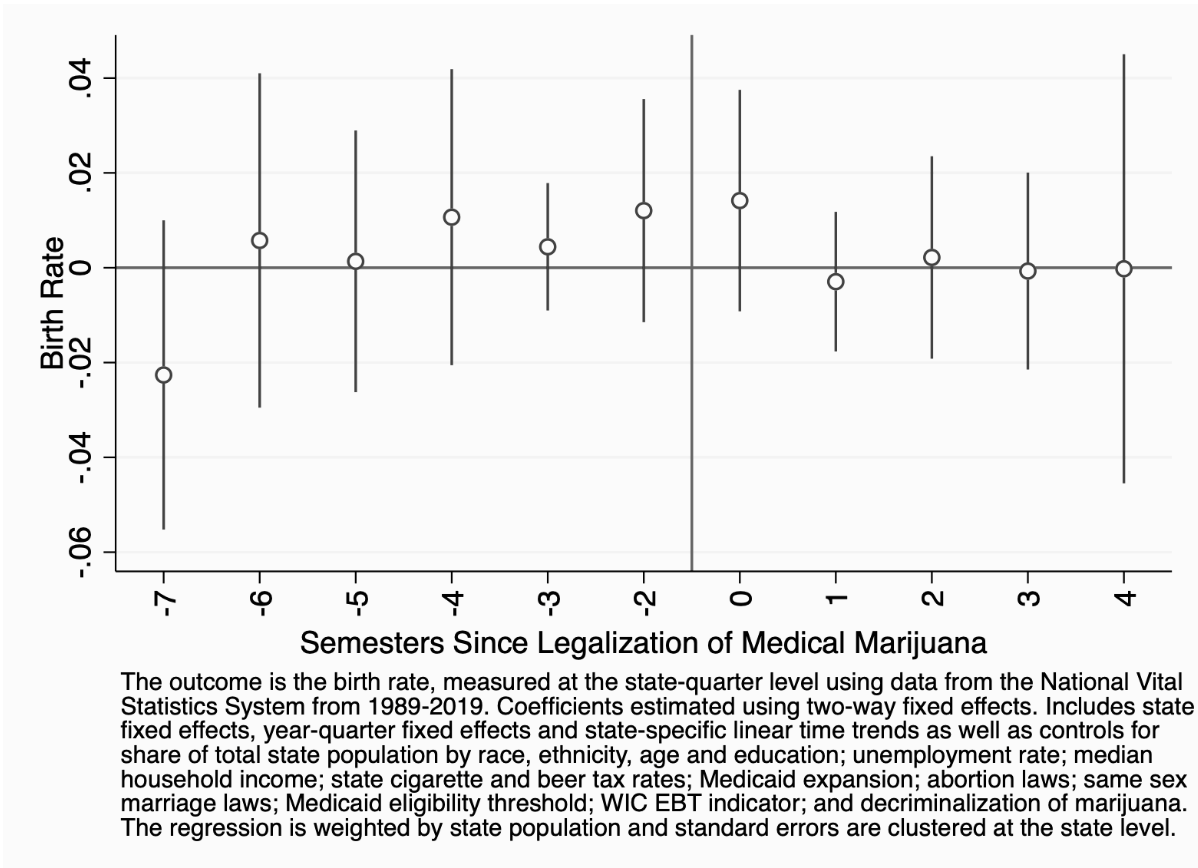

Figure 3.9: MML Event Study for Birth Rate (Women 21-30)

Previous literature has found that MMLs increase the birth rate [93]. This paper expands on previous work studying the effects of MMLs on fertility by including additional years of data and using the estimator from [90] to determine whether results are biased by negative weights. I examine several potential sources of the difference between this paper's result and the result from the previous literature and conclude that bias from negative weights most likely drove the finding of an increase in the birth rate in the previous literature. Bias from negative weights is diminished in the longer sample used for this paper. A detailed discussion of this investigation is available in the Appendix.

Finally, I examine how MMLs affect sexual activity and gonorrhea cases. Figure 3.10

displays the event study for the effects of MMLs on the number of sexual encounters in the past month. Pre-treatment estimates are statistically insignificant and close to zero, which provides support for the parallel trends assumption, although one pre-treatment estimate has a large standard error. All but one of the dynamic effects are positive, and one is statistically significant, but the dynamic effect estimates are imprecise. This event study suggests that MMLs increase the number of sexual encounters per month. The two-way fixed effects estimate displayed in Table 3.6 supports this conclusion: MMLs are estimated to increase the number of sexual encounters in the past month by 1.56, and this estimate is statistically significant at the 5% level.

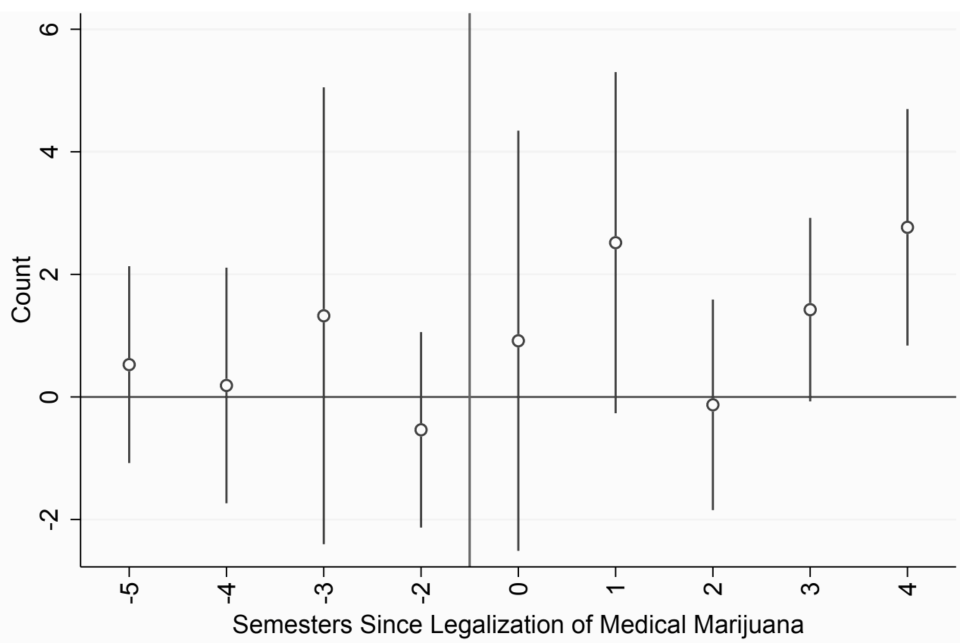

Number of sexual encounters in the last month is the outcome, measured annually at the individual level using data from the National Longitudinal Survey of Youth from 1997-2017. Coefficients are estimated using two-way fixed effects. Includes state fixed effects, year-quarter fixed effects and state-specific linear time trends as well as controls for share of total population by race, ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; decriminalization of marijuana; and the respondent's race, age, sex, education and marital status.

Figure 3.10: MML Event Study for Number of Sexual Encounters Last Month

Figure 3.11 displays the event study for the effect of MMLs on gonorrhea cases. All pre-treatment estimates are very close to zero, providing support for the parallel trends assumption. The dynamic effects suggest that MMLs lead to a slight increase in gonorrhea cases: all but one of the dynamic effect estimates are positive, and three are larger in magnitude than any pre-treatment estimates. However, the event study estimates are imprecise. Table 3.6 shows the two-way fixed effects estimate: an increase of 0.75 cases per 100,000 population, which is statistically insignificant with a large standard error. While this estimate has the same direction as the estimated effect of RMLs, the magnitude of the estimated effect of RMLs is approximately eight times larger.

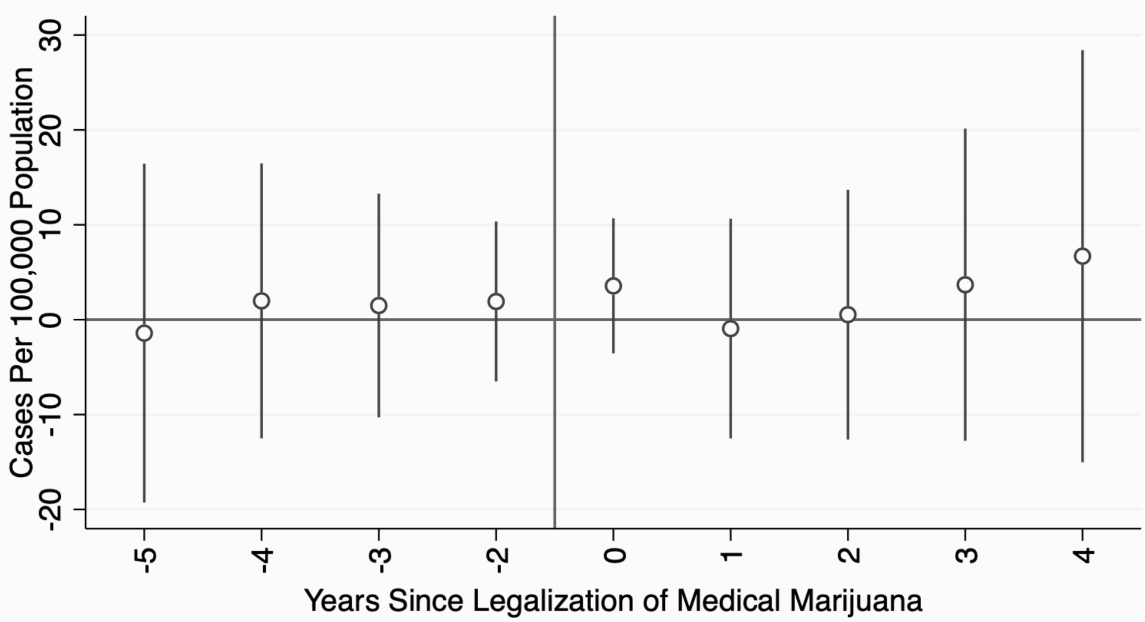

The outcome is the number of gonorrhea cases among men per 100,000 population, measured at the state- year level using data from the Centers for Disease Control and Prevention from 2000-2019. Coefficients estimated using two-way fixed effects. Includes state fixed effects, year fixed effects and state-specific linear time trends as well as controls for share of total state population by race,ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; and decriminalization of marijuana. The regression is weighted by state population and standard errors are clustered at the state level.

Figure 3.11: MML Event Study for Gonorrhea Cases

Together, these results suggest that MMLs lead to small increases in marijuana use and sexual activity but have at most a small negative effect on the birth rate.

3.7 Robustness Checks

3.7.1 The DID_M Estimator

Figure 3.12 displays the event study for the effects of RMLs on the birth rate among women aged 14-44 using the DID_M estimator from [90]. While the point estimates follow a downward trend after legalization, fewer dynamic effects are statistically significant due to the loss of efficiency from using this estimator. Estimated average effects using the DID_M estimator, which are the most comparable to the two-way fixed effects estimate, appear in Table 3.4 and are also negative, similar in magnitude to the two-way fixed effects estimates, and within the 95% confidence interval of the two-way fixed effects estimates. The DID_M estimates are statistically insignificant due to the estimator's inefficiency. The similarities between the DID_M and two-way fixed effects estimates support the use of the two-way fixed effects estimates as the main results of the paper.[7] The similarities between the estimates also provide evidence that the two-way fixed effects estimates of the effects of RMLs are not substantially biased by contamination from MMLs, as the DID_M estimator is robust to contamination effects.

[7] The event study using the DID_M estimator for the effect of RMLs on the birth rate to women 21-30 appears in the Appendix and shows a nearly identical pattern. Similarly, results for the effect of MMLs on the birth rate are robust to use of the DID_M estimator. Details of this robustness check for the effects of MMLs are provided in the Appendix.

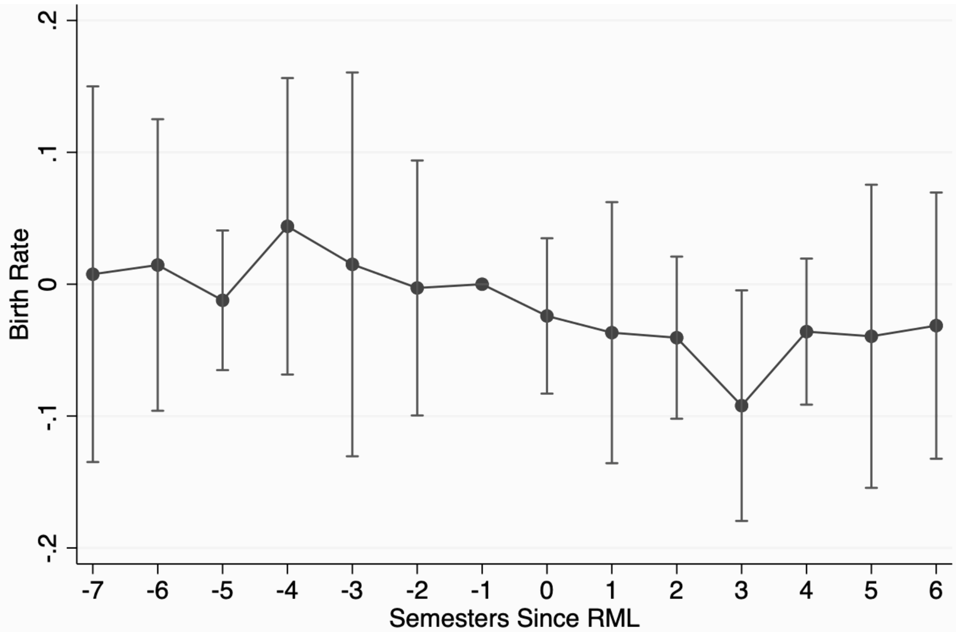

Birth rate is the outcome, measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. Coefficients are estimated using DID_M. Includes state fixed effects and year-quarter fixed effects as well as controls for share of the total population by race, ethnicity, age and education, as well as unemployment rate, median household income, state cigarette and beer tax rates, Medicaid expansion, abortion laws, same sex marriage laws, Medicaid eligibility threshold, WIC EBT indicator, and decriminalization of marijuana.

Figure 3.12: RML Event Study for Birth Rates with DID_M Estimator (Women 14-44)

DID_M estimates of average effects and 95% confidence intervals for marijuana use results appear in Table 3.3. The two-way fixed effects estimates are in the 95% confidence intervals of the DID_M estimates for all marijuana use results.[8] For gonorrhea cases, the average effect of RMLs from the DID_M estimator is an increase of 10.8. Like the two-way fixed effects estimate, the DID_M estimate is positive and statistically insignificant. The similarities between the DID_M estimates and the two-way fixed effects estimates

[8]The DID_M estimator does not allow controls to be used with the NLSY data, nor does it allow a second treatment to be included. The estimator is also unable to compute more than three dynamic effects of RMLs, and because this method requires the number of dynamic and pre-treatment effects to be equal, only three pre-periods can be used for event studies with the DID_M estimator. Due to these limitations, substantial differences between DID_M and two-way fixed effects estimates could arise even in the absence of bias from negative weights.

indicates that the two-way fixed effects estimates are not substantially biased by negative weights.

3.7.2 Dispensary Openings

The date when marijuana laws became effective is used to define treatment in this paper. An alternative measure is the date when dispensaries opened. Dispensary opening dates were obtained from [104]. Estimated effects of recreational dispensary openings on birth rates are negative and similar in magnitude to the estimated effects of RMLs, but the estimated effect on birth rates to women 14-44 is significant only at the 10% level. The estimated effect on birth rates to women 21-30 remains nearly identical in magnitude to the main results but loses statistical significance.[9] The changes in statistical significance may result from more limited data on the effects of recreational dispensaries: four of the 11 states with RMLs either did not open dispensaries or opened them after 2019, and dispensary openings occurred much closer to the end of the sample period than RMLs, leaving few observations of birth rates post-recreational marijuana dispensary openings. Estimated effects of medical dispensaries remain statistically insignificant, although the magnitudes increase compared to the estimated effects of MMLs.

The estimated effect of RMLs on marijuana use also remain similar to the main results when dispensary opening dates are used to define treatment. I estimate that recreational dispensary openings increase days of marijuana use by 0.684. This estimate is similar in magnitude to the main result and is statistically significant at the 10% level. The estimated effect of medical dispensaries on days of marijuana use is negative and statistically insignificant. The data limitations discussed above are exacerbated for the NLSY data, which ends two years earlier than the birth data. Only five states opened recreational dispensaries during the NLSY sample period.

[9] This table of results is available in the Appendix.

3.7.3 Individual Fixed Effects with NLSY Data

Because the NLSY follows the same individual over time, it allows for the use of individual fixed effects. As a robustness check, I repeat the NLSY analysis with individual fixed effects included. When individual fixed effects are added, the effect of RMLs on days of marijuana use increases to 1.18 and remains significant at the 1% level. The estimated effect of MMLs on days of marijuana use remains positive but shrinks to 0.272 and becomes insignificant at the 5% level.[10] Estimated effects of both types of laws on any marijuana use in the past month are statistically insignificant and have even smaller magnitudes than the main results, supporting the finding that marijuana use does not increase at the extensive margin in response to legalization.

3.8 Conclusion

Marijuana legalization has many consequences that are only partially understood. This paper improves our understanding of these consequences by investigating how marijuana legalization affects fertility. I find that RMLs decrease the birth rate by 2.78%. The decline in the birth rate occurs despite an increase in sexual activity: I also find suggestive evidence that RMLs increase the probability that an individual is sexually active by 3.6 percentage points. Together, these two findings demonstrate that marijuana's negative physical effects on fertility dominate its behavioral effects.

Understanding marijuana legalization's effect on fertility is important for policymakers. The decline in the national birth rate has raised economic concerns, and as more states pass RMLs, they are likely to experience a decline in fertility. This decline may increase the urgency of enacting policies that enable people who would like to have children to do so. Examples of such policies could include paid maternity leave or subsidized

[10] The table of results is available in the Appendix.

childcare.

Although this paper is able to provide new insights into the relationship between marijuana legalization and fertility, some questions remain unanswered. Additional mechanisms, such as changes in the seriousness of romantic relationships when marijuana use increases and the effects of fewer people being imprisoned for marijuana possession, are promising areas for future research.

Appendix A

Appendix to Chapter 1

A.1 Distance Traveled to Walmart

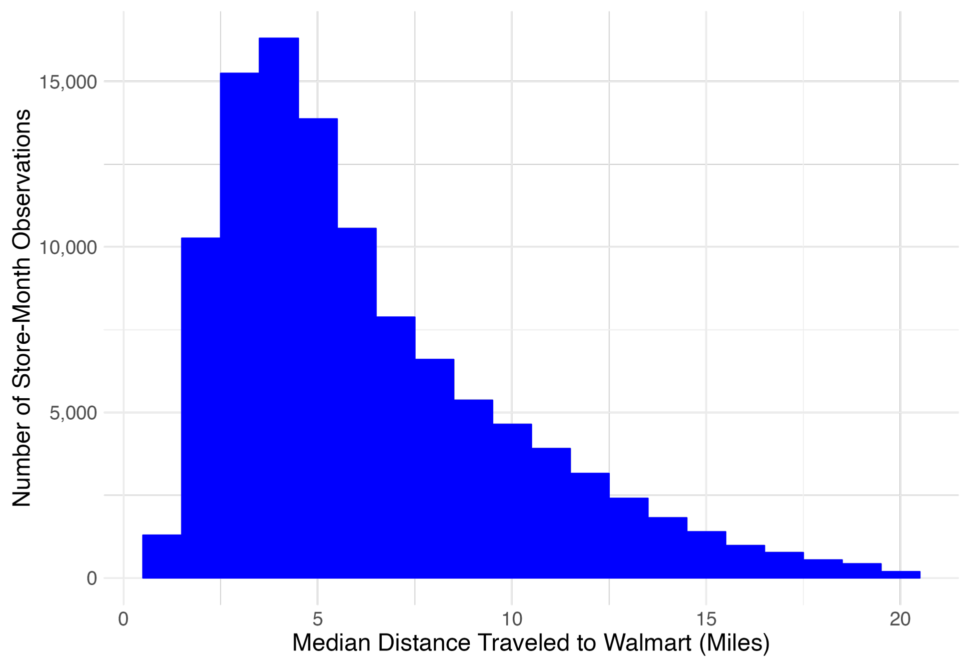

This histogram shows the distribution of distances that people traveled to shop at Walmart in the United States from 2019-2021. Data source: SafeGraph. Each observation is the median distance traveled from home to visit a Walmart store in a given month. The sample used to make this histogram is restricted to distances of twenty miles or less. The median distance traveled in the full sample is 5.37 miles.

Figure A.1: Histogram of Distance Traveled to Walmart

The identification strategy assumes that individuals who live near a Walmart store are more likely to shop at Walmart in person than people who live far away from a Walmart store. I provides descriptive evidence to support this assumption using data from SafeGraph on the distance customers travel to shop at Walmart. Each observation in the SafeGraph data is the median distance traveled to visit one Walmart store in one month. The data spans from 2019 to 2021. Figure A.1 shows a histogram of these median distances. The data used to make the histogram is restricted to median distances of less than 20 miles. The median distance in the full sample is 5.37 miles. The share of store-month level observations in each bin increases until the median distance reaches four miles and then steadily declines as distance increases, creating a right-skewed histogram. This histogram suggests that individuals who live far from a Walmart store are less likely to shop there in person than individuals who live near a Walmart store.

A.2 Additional Heterogeneity

Table 1.21 shows the results of heterogeneity analysis by marital status, education and home ownership status. I find evidence of slightly stronger effects on college- than high school-educated individuals. However, heterogeneity by education should be interpreted with caution: the education variable is missing for approximately half of the sample, and it is more likely to be missing for an individual without a college education because student loans are reported to credit bureaus. Heterogeneity by marital status does not show a clear pattern of stronger effects among single or married people. Effects among home owners are similar to the main results in direction and magnitude. I was unable to perform heterogeneity analysis on a sample of only renters due to limited data on whether an individual is a renter.

Extensive Margin Results						
	Any Credit Cards	**Any Open Balance**	**Any Monthly Payment**	**Any Past Due**	**Any Current Delinquencies**	**Any Public Records**
BNPL Access	0.00879***	0.00147	0.00324***	-0.00361***	-0.00359***	-0.00122
	(0.00129)	(0.00103)	(0.00112)	(0.000574)	(0.000580)	(0.000826)
Observations	9,929,100	9,929,100	9,929,100	9,929,100	9,929,100	9,929,100

*: p<0.1, **: p<0.05, ***: p<0.01. Clustered standard errors in parentheses. All regressions include quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and a bank desert control. Five miles from Walmart is used as the cutoff for treatment. Outcomes are indicator variables.

Table A.1: Extensive Margin Results

A.3 Extensive Margin Results and Normalization

Normalized Results						
	Number of Open Credit Cards	**Total Open Balance**	**Average Monthly Payment**	**Total Past Due**	**Number of Current Delinquencies**	**Total Public Records**
BNPL Access	0.0380***	0.0112*	-0.0310**	-0.0232	-0.0556***	-0.0112
	(0.00375)	(0.00583)	(0.0148)	(0.0669)	(0.0157)	(0.0142)
Observations	9,929,100	9,929,100	9,929,100	9,929,100	9,929,100	9,929,100

*: p<0.1, **: p<0.05, ***: p<0.01. Clustered standard errors in parentheses. All regressions include quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and a bank desert control. Five miles from Walmart is used as the cutoff for treatment. Outcomes are normalized using the pre-treatment average for the full sample.

Table A.2: Normalized Results

A.4 Restriction to Exclude Treatment Switchers

Individuals who move during the post-treatment period can experience a change in treatment status if they move from an area within five miles of Walmart to an area more than five miles from Walmart or vice versa. To address the possibility that changes in an individual's treatment status could affect the main results, I repeat the main analysis using a sample restricted to exclude individuals whose treatment status switches in the post-treatment period. Results using this restricted sample appear in Table A.3 and are nearly identical to the main results in terms of both magnitude and significance.

	Results Without Switchers					
	Number of Open Credit Cards	**Total Open Balance**	**Average Monthly Payment**	**Number of Current Delinquencies**	**Total Past Due**	**Total Public Records**
BNPL Access	0.0230***	0.0485***	0.00926	-0.00313***	-0.0263***	0.000407
	(0.00218)	(0.0110)	(0.00591)	(0.000569)	(0.00428)	(0.00569)
Observations	8,465,256	8,465,256	8,465,256	8,465,256	8,465,256	8,465,256

*: p<0.1, **: p<0.05, ***: p<0.01. Clustered standard errors in parentheses. All regressions include quarter-year, census block and individual fixed effects as well as individual-level controls for age and education and a bank desert control. Five miles from Walmart is used as the cutoff for treatment. Individuals whose treatment status switched in the post-treatment period are removed from the sample.

Table A.3: Results Excluding Treatment Switchers

Appendix B

Appendix to Chapter 2

B.1 Matching Donors

This section provides a more detailed description of the procedure for and results from using names to identify multiple contributions from the same donor. I used the donor's name to identify donations by the same individual in multiple years. Although the donor's address is available for non-voucher contributions, it is not available for voucher contributions. Furthermore, when two donors have the same name, the address does not allow me to distinguish between two different individuals at different addresses and the same individual moving to a new address. Name-based matching is imperfect, as multiple donors may have the same name or may change the way they write their name ("J. Smith", "Jen Smith" or "Jennifer Smith"), and two different voucher donors may have the same name. I removed punctuation and capitalization before matching, but that does not eliminate all error from changes in the way a donor writes her name. Another caveat is that contributions of $25 or less are sometimes reported simply as "Small Contributions", and other contributions are reported only as "Miscellaneous Receipts" or as anonymous. Therefore names are not available for all donors. An important consideration in interpreting the results of matching is that using only names leads to an artificially large number of matches, as King County has a large enough population for individuals to share names.

I identified 133,952 individuals who made a non-voucher donation between 2009 and 2021. I will refer to these donations as "cash" donations. I identified 74,151 individuals who made a voucher donation between 2017 and 2021. Using names to match voucher and cash donors, I identified 61,799 individuals who made both a cash donation and a voucher donation between 2009 and 2021. King County has a population of 2.252 million people, indicating that 9.2% of the population ever donated to a campaign between 2009 and 2021.

I began by examining whether individuals tend to donate in multiple elections. Of the seven elections in my data, the average cash donor gave a donation in 1.29 elections. Of the three elections when Democracy Vouchers were available, the average voucher donor donated a voucher in 1.27 elections. Even among the 9.2% of the population that is politically active enough to ever donate to a campaign, donating in multiple elections is a rare behavior.

The data set that matches cash and voucher donors allows me to examine how cash donations changed when an individual who had donated cash in the past started donating vouchers. 88% of individuals who donated both cash and vouchers at some point from 2009-2021 had never donated before 2017. 7,171 individuals who had donated cash before 2017 also donated Democracy Vouchers. For these individuals, the average contribution barely increased from \$141.23 to \$145.59 when vouchers became available. Therefore this group does not descriptively appear to drive the decrease in contributions over \$100, although it is possible that their cash donations would have been larger from 2017-2021 if Democracy Vouchers had not become available.

For the 54,628 individuals who donated both vouchers and cash from 2017-2021 and had never donated before 2017, the average cash donation from 2017-2021 was \$45.74 and the average voucher donation was \$41.83. Although these individuals had never donated cash before Democracy Vouchers were available, on average they donated slightly more cash than vouchers once vouchers became available.

This descriptive evidence might suggest a "crowding in" effect in which these individuals would not have donated cash from 2017-2021 in the absence of the Democracy Vouchers program. However, the fact that the typical donor only donated in one election from 2009-2021 casts doubt on this idea. If these individuals are politically engaged enough to care about donating vouchers, they may have been inspired to donate cash for the first time at some point between 2017 and 2021 even in the absence of vouchers. In fact, they may have even made larger cash donations in the absence of the program, so that donations would have been crowded out rather than crowded in.

The final group of Seattle donors consists of individuals who donated cash at some point between 2009 and 2021 but never gave a voucher donation. For these individuals, the average donation appears unchanged by the Democracy Vouchers program: their average donation was \$44.29 from 2009-2015 and \$45.22 from 2017-2021. As with the other two groups of donors, it is possible that their donations from 2017-2021 would have been larger or smaller in the absence of the voucher program, so a causal effect cannot be determined from this descriptive evidence.

For donors in King County outside Seattle, the average donation per donor per election increased from \$50.14 from 2009-2015 to \$62.81 from 2017-2021, a 25.3% increase. If the individuals who donated both vouchers and cash before 2017 increased their donations by 25.3%, their average cash donation would have been \$176.96 from 2017-2021 instead of \$145.59.

B.2 Checking for Spillovers

Other cities in King County besides Seattle are used as a control group for Seattle in this paper. One possibility is that the Democracy Vouchers program leads to spillovers, in which the Democracy Vouchers program affects contribution behavior in the control cities. One type of spillover would occur if the program affects the likelihood that individuals outside Seattle donate to campaigns in Seattle rather than in their own city, or the likelihood that individuals in Seattle donate to campaigns outside Seattle. I investigate the extent of these spillover effects before and after the voucher program went into effect. I find that these effects are small and unlikely to drive the main results.

I use the individual contribution data from 2009-2021 to examine the extent of spillovers. Each observation in this data set is one donation by one individual to one campaign in one election year. The non-voucher donation data includes the donor's home city for 638,794 observations. I removed contributions from outside King County and contributions from businesses. I categorized a contribution as an "inflow" if it came to a candidate in Seattle from a donor outside Seattle, and an "outflow" if it came to a candidate outside Seattle from a donor in Seattle. I kept only observations with a unique combination of donor city, recipient city, donor name, and election year to avoid distortions from multiple donations by the same individual. This process left me with 288,064 observations.

Over the 2009-2021 period, 97% of donors to Seattle candidates lived in Seattle. Of all contributions in the sample, 2.1% were outflows from Seattle and 3.1% were inflows to Seattle. In dollars, 5.3% of all dollars contributed were inflows to Seattle and 3% of all dollars contributed were outflows from Seattle. These statistics show that out-of-city contributions did not exert a major influence on campaign finance in Seattle during this time period.

How did these percentages change when Democracy Vouchers became available? Before Democracy Vouchers, 10.6% of all dollars contributed were inflows to Seattle. When vouchers were available, 2.8% of all dollars contributed were inflows to Seattle. The percent of all dollars contributed that were outflows from Seattle remained fairly stable over this period, slightly decreasing from 3.1% to 3% when vouchers became available. While I do find some evidence of spillover donations across city boundaries, these spillovers account for a very small percentage of overall contributions and are therefore unlikely to be a major driver of the main results.

B.3 Number of Candidates and Number of Donors

The Democracy Vouchers program has the potential to affect the number of candidates running for office. An individual who is interested in holding public office may be more likely to campaign when Democracy Vouchers are available, as they provide a source of funding for candidates without connections to wealthy donors. At the same time, well-connected individuals and incumbents likely still have better chances of attracting Democracy Voucher donations due to name recognition. Potential candidates for office may be skeptical that Democracy Vouchers will level the playing field enough for them to be competitive. Furthermore, the decision of whether to run for office may be driven more by other factors than by whether the individual has access to a donor base.

The program could also affect the number of individuals who donate to campaigns. The Democracy Vouchers program could inspire individuals who have not donated in the past to make a donation, either by giving vouchers, a cash donation, or both. The matching exercise above indeed shows that many voucher donors had never given to a campaign before vouchers were available, and that many of them made cash donations after donating vouchers for the first time.

Table 9 in the main text shows the program's estimated effects on the number of campaigns and the number of donors. I estimate a statistically insignificant decrease of 1.5 candidates running for office per election when Democracy Vouchers became available. However, this estimate should be interpreted with caution, as event studies cast doubt on the plausibility of the parallel trends assumption for this outcome. The estimated effect on the number of donors is an increase of 38,140 donors per year, which represents a 315% increase from the pre-treatment mean. This estimate is consistent with descriptive evidence that the number of donors increased when vouchers became available, but the event study for this outcome also does not support a causal interpretation of the estimate.

B.4 Census Tract Level Analysis

An alternative to measuring campaign finance outcomes at the city level is to measure them at a smaller geographic level, such as census tract or precinct. I chose the census tract level because it allows me to use census tract-level demographic controls without estimating precinct-level demographics. A limitation of using geographic variation below the city level is that Democracy Voucher data does not include the contributor's address, and therefore the tract-level analysis excludes Democracy Voucher donations. Political participation outcomes are not included in this section because they are already measured at the precinct level rather than the city level.

Outcome	DiD Estimate	Ferman & Pinto P-Value	N
Under \$100	0.1531	.7918	18,594
Over \$100	-0.7959	.0464	18,222
Under \$250	-0.3340	.6739	18,594
Over \$250	-0.8456	.0002	17,874
Between \$100 and \$250	-0.5751	.0012	17,833
Total Contributions	-0.7460	.4209	18,597

All dependent variables are logged. Analysis is at the census tract level. Census tract-level demographic controls are included in all regressions. The unit of analysis is the city.

Table B.1: Census Tract Level Results

Table B.1 shows that as with the main results, I find decreases in large contributions: the estimated effects on contributions over \$100, contributions over \$250, and contributions between \$100 and \$250 are negative and statistically significant. However, I do not find increases in small or total contributions with census tract-level analysis. Estimated coefficients are negative for total contributions and contributions under \$250, with a positive estimate for contributions under \$100. None of these three estimates are statistically significant. This difference from the main results can likely be explained by the exclusion of Democracy Vouchers from this analysis, since Democracy Vouchers account for much of the increase in small and total donations in the main analysis. Figure **??** shows the results of event studies using the census tract-level data. Despite the larger sample size, census tract-level data does not provide tighter confidence intervals in the event studies than the city-level data, and census tract-level event studies also show weaker evidence of parallel trends.

One advantage of census tract-level analysis is the possibility of propensity score trimming. Seattle differs demographically from the other cities in King County, and propensity score trimming can increase the plausibility of the parallel trends assumption

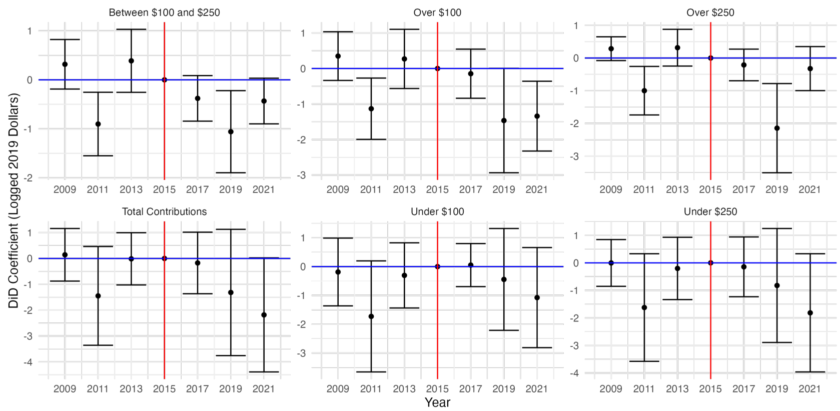

Figure B.1: Census Tract Level Event Studies

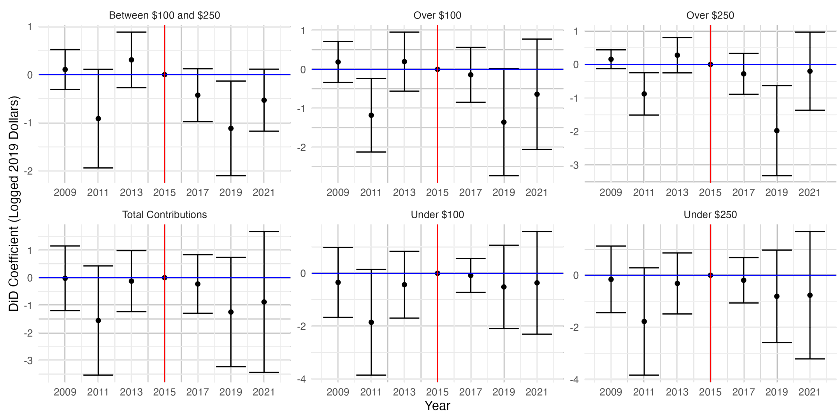

Figure B.2: Census Tract Level Event Studies (Propensity Score Trimmed)

by creating a sample of treated and control census tracts that are more demographically similar. Following the procedure established in [50], I began by estimating each census tract's propensity to be treated using a probit regression. Predicted values from this probit regression are the propensity scores, and I removed census tracts with scores above 0.9 or below 0.1. This procedure reduced the sample size for each regression by approximately 4,000 observations. Event studies using the propensity score-trimmed sample do not show stronger evidence in favor of the parallel trends assumption, as seen in Figure **??**.

B.5 Synthetic Control: Brodersen et al. (2015) Approach

I use the synthetic control approach from [78] as an alternative estimation strategy. I chose the [78] method for two reasons. First, it does not require the treated unit's outcome to be in the convex hull of the control units' outcomes. The convex hull requirement is common in other synthetic control methods and is not met in my setting because Seattle is the largest city in King County and therefore often has a larger value for an outcome than any of the control cities. Second, the [78] method does not rely on asymptotic results and is therefore useful in a setting with few pre-treatment periods. I only have four pre-treatment periods and therefore cannot rely on a method that relies on asymptotic results. As for the main results, I use data from the odd-numbered years from 2009-2021.

Outcome	Absolute Effect	Lower Bound	Upper Bound
Under $100	2.24	-1.44	6.35
Over $100	-0.31	-2.80	1.98
Under $250	1.44	-1.45	4.74
Over $250	-0.50	-2.34	1.87
Between $100 And $250	-0.10	-2.36	2.11
Total Contributions	-0.07	-2.39	2.78
Total Expenditures	0.02	-2.45	3.34
Registration	13.53	-128.79	141.21
Turnout	3.53	-16.71	21.23

All dependent variables are logged except voter registration and voter turnout. Effects are estimated using the synthetic control approach from Brodersen et al. (2015). Upper and lower bounds represent the bounds of the 95% credibility interval.

Table B.2: Synthetic Control Results

Table B.2 shows the results. The directions of the synthetic control estimates are the same as the directions of the main results for all outcomes except total contributions. However, zero is in the 95% credibility interval for all outcomes, and the credibility intervals are very wide. Although the estimated effect on total contributions is negative, the point estimate from the main results is within the 95% credibility interval.

I use placebo tests to determine whether the synthetic control provides a reliable counterfactual for Seattle. For each placebo test, I restrict the sample to the years before treatment (2009-2015) and artificially move the treatment from 2017 to 2015. I chose 2015 because the [78] method requires a minimum of three pre-treatment periods to provide a synthetic control estimate. Because treatment did not occur in the 2009-2015 period, the synthetic control should estimate a precise zero in the placebo test. A statistically significant result or a wide credibility interval containing zero would suggest that the synthetic control is not a reliable counterfactual for the treated group. The placebo test results appear in Table **??**.

Outcome	Absolute Effect	Lower Bound	Upper Bound
Under $100	1.89	-5.82	9.85
Over $100	0.63	-4.27	5.86
Under $250	1.46	-5.44	7.97
Over $250	0.48	-4.27	4.90
Between $100 And $250	0.90	-4.59	5.63
Total Contributions	1.09	-4.75	7.72
Total Expenditures	1.11	-5.14	7.58
Registration	-7.15	-193.63	133.38
Turnout	-7.35	-38.58	22.02

All dependent variables are logged except voter registration and voter turnout. This table displays the results of placebo tests for the synthetic control approach from Brodersen et al. (2015). Upper and lower bounds represent the bounds of the 95% credibility interval. Three pre-treatment periods and one post-treatment period are used for the placebo tests.

Table B.3: Synthetic Control Placebo Test Results

While zero is in the 95% credibility interval for all placebo tests, the estimates are not precise zeros: the point estimates are not close to zero and the credibility intervals are wide. These placebo tests show that the synthetic control is not a reliable estimate of the treated group's outcomes in the absence of treatment. The reason for the synthetic control's poor performance in this case is most likely the short pre-treatment period. Although the [78] approach does not rely on asymptotic results, its performance improves as the number of pre-treatment periods increases.

B.6 Conventional Synthetic Control

As an additional robustness check, I implement synthetic control using the conventional method from [79]. To address the problem of Seattle's outcomes being outside the convex hull of the control units' outcomes for the campaign finance outcomes, I index the campaign finance data using the 2009 value of each outcome. Although Seattle's outcomes are in the convex hull of the control units' outcomes when the data is indexed, indexing does not resolve the issue of the small number of pre-treatment time periods.

Unlike the [78] approach, the approach from [79] does rely on asymptotic assumptions regarding the the number of pre-treatment time periods and is therefore unlikely to perform well in this setting.

For each campaign finance outcome, I begin the indexing process by dropping cities with a value of zero for that outcome in 2009. I drop cities with a value of zero in 2009 because indexing requires dividing post-2009 values of the outcome by the 2009 value. After dropping these cities, I have a sample of 21 cities for all campaign finance outcomes.

The political participation outcomes, voter registration and voter turnout, are measured at the precinct level rather than the city level. Because synthetic control is designed for a single treated unit, I use the average values of voter registration and voter turnout in a given city-year as my outcomes for synthetic control analysis. Average voter registration and voter turnout in Seattle are in the convex hull of the donor pool in every election from 2009 to 2021, and therefore I do not index these outcomes. Because I do not drop any cities in King County from the political participation analysis, I have a sample of 39 cities for both voter registration and voter turnout.

Table 4: Conventional Synthetic Control

Outcome	Average Effect	P-Value
$100 and Under	5.84	0.15
Over $100	-0.87	0.19
$250 and Under	2.56	0.10
Over $250	-1.24	0.38
Between $100 and $250	-0.42	0.19
Contributions	-0.16	0.43
Expenditures	0.23	0.48
Voter Registration	19.12	0.15
Voter Turnout	3.57	0.79

All dependent variables are indexed using 2009 values except voter registration and voter turnout. Effects are estimated using the synthetic control approach from Abadie and Gardeazabal (2003).

Table B.4: Conventional Synthetic Control Results

Table B.4 displays the results. The estimated effects are consistent with both the results obtained using the [78] approach and the main results of the paper. I estimate substantial increases in small contributions, decreases in large contributions, and increases in voter registration and voter turnout. The estimated effects on voter registration and turnout, in particular, have very similar magnitudes whether they are obtained using the [54] method, the [78] method or the [79] method.

Although p-values are displayed alongside the estimates in Table B.4, statistical significance is very difficult to obtain due to the relatively small size of the donor pool. The calculation of p-values for conventional synthetic control involves ranking the treated unit in the distribution of the control units. For the campaign finance outcomes, with only 21 cities in the data, Seattle would have to be first or last in the distribution to achieve a p-value of 0.05. For the political participation outcomes, Seattle would have to be in the top two to achieve a p-value of 0.05. Therefore this robustness check is more useful for assessing estimated effects than for determining statistical significance.

As for the [78] approach, I use placebo tests to assess whether the synthetic control method from [79] performs well in this setting. Placebo tests that show precisely estimated null effects would indicate that synthetic control performs well and can be reliably used for causal inference.

Outcome	Placebo Effect	P-Value
$100 and Under	0.41	0.30
Over $100	0.19	0.33
$250 and Under	0.18	0.48
Over $250	0.20	0.29
Between $100 and $250	0.07	0.52
Contributions	0.40	0.43
Expenditures	0.54	0.33
Voter Registration	-2.10	0.36
Voter Turnout	8.58	0.77

All dependent variables are indexed using 2009 values except voter registration and voter turnout. Effects are estimated using the synthetic control approach from Abadie and Gardeazabal (2003). Three pre-treatment periods and one post-treatment period are used for the placebo tests.

Table B.5: Conventional Synthetic Control Placebo Test Results

Table B.5 shows the results of the placebo tests. The placebo tests do not produce estimates close to zero for any outcome except voter registration, indicating that the conventional synthetic control method does not perform well in this setting. The most likely explanation for the method's poor performance is the small number of pre-treatment periods.

B.7 Cities Included in Each Control Group

The political participation results include all 39 cities in King County. The pool of potential control cities for the campaign finance outcomes consists of the 31 cities aside from Seattle that hold municipal elections in odd-numbered years, with the exact control group varying by outcome as described in the data section. Here I present a table showing which cities were included in the control group for each campaign finance outcome.

City	Total Contributions	Total Contributions (City Council)	Total Expenditures	Total Expenditures (City Council)	$100 and Under	$100 and Under (City Council)	$250 and Under	$250 and Under (City Council)	Between $100 and $250	Between $100 and $250 (City Council)	Over $100	Over $100 (City Council)	Over $250	Over $250 (City Council)
Auburn	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Bellevue	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Black Diamond	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Bothell	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Burien	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Carnation	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Covington	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Des Moines	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Duvall	X		X		X		X		X		X		X	
Enumclaw	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Federal Way	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Issaquah	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Kenmore	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Kent	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Kirkland	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Lake Forest Park	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Maple Valley	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Medina	X	X	X	X							X	X	X	X
Mercer Island	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Newcastle	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Normandy Park	X	X	X	X	X	X	X	X	X	X	X	X	X	X
North Bend	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Pacific	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Redmond	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Renton	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Sammamish	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Seatac	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Shoreline	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Snoqualmie	X		X		X		X		X		X		X	
Tukwila	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Woodinville	X	X	X	X	X	X	X	X	X	X	X	X	X	X

X indicates that the city was in the control group for the outcome. Political participation outcomes are not listed because all 39 cities in King County were included in analyses for those outcomes.

Table B.6: Cities in Each Control Group

Appendix C

Appendix to Chapter 3

C.1 Additional Robustness Checks

C.1.1 The DID_M Estimator for MMLs

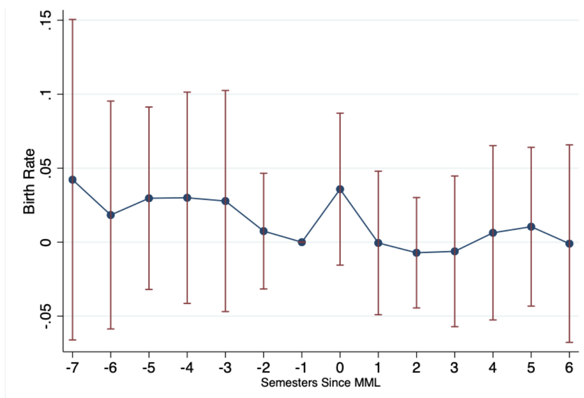

Birth rate is the outcome, measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. Coefficients are estimated using DIDM. Includes state fixed effects and year-quarter fixed effects as well as controls for share of total population by race, ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; and decriminalization of marijuana.

Figure C.1: MML Event Study for Birth Rates with DID_M Estimator (Women 14-44)

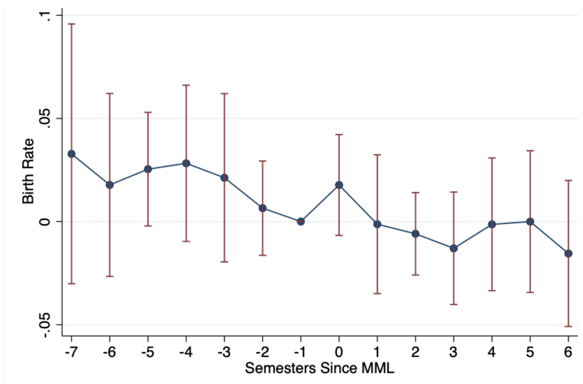

Birth rate is the outcome, measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. Coefficients are estimated using DIDM. Includes state fixed effects and year-quarter fixed effects as well as controls for share of total population by race, ethnicity, age and education, as well as unemployment rate, median household income, state cigarette and beer tax rates, Medicaid expansion, abortion laws, same sex marriage laws, Medicaid eligibility threshold, WIC EBT indicator, and decriminalization of marijuana.

Figure C.2: MML Event Study for Birth Rates with DID_M Estimator (Women 21-30)

Figure C.1 and Figure C.2 display the results of event studies using the DID_M estimator to estimate the effects of MMLs on birth rates. Like the standard event studies, these do not show statistically significant dynamic effects. The estimated average effects using DID_M are negative and statistically insignificant, as shown in Table 4; the DID_M estimate and the two-way fixed effects estimate for women 21-30 are very similar in magnitude, while the DID_M estimate for women 14-44 is closer to zero than the two-way fixed effects estimate. Results from both DID_M and two-way fixed effects lead to the conclusion that MMLs have at most a small negative effect on births.

C.1.2 Wild Cluster Bootstrap

	Estimate	P-Value
Women 14-44		
Any RML Effective	-0.097	0.048
Any MML Effective	-0.027	0.600
Women 21-30		
Any RML Effective	-0.057	0.013
Any MML Effective	-0.004	0.796

The dependent variable is the birth rate, measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. Standard errors obtained with the wild cluster bootstrap, with clustering at the state level. Regressions include state and year-quarter fixed effects and a state-specific time trend. Regressions are weighted by state population. Regressions include controls for share of total population by race, ethnicity, age, and education, decriminalization of marijuana, unemployment rate, median household income, state cigarette and beer tax, Medicaid expansion and abortion laws, same sex marriage laws and a WIC EBT indicator.

Table C.1: Birth Rate Results with Wild Cluster Bootstrap

Because relatively few states implemented RMLs, I repeat my analysis using the wild cluster bootstrap for inference. The findings support the main results. The results for birth rates appear in Table C.1. With the wild cluster bootstrap, estimated effects of RMLs on birth rates among women 14-44 and 21-30 are statistically significant at the 5% level, while estimated effects of MMLs on birth rates in both groups remain statistically insignificant. Results for marijuana use appear in Table C.2. The estimated effect of RMLs on days of marijuana use is significant at the 5% level when the wild cluster bootstrap is used, while the estimated effect of MMLs on days of marijuana use becomes statistically insignificant. Estimated effects of both laws on any marijuana use remain statistically insignificant.

C.1.3 Removing MML-Only States

	Days of Marijuana Use	Any Marijuana Use
Panel A: Without MML-Only States		
Any RML EFfective	0.359	-0.0384*
	(0.447)	(0.0194)
Any MML Effective	0.502*	0.0202
	(0.272)	(0.0308)
Observations	64,122	70,286
Panel B: Bootstrapped Standard Errors		
Any RML EFfective	0.762**	-0.0157
	(0.011)	(0.319)
Any MML Effective	0.425	0.00237
	(0.112)	(0.850)
Observations	102,940	112,558
Panel C: With Individual Fixed Effects		
Any RML EFfective	1.181***	-0.000804
	(0.263)	(0.0178)
Any MML Effective	0.272*	-0.00526
	(0.139)	(0.00649)
Observations	102,331	111,949
Panel D: Poisson Regression		
Any RML EFfective	0.192	
	(0.119)	
Any MML Effective	0.166**	
	(0.081)	
Observations	102,940	
Panel E: Dispensary Openings		
Recreational Dispensaries Open	0.829	0.346
	(0.631)	(0.0449)
Medical Dispensaries Open	-1.980	-0.0699*
	(1.429)	(0.0371)
Observations	7,319	10,365

*: p<0.1 **: p<0.05 ***: p<0.01. Standard errors clustered at the state level in parentheses except in Panel B, which displays p-values from the wild cluster bootstrap in parentheses. Outcomes are number of days of marijuana use per month and an indicator for any marijuana use in the last month. Outcomes are measured at the individual level using data from the National Longitudinal Survey of Youth from 1997-2017, which surveys individuals once per year. Coefficients estimated using two-way fixed effects. Regressions include state fixed effects, year-quarter fixed effects and state-specific linear time trends, as well as controls for share of total state population by race, ethnicity, age, and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; decriminalization of marijuana; and the respondent's age, sex, education and marital status. RML: Recreational Marijuana Law. MML: Medical Marijuana Law.

Table C.2: Marijuana Use Robustness Checks

	Women 14-44	Women 21-30
Panel A: 2005-2019		
Any RML Effective	-0.045**	-0.024**
	(0.020)	(0.009)
Any MML Effective	0.004	-0.006
	(0.012)	(0.007)
Observations	2,907	2,907
Panel B: 2005-2014		
Any RML Effective	-0.019	-0.020
	(0.020)	(0.020)
Any MML Effective	0.030**	0.004
	(0.013)	(0.007)
Observations	2,040	2,040
Panel C: Dropping states with MMLs after 2014		
Any RML Effective	-0.093**	-0.047**
	(0.038)	(0.018)
Any MML Effective	-0.051	-0.007
	(0.039)	(0.018)
Observations	5,456	5,456

*: p<0.1 **: p<0.05 ***: p<0.01. The outcome is the birth rate, which is measured at the state-quarter level using data from the National Vital Statistics System. Data from 2005-2019 is used for Panel A, 2004-2014 for Panel B. Panel C uses data from 1989-2019 and drops states that legalized medical marijuana after 2014. All regressions include controls for share of total population by race, ethnicity, age and education, as well as unemployment rate, median household income, state cigarette and beer tax rates, Medicaid expansion, abortion laws, same sex marriage laws, Medicaid eligibility threshold, WIC EBT indicator, and decriminalization of marijuana. State fixed effects, year fixed effects and a state-specific time trend are included, and regressions are weighted by state population. Standard errors are clustered at the state level. Standard errors in parentheses. RML: recreational marijuana law. MML: medical marijuana law.

Table C.3: Birth Rate Robustness Checks

As an additional robustness check, I replicate the analyses on birth rates and marijuana use without the states that passed MMLs but not RMLs. This sample restriction removes 23 states, as can be seen in Table 1. The results for birth rates appear in Table C.3. The point estimates for the effects of RMLs remain negative and are similar in magnitude to the main results. Estimated effects of MMLs remain negative and statistically insignificant, but have larger magnitudes when MML-only states are dropped. The results for marijuana use appear in Table C.2. Estimated effects on days of marijuana use remain positive and similar in magnitude when MML-only states are removed from the sample. Results for both birth rates and marijuana use lose statistical significance when the MML-only states are removed. This loss of statistical significance is likely due to the decrease in power from removing nearly half of the sample.

C.1.4 Removing State-Specific Linear Trends

Next, I assess whether the main results are robust to not including state-specific linear trends. Figure C.3 and Figure C.4 show event studies without state-specific linear trends for the effects of RMLs on birth rates among women aged 14-44 and 21-30, respectively. Similar to the event studies that include state-specific linear trends, both event studies show small and statistically insignificant pre-treatment estimates that are consistent with the parallel trends assumption. Both event studies also show a downward trend in the birth rate following RML implementation. All dynamic effects in both event studies are negative, and all but one are statistically significant. Two-way fixed effects estimates without state-specific linear trends, shown in Table C.4, are very similar to the main results: estimated effects of RMLs on birth rates among women aged 14-44 and 21-30 are negative, statistically significant and similar in magnitude to the main results, while estimated effects of MMLs on birth rates among women in both groups are negative, statistically insignificant and smaller in magnitude than the effects of RMLs. These findings indicate that the main results are robust to not including state-specific linear trends.

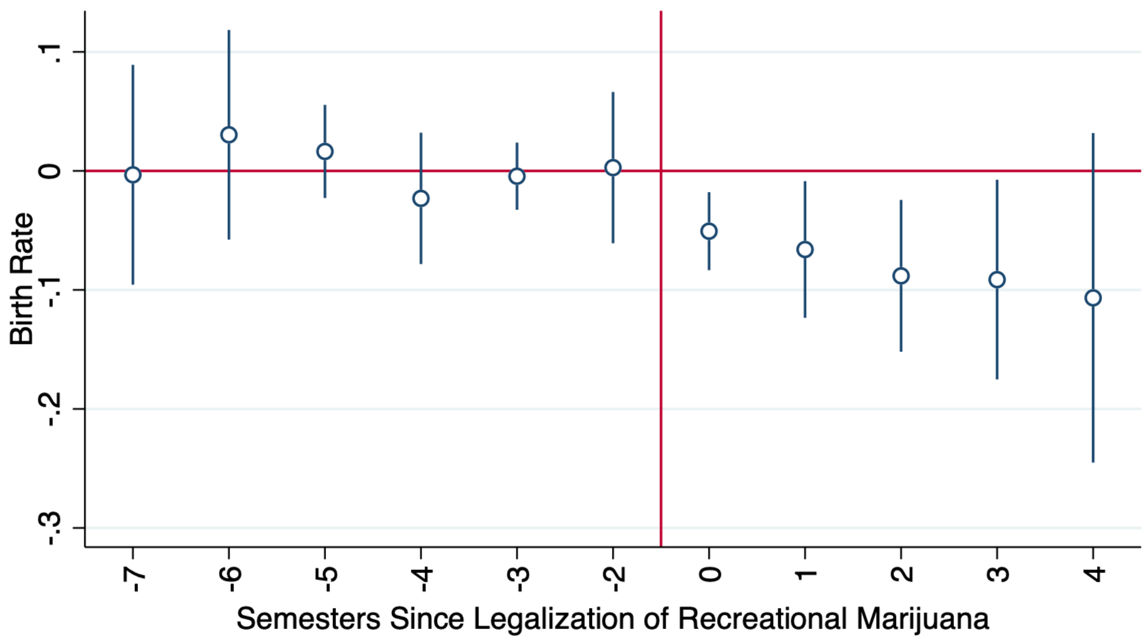

The outcome is the birth rate, measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. Coefficients estimated using two-way fixed effects. Includes state fixed effects and year-quarter fixed effects as well as controls for share of total state population by race, ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; and decriminalization of marijuana. The regression is weighted by state population and standard errors are clustered at the state level.

Figure C.3: RML Event Study for Birth Rates Without Time Trend (Women 14-44)

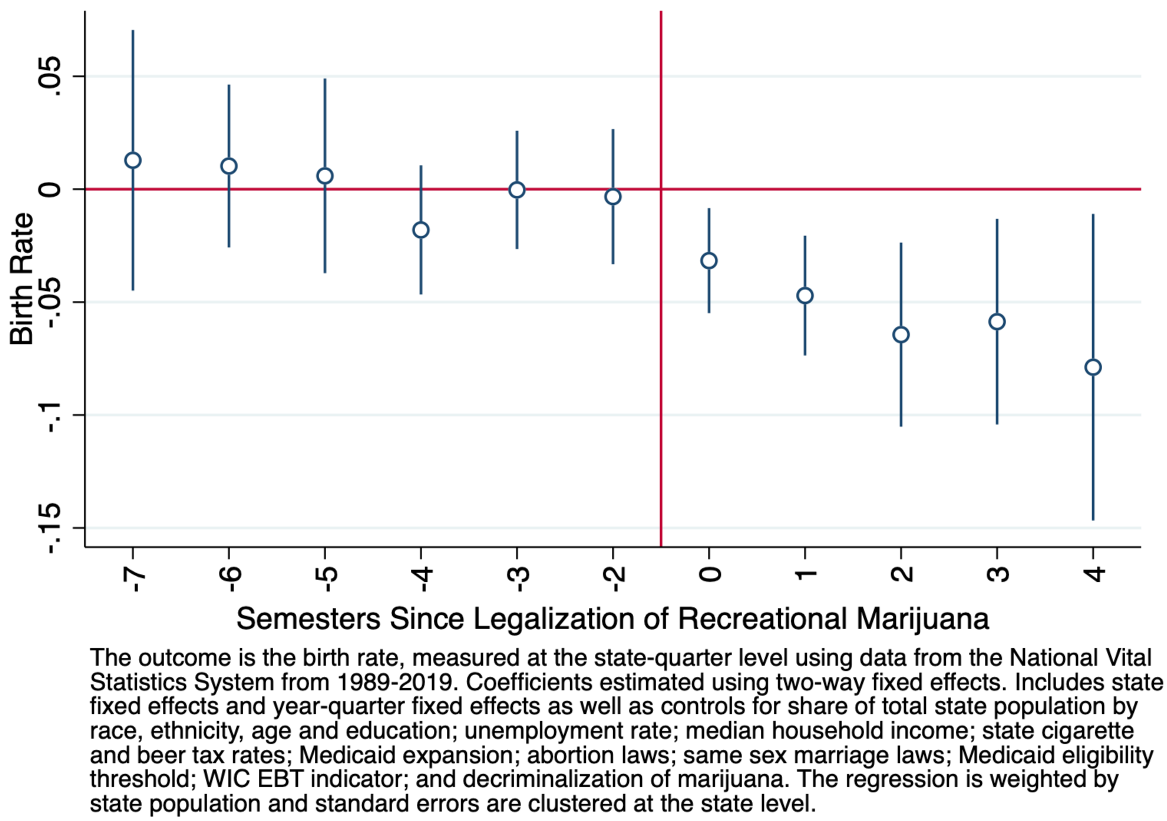

Figure C.4: RML Event Study for Birth Rates Without Time Trend (Women 21-30)

	Women 14-44	Women 21-30
Any RML Effective	-0.108*	-0.081**
	(0.058)	(0.031)
Any MML Effective	-0.051	-0.031
	(0.042)	(0.030)
Observations	6,324	6,324

Birth rate is the outcome. All regressions include controls for share of total population by race, ethnicity, age and education, as well as unemployment rate, median household income, state cigarette and beer tax rates, Medicaid expansion, abortion laws, same sex marriage laws, Medicaid eligibility threshold, WIC EBT indicator, and decriminalization of marijuana. Standard errors in parentheses. *: $p<0.1$, **: $p<0.05$, ***: $p<0.01$

Table C.4: Birth Rate Results Without State-Specific Linear Trends

C.1.5 Birth Counts

As an additional robustness check, I use a Poisson regression rather than a linear regression to estimate effects of marijuana legalization on days of marijuana use. The results appear in Table C.2. Point estimates are positive for both MMLs and RMLs. While the estimated effect of RMLs is larger in magnitude, only the estimated effect of MMLs is statistically significant at the 5% level due to a larger standard error on the estimated effect of RMLs.

C.1.6 Birth Counts

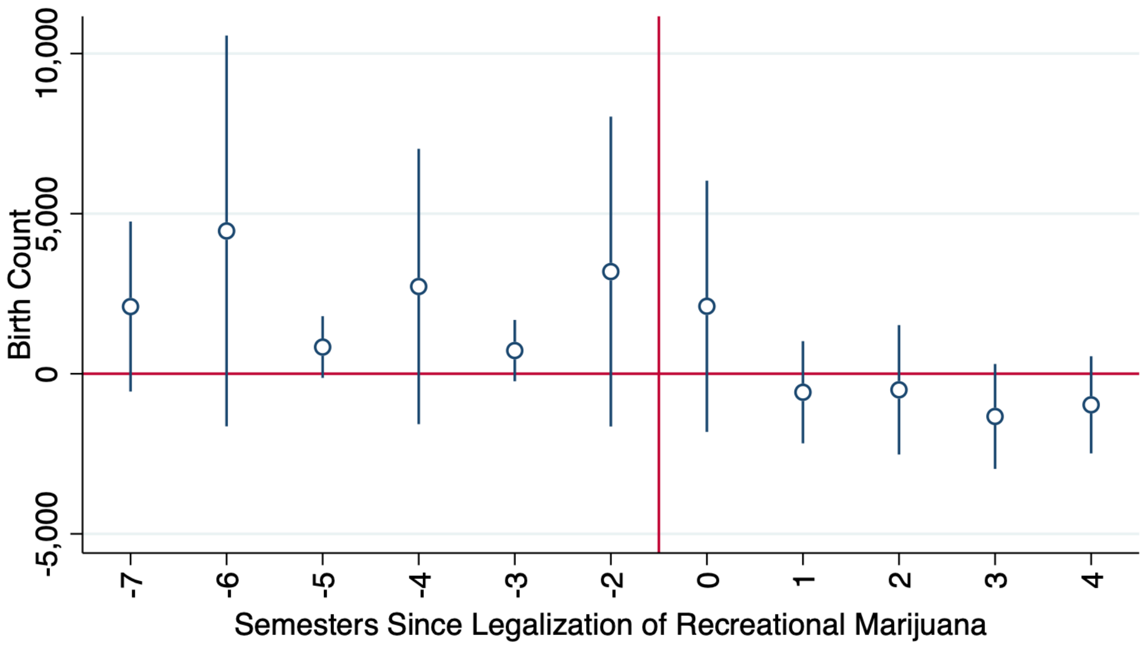

The outcome is the birth count, measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. Coefficients estimated using two-way fixed effects. Includes state fixed effects, year-quarter fixed effects and state-specific linear time trends as well as controls for share of total state population by race, ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; and decriminalization of marijuana. The regression is weighted by state population and standard errors are clustered at the state level.

Figure C.5: Event Study for the Effect of RMLs on Birth Counts

Birth rates, which are used for the main results, can suffer from measurement error as the population measure in the denominator is an estimate. I therefore use birth counts rather than birth rates as a robustness check to eliminate this measurement error. The event study for the effect of RMLs on birth rates appears in Figure C.5. None of the pre-treatment estimates are statistically significant, which is consistent with the parallel trends assumption. However, four of the pre-treatment estimates have large magnitudes. While the post-treatment estimates show a downward trend that is consistent with the main finding of a decrease in births, the post-treatment estimates are statistically insignificant. All estimates in the event study have fairly large standard errors.

The two-way fixed effects estimates appear in Table C.3. I estimate that RMLs decrease birth counts by 2,581 among women aged 14-44 and 1,912 among women aged 21-30. Both estimates are significant at the 5% level. These estimates support the main results by showing negative and statistically significant effects of RMLs on births. The average birth count among women 14-44 for states that ever implemented RMLs is 22,851, making the estimated effect on the birth count larger as a percent of the mean than the estimated effect on the birth rate. However, the standard errors are fairly large: 994 for women 14-44 and 728 for women 21-30. For women 14-44, the upper bound of the 95% confidence interval is a decrease equivalent to 2.77% of the mean, which is nearly identical to the estimated effect on the birth rate.

The imprecision of the birth count estimates demonstrates the trade-off of using birth rates rather than birth counts: while using birth counts eliminates population measurement error, birth counts vary more than birth rates and produce noisier estimates. Furthermore, the event study using birth rates provides stronger evidence in favor of the parallel trends assumption, as pre-treatment estimates in that event study are closer to zero. Birth rates are therefore preferred as the main outcome measure for this paper. Estimated effects of MMLs on birth counts are statistically insignificant and smaller in magnitude than estimated effects of RMLs, which is consistent with the main results.

C.1.7 Sexual Activity Robustness

Table C.5 shows the results of robustness checks for estimated effects on sexual activity.

	Number of Sexual Encounters Last Month	Any Sex Last Month	Sex with a Stranger
Panel A: Dispensary Openings			
Recreational Dispensaries Open	-0.266	0.00423	-0.00643
	(1.119)	(0.0121)	(0.0132)
Medical Dispensaries Open	0.0112	-0.00931	0.00833*
	(0.537)	(0.0108)	(0.00463)
Observations	112,557	112,558	82,217
Panel B: With Individual Fixed Effects			
Any RML Effective	1.127	0.0326***	0.00755
	(0.821)	(0.0109)	(0.0126)
Any MML Effective	1.523*	0.00115	-0.0135***
	(0.839)	(0.00846)	(0.00493)
Observations	111,948	111,948	81,832

*: $p<0.1$ **: $p<0.05$ ***: $p<0.01$. Standard errors clustered at the state level in parentheses. Outcomes are number of sexual encounters per month and indicators for being sexually active in the last month and having sex with a stranger. Outcomes are measured at the individual level using data from the National Longitudinal Survey of Youth from 1997-2017, which surveys individuals once per year. Coefficients estimated using two-way fixed effects. Regressions include state fixed effects, year-quarter fixed effects and state-specific linear time trends, as well as controls for share of total state population by race, ethnicity, age, and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; decriminalization of marijuana; and the respondent's age, sex, education and marital status. RML: Recreational Marijuana Law. MML: Medical Marijuana Law.

Table C.5: Sexual Activity Robustness Checks

C.2 Additional Tables and Figures

Table C.6 shows demographic differences between states that legalized medical marijuana and all other states.

	All Other States		Medical Marijuana Legalizers			
	Mean	Std. Dev.	Mean	Std. Dev.	Diff. in Means	Std. Error
Birth Rate	3.43	0.51	3.45	0.88	0.02	0.02
Days of Marijuana Use	1.60	5.88	2.03	6.57	-0.43	0.036
Any Marijuana Use	0.22	0.41	0.24	0.43	-0.03	0.002
Number of Sexual Encounters	6.06	35.44	6.29	38.80	-0.23	0.210
Sexually Active	0.46	0.50	0.45	0.50	0.01	0.003
Sex with a Stranger	0.06	0.23	0.06	0.24	0.00	0.001
Gonorrhea Cases	4,117.98	4,365.56	3,736.50	5,978.82	381.48	346.58
Total Population	5,642,151.14	4,948,802.81	5,687,617.57	7,068,295.59	45,466.43	151,611.15
Unemployment	5.28	1.79	5.65	1.89	0.37	0.05
Median Household Income	40,470.50	10,182.37	45,652.82	13,329.71	5,182.33	297.97
Percent White	83.34	10.41	81.57	15.59	-1.78	0.33
Percent Black	13.64	10.77	10.45	11.83	-3.19	0.29
Percent Native American	3.08	2.09	8.09	12.37	5.01	0.21
Percent Hispanic	7.53	8.12	12.14	11.12	4.61	0.24
Percent Aged 0-19	28.26	1.92	27.20	2.86	-1.06	0.06
Percent Aged 20-39	28.38	2.36	28.83	3.25	0.46	0.07
Percent Aged 40-64	30.21	2.86	30.63	3.44	0.42	0.08
Percent Aged 65+	13.15	1.67	13.33	2.47	0.18	0.05
Percent High School	79.81	8.35	82.07	6.87	2.27	0.20
Percent College	21.75	5.17	25.54	7.02	3.78	0.15

Table C.6: MML Summary Stats

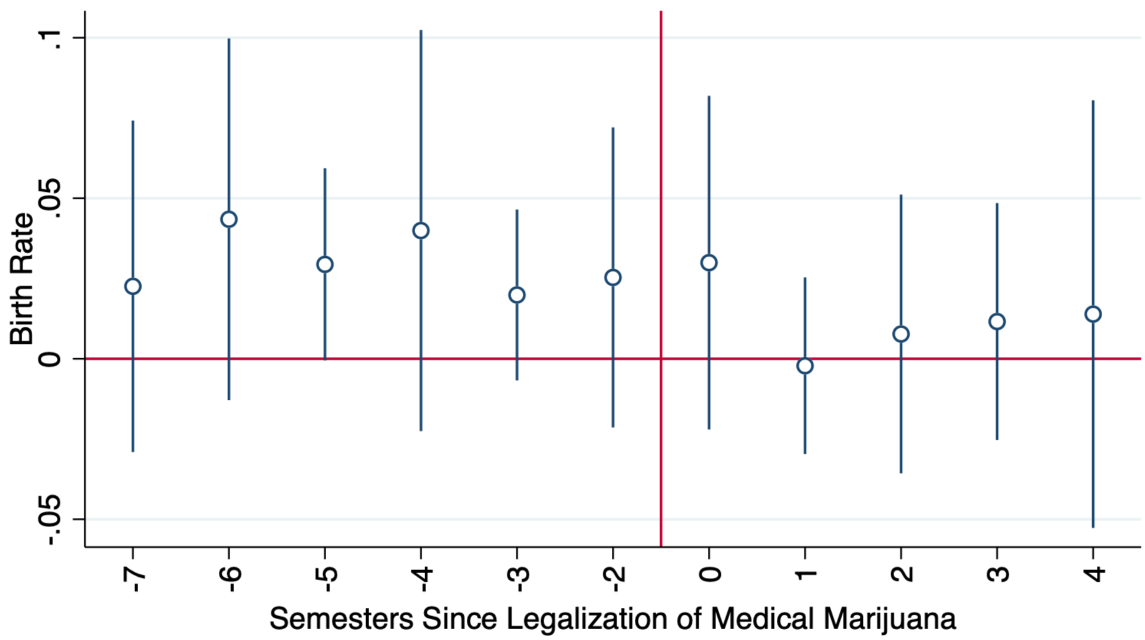

The outcome is the birth rate, measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. Coefficients estimated using two-way fixed effects. Includes state fixed effects, year-quarter fixed effects and state-specific linear time trends as well as controls for share of total state population by race, ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; and decriminalization of marijuana. The regression is weighted by state population and standard errors are clustered at the state level.

Figure C.6: MML Event Study for Birth Rate (Women 14-44)

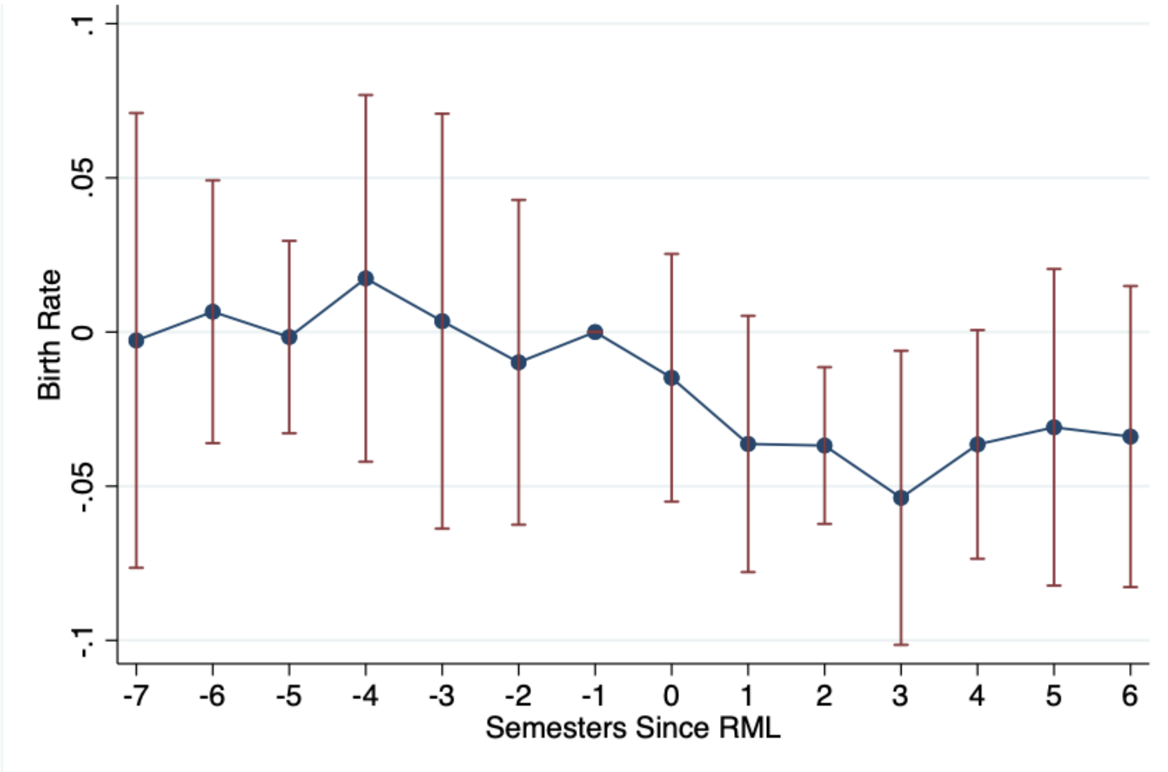

Birth rate is the outcome, measured at the state-quarter level using data from the National Vital Statistics System from 1989-2019. Coefficients are estimated using DIDM. state fixed effects and year-quarter fixed effects as well as controls for share of total population by race, ethnicity, age and education, as well as unemployment rate, median household income, state cigarette and beer tax rates, Medicaid expansion, abortion laws, same sex marriage laws, Medicaid eligibility threshold, WIC EBT indicator, and decriminalization of marijuana.

Figure C.7: RML Event Study for Birth Rates with DID_M Estimator (Women 21-30)

C.3 Additional Heterogeneity

C.3.1 Marijuana Use Heterogeneity

Table C.7 shows heterogeneity in effects of marijuana legalization on marijuana use by race. Black respondents experience the largest increase in marijuana use from RMLs, with an estimated increase of 3 days of marijuana use per month. Respondents who are neither Black nor Hispanic report the largest increase in marijuana use from MMLs: an estimated increase of 0.5 days per month.

	Black	Hispanic	Neither Black nor Hispanic
Panel A: Days of Marijuana Use			
Any RML Effective	2.951***	1.129	0.418
	(1.320)	(1.057)	(0.388)
Any MML Effective	-0.0727	-0.0533	0.529***
	(0.411)	(0.464)	(0.191)
Observations	27,053	21,259	53,627
Panel B: Any Marijuana Use			
Any RML Effective	0.002	0.007	-0.022
	(0.031)	(0.022)	(0.018)
Any MML Effective	0.002	-0.004	0.004
	(0.017)	(0.019)	(0.011)
Observations	29,741	23,193	58,525

[a] *: p<0.1 **: p<0.05 ***: p<0.01. Standard errors clustered at the state level in parentheses. Outcomes are days of marijuana use last month and an indicator for any marijuana use last month. Outcomes are measured at the individual level using data from the National Longitudinal Survey of Youth from 1997-2017. Coefficients estimated using two-way fixed effects. Regressions include state fixed effects, year-quarter fixed effects and state-specific linear time trends, as well as controls for share of total state population by race, ethnicity, age, and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; decriminalization of marijuana; and the respondent's age, sex, education and marital status.

Table C.7: Marijuana Use Heterogeneity Results

C.3.2 Birth Rate Heterogeneity by Race and Education

Table 5 shows heterogeneity by race, ethnicity and education level. As with the overall results, all the point estimates are negative. However, the effects of RMLs are not statistically significant for any of these sub-groups. The point estimates with the largest magnitudes are for Black and college-educated women, which is consistent with the fact that Black NLSY respondents experienced the largest increase in marijuana use due to RMLs. MMLs only have statistically significant effects on birth rates among white women and women who have a high school education.

C.3.3 Sexual Activity Heterogeneity

Table 6 shows heterogeneity in sexual activity by sex. This table shows that men increase their sexual activity more than women in response to marijuana legalization. Men are 4.9 percentage points more likely to be sexually active in the past month in response to RMLs, while the effect for women is statistically insignificant. MMLs increase men's average number of sexual encounters in the past month by 3.4 and their probability of being sexually active in the past month increases by 2.5 percentage points. MMLs do not have a statistically significant effect on either outcome for women. The heterogeneity analysis shows that RMLs decrease the probability of having sex with a stranger by 3.4 percentage points for men, while MMLs decrease the probability of having sex with a stranger by 0.8 percentage points for women.

C.4 Additional Outcomes

C.4.1 Effects on Fetal Deaths

In this section, I examine the effects of marijuana legalization on fetal deaths. I use data on fetal deaths from the Centers for Disease Control and Prevention (CDC). This data includes the the total number of fetal deaths at 20 weeks of gestation or more at the state-month level from 2005-2021. Fetal death is defined by the CDC as "the spontaneous intrauterine death of a fetus at any time during pregnancy" and includes stillbirths.

I aggregate the data to the state-quarter-year level to match the birth rate analysis, creating a sample of 2,741 observations. Removing data from states that did not report data in at least one quarter reduces the sample to 2,380 observations. I also remove observations from after the birth data ends, which leaves a sample of 1,995 observations from 35 states. I include the same demographic and policy controls that I use for the birth rate analysis, and I use the same regression specification. The outcome is logged number of fetal deaths.

Estimated effects of both RMLs and MMLs on fetal deaths are statistically insignificant and close to zero. The estimated effect of RMLs is a 0.8% decrease in fetal deaths, while the estimated effect of MMLs is a 1.2% decrease in fetal deaths. Compared to the average of 166 fetal deaths, these decreases represent 1.32 and 1.99 fewer fetal deaths per state-quarter-year, respectively. Even if the true effects are at the bounds of the 95% confidence intervals, they would represent changes of fewer than ten fetal deaths compared to the average. These results suggest that marijuana legalization does not affect fetal deaths.

Treatment	Estimate	SE	N
RML	-0.008	0.024	1,995
MML	-0.012	0.021	1,995

The outcome is the logged number of fetal deaths at the state-quarter-year level. Regressions include state and year-quarter fixed effects, a state-specific time trend, and controls for demographic characteristics and state policies. Regressions are weighted by state population and standard errors are clustered at the state level. The average number of fetal deaths per state-quarter-year is 166.

Table C.8: Fetal Death Results

Event studies suggest that the parallel trends assumption is plausible for fetal deaths, with zero in the 95% confidence interval for all pre-treatment estimates in both event studies. Estimated coefficients in the pre-treatment period have small magnitudes.

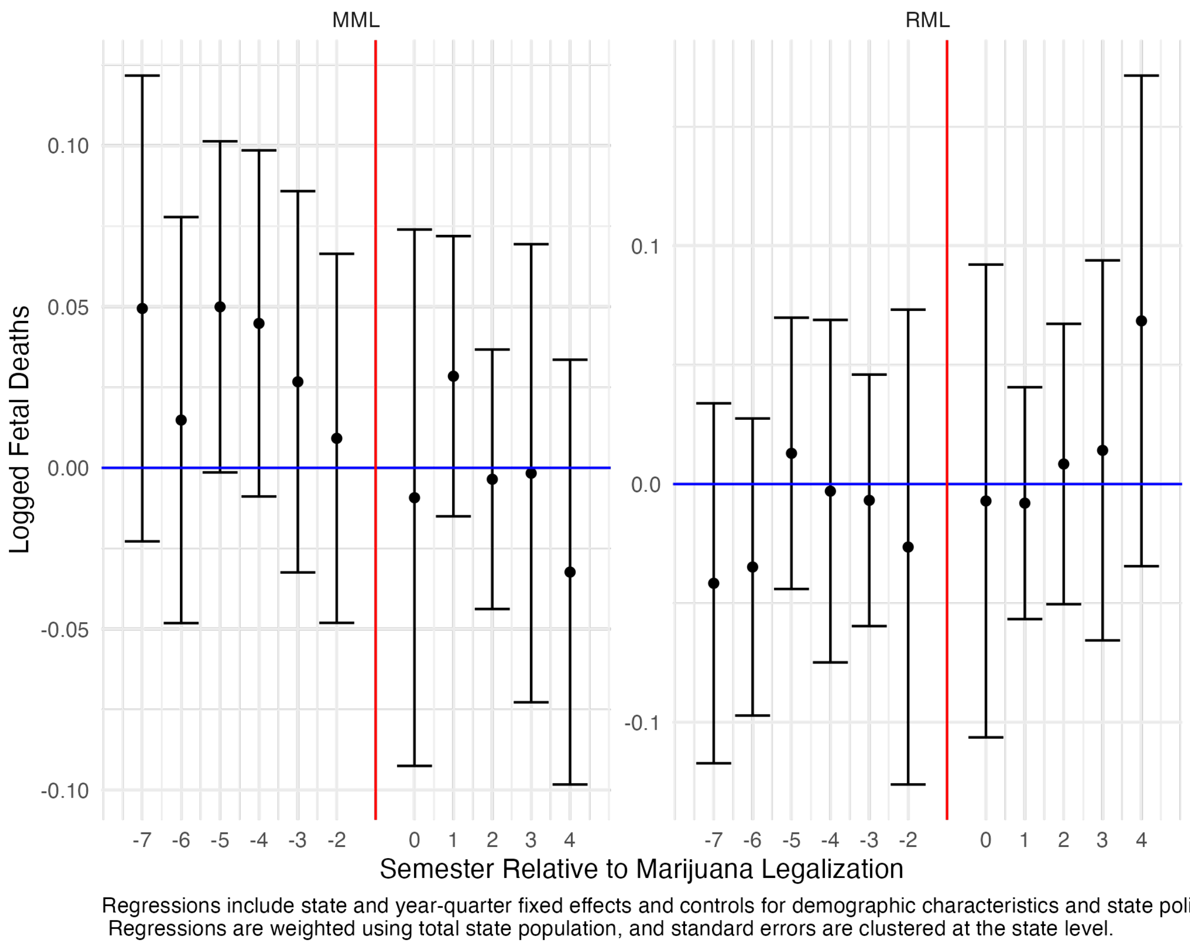

Figure C.8: Fetal Death Event Studies

C.4.2 Effects on Abortion

In this section, I estimate the effects of marijuana legalization on abortion. I use state-year level abortion counts from the CDC for the years 2011-2019. The data includes the number of abortions performed in each state as well as the number of residents of each state who received an abortion. I use the number of residents of each state who received an abortion for analysis, because women may travel to a different state to receive an abortion due to the substantial state-level variation in abortion access.

The sample includes 458 observations, one for each of the 50 states and Washington, D.C. for the years 2011-2019, with the exception of one missing observation for Colorado in 2018. I use two outcomes: the number of abortions and the abortion rate, measured as the number of abortions per 1,000 women between the ages of 14 and 44. One limitation of the data is that states and sub-state areas are not required to to report their abortion statistics to the CDC. Abortions may be missing from the data, and the extent of this reporting error likely varies across states and years.

I use the same two-way fixed effects identification strategy as for the main results. Many of the same factors that affect the birth rate, such as unemployment rate, income, and restrictiveness of abortion laws, are likely to affect the number of abortions. I therefore use the same set of controls that I use for the main results.

Table C.9 shows the results. The two-way fixed effects estimates are imprecise for both the effects of RMLs and MMLs on abortion counts. The estimated effect of RMLs is a decrease of 514 abortions, but the standard error is much larger than the estimate at 1,651. The estimated effect of MMLs is even less precise, with an estimate of 13,457 and a standard error of 11,119. Neither estimate is statistically significant. Estimated effects on the abortion rate are similarly imprecise. I estimate that RMLs decrease the abortion rate by 0.38 and MMLs increase the abortion rate by 2.86. Both estimates have

Treatment	Estimate	SE	N
Abortion Count			
RML	-514.00	1,651.32	458
MML	13,457.34	11,119.16	458
Abortion Rate			
RML	-0.38	0.45	458
MML	2.86	2.49	458

The outcome is the number of abortions or abortion rate per 1,000 women aged 14-44 for residents of the state, regardless of where the abortion was performed. Regressions include state and year fixed effects, a state-specific time trend, and controls for demographic characteristics and state policies. Regressions are weighted by state population and standard errors are clustered at the state level.

Table C.9: Abortion Results

large standard errors and are statistically insignificant.

I use event studies to assess the plausibility of the parallel trends assumption for these outcomes. The results appear in Figure C.9. The point estimates show a downward trend in abortion counts following an RML and an upward trend in abortion counts following an MML. Both event studies show point estimates that are closer to zero and have smaller confidence intervals before marijuana legalization than after, and patterns are very similar for abortion counts and abortion rates. However, as with the main results, all estimated effects are imprecise with very wide confidence intervals. For MMLs, for example, the 95% confidence interval for the effect on abortion counts in the fourth year after the MML became effective ranges from -10,389 to 54,011.

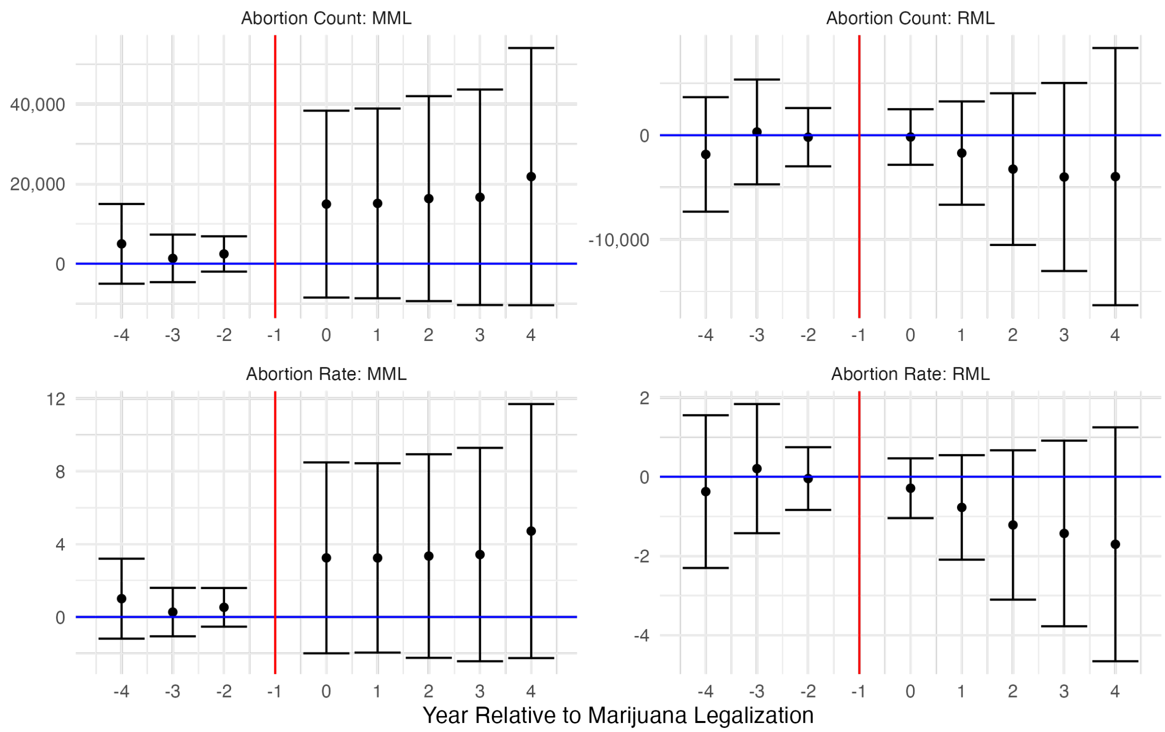

The outcome is abortion count for residents of the state, regardless of where the abortion was performed. Regressions include state and a state-specific time trend, and controls for demographic characteristics and state policies. Regressions are weighted using total state population, and standard errors are clustered at the state level.

Figure C.9: Abortion Event Studies

C.4.3 Effects on Birth Control

As explained in the Mechanisms section, marijuana legalization and the corresponding increase in marijuana use could affect birth control use. This section discusses estimated effects of marijuana legalization on birth control use. Estimates are obtained using data from the NLSY. The outcome is the number of times the individual used birth control in the past year. I estimate that RMLs decrease use of birth control by 8.3 uses per year and MMLs decrease use of birth control by 2.4 uses per year. These estimates are consistent with the past associations found in the medical literature between marijuana use and a lower likelihood of using birth control [89, 119, 120].

Due to data limitations, these findings should be interpreted with caution. Suppose an individual has twenty sexual encounters and uses birth control ten times in year one. In year two, the individual has two sexual encounters and uses birth control two times. In year two the individual is more likely to use birth control in each sexual encounter, but the estimated effect on her use of birth control is a decrease of eight uses per year. The natural solution would be to divide the number of birth control uses by the number of sexual encounters in the past year. However, both birth control use and number of sexual encounters in the past year have many missing observations. In fact, there are zero observations in states with RMLs in the post-RML period for which birth control use and number of sexual encounters in the past year are both available, and therefore the effect of RMLs on the probability of using birth control in each sexual encounter cannot be estimated using this data.

	Frequency of Birth Control Use Last Year
Any RML Effective	-8.319**
	(3.725)
Any MML Effective	-2.391
	(3.193)
Observations	51,239

*: $p<0.1$ **: $p<0.05$ ***: $p<0.01$. The outcome is the number of times the individual used birth control in the past year, measured at the individual level using data from the National Longitudinal Survey of Youth from 2005-2011. Coefficients estimated using two-way fixed effects. The regression includes state, year-quarter and state-quarter fixed effects and state-specific linear time trends, as well as controls for share of total state population by race, ethnicity, age, and education; the unemployment rate, which is also interacted with the demographic controls; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; decriminalization of marijuana. The regression includes controls for the respondent's age, sex, education and marital status.

Table C.10: Birth Control Results

C.5 Comparison with Previous Literature

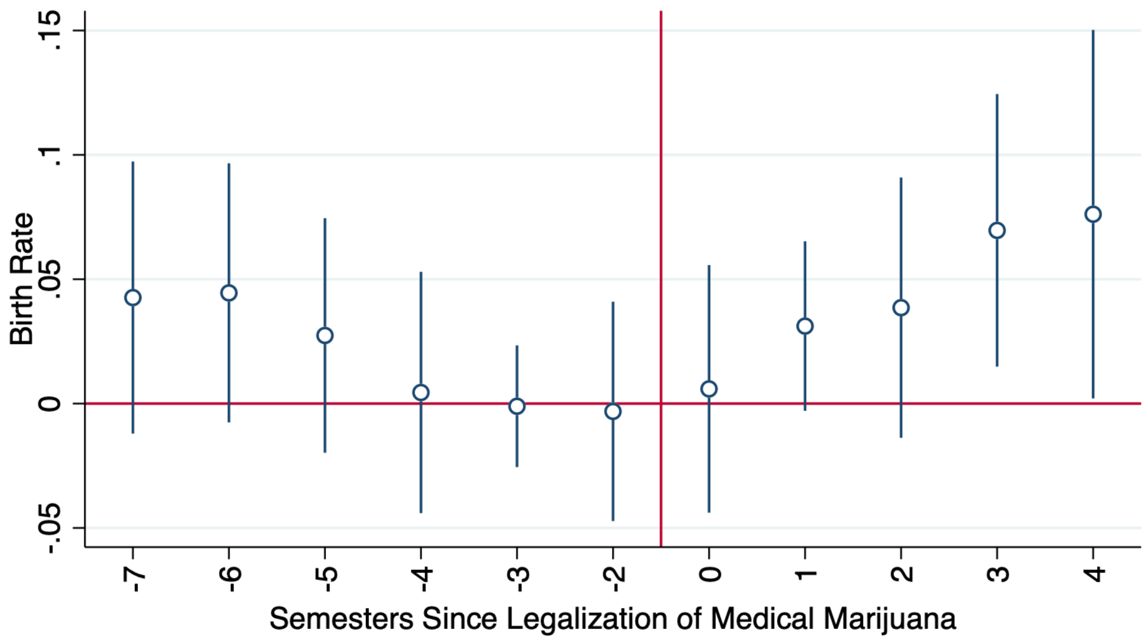

The outcome is the birth rate, measured at the state-quarter level using data from the National Vital Statistics System from 2005-2014. Coefficients estimated using two-way fixed effects. Includes state fixed effects, year-quarter fixed effects and state-specific linear time trends as well as controls for share of total state population by race, ethnicity, age and education; unemployment rate; median household income; state cigarette and beer tax rates; Medicaid expansion; abortion laws; same sex marriage laws; Medicaid eligibility threshold; WIC EBT indicator; and decriminalization of marijuana. The regression is weighted by state population and standard errors are clustered at the state level.

Figure C.10: MML Event Study for Birth Rates (2005-2014)

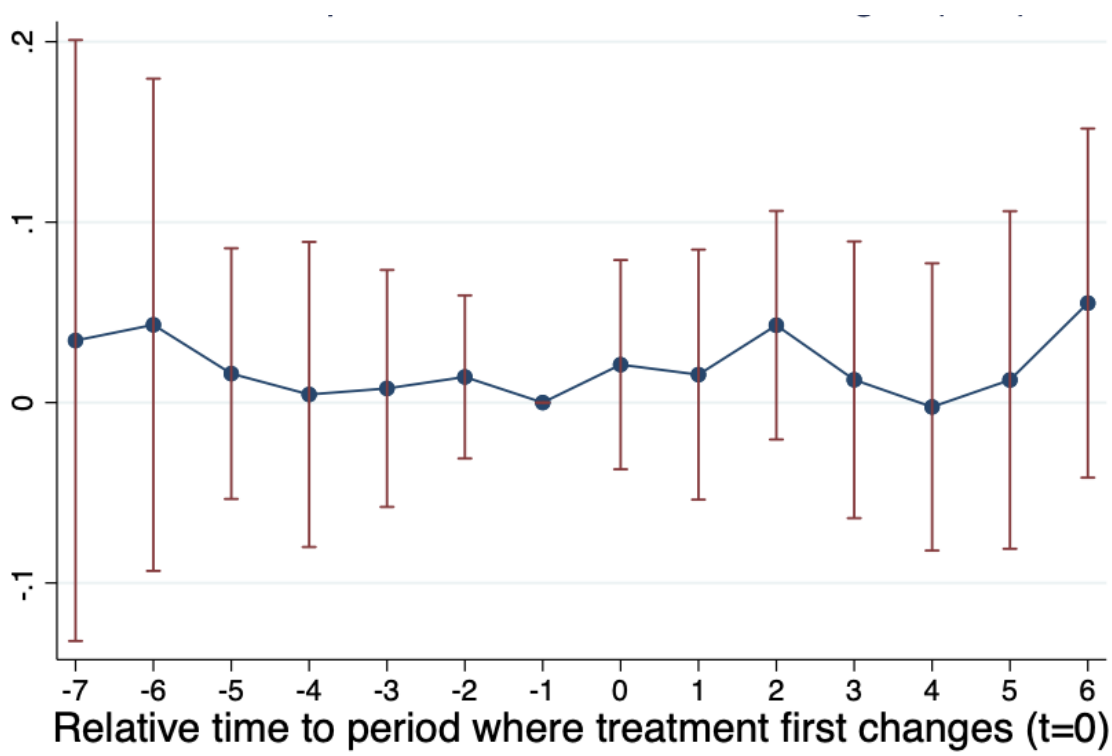

Birth rate is the outcome, measured at the state-quarter level using data from the National Vital Statistics System from 2005-2014. Coefficients are estimated using DIDM. Includes state fixed effects, year-quarter fixed effects and state-specific linear time trends as well as controls for share of total population by race, ethnicity, age and education, as well as unemployment rate, median household income, state cigarette and beer tax rates, Medicaid expansion, abortion laws, same sex marriage laws, Medicaid eligibility threshold, WIC EBT indicator, and decriminalization of marijuana.

Figure C.11: MML Event Study for Birth Rates with DID_M Estimator (2005-2014)

	Women 14-44	Women 21-30
Panel A: 2005-2019		
Any RML Effective	-0.045**	-0.024**
	(0.020)	(0.009)
Any MML Effective	0.004	-0.006
	(0.012)	(0.007)
Observations	2,907	2,907
Panel B: 2005-2014		
Any RML Effective	-0.019	-0.020
	(0.020)	(0.020)
Any MML Effective	0.030**	0.004
	(0.013)	(0.007)
Observations	2,040	2,040
Panel C: Dropping states with MMLs after 2014		
Any RML Effective	-0.093**	-0.047**
	(0.038)	(0.018)
Any MML Effective	-0.051	-0.007
	(0.039)	(0.018)
Observations	5,456	5,456

*: p<0.1 **: p<0.05 ***: p<0.01. The outcome is the birth rate, which is measured at the state-quarter level using data from the National Vital Statistics System. Data from 2005-2019 is used for Panel A, 2004-2014 for Panel B. Panel C uses data from 1989-2019 and drops states that legalized medical marijuana after 2014. All regressions include controls for share of total population by race, ethnicity, age and education, as well as unemployment rate, median household income, state cigarette and beer tax rates, Medicaid expansion, abortion laws, same sex marriage laws, Medicaid eligibility threshold, WIC EBT indicator, and decriminalization of marijuana. State and year-quarter fixed effects and a state-specific time trend are included, and regressions are weighted by state population. Standard errors are clustered at the state level. Standard errors in parentheses. RML: recreational marijuana law. MML: medical marijuana law.

Table C.11: Additional Birth Rate Analysis

This paper is most closely related to [93], which studies the effects of MMLs on fertility. My analysis of the effects of MMLs builds on [93] in multiple ways. First, I use a longer time frame: [93] use the years 2005-2014 for all outcomes, while I use the years 1989-2019 for birth rates, 1997-2017 for marijuana use and sexual activity, and 2000-2019 for gonorrhea rates. Second, I verify that my results are robust to heterogeneous treatment effects using the method from [90]. In this section, I investigate why the results I estimate for MMLs differ from the results of [93].

C.5.1 Birth Rates

Although [93] estimate that MMLs increase birth counts by about 2%, I find that MMLs lead to statistically insignificant decreases in birth rates. One potential explanation is that the effect of MMLs from 1989-2004, the years in my data that are before the period studied in [93], was different than the effect from 2005-2014. I used only the years 2005-2019 for my analysis to determine whether effects from 1989-2004 caused the difference in results. However, Table C.11 shows that results from 2005-2019 resemble my main results much more closely than the results from [93], with estimated coefficients that are statistically insignificant and close to zero. These estimates also serve as a robustness check for my main results, showing that using a shorter pre-treatment period does not substantially change the results. The long pre-treatment period for RMLs could be a source of concern in the absence of this robustness check.

Another possibility is that the effect of MMLs changed after 2014. Table C.11 shows that when the states that passed MMLs after 2014 are dropped from my sample, the results are very similar to the results I find with the full sample. Therefore, the difference between my results and the results in [93] do not appear to arise from the effects of MMLs changing between their sample period and mine.

I repeat my analysis using only the years 2005-2014, which are used in [93]. Table C.11 shows that my results in this period are similar to the results in [93], with positive coefficients, one of which is statistically significant. Estimates are not expected to be identical, as [93] uses a Poisson regression to obtain their main results. The event study also shows an upward trend after legalization, as shown in Figure C.10. The fact that results in the 2005-2014 period are very similar between the two papers demonstrates that the substantive difference in results does not come from [93] using birth counts with a Poisson regression, while I use birth rates with OLS; nor does it arise from a difference in data cleaning or controls. I also replicate the results from [93] using birth counts and a Poisson regression and find nearly identical results, which are available upon request.

When I use the DID_M estimator in the 2005-2014 period, Figure C.11 shows that the event study has a flat trend after a MML becomes effective. The average effects from the DID_M estimator from 2005-2014 are statistically insignificant and close to zero, and one of the two is negative. The discrepancy between the DID_M and two-way fixed effects estimates from 2005-2014 arises from negative weights creating upward bias in the two-way fixed effects estimate. Negative weights account for 30% of all weights from 2005-2014, but only 19% of all weights from 1989-2019. The larger share of negative weights from 2005-2014 creates upward bias in the results in [93], but in the longer time period that I use, negative weights are a smaller share of all weights and they induce less bias. Therefore, the DID_M estimator should be used in the 2005-2014 period but not in this paper's period of study from 1989-2019. With use of the DID_M estimator from 2005-2014 and standard two-way fixed effects in the full period from 1989-2019, results are consistent across periods of study.

C.5.2 Other Outcomes

The remaining differences in our results are more minor and most likely arise from my use of a longer time frame. My results for the effect of MMLs on days of marijuana use in the past month are nearly identical to the finding in [93]. However, [93] find that MMLs increase the probability of any marijuana use, while I find no effect. [93] find that MMLs increase the probability of being sexually active by 4.3 percentage points but do not affect the frequency of sexual activity. By contrast, I estimate that that MMLs increase the frequency of sexual activity in the past month by 1.8 sexual encounters but do not affect the probability of being sexually active. My result for gonorrhea is similar to [93], as they also found a positive point estimate with large standard errors. But unlike [93], I do not find a statistically significant effect of MMLs.

Bibliography

[1] B. of Governors of the Federal Reserve System, *Economic Well-Being of U.S. Households in 2022*, *Federal Reserve* (May, 2023).

[2] C. Karma, *Buy now pay later surges throughout pandemic, consumers' credit takes a hit*, Sept., 2021.

[3] CFPB, *Buy Now, Pay Later: Market trends and consumer impacts*, *Consumer Financial Protection Bureau* (Sept., 2022) 83.

[4] S. Smith, *2021 Top 100 Retailers*, 2021.

[5] A. Carrns, *'Buy Now, Pay Later' Loans May Soon Play Bigger Role in Credit Scores*, *The New York Times* (Dec., 2021).

[6] A. Meola, *Almost 75% of BNPL users in the US are millennials or Gen Z*, 2021.

[7] TransUnion, *Point of Sale Loan Applicants Generally Use Other Forms of Credit More Actively*, Sept., 2021.

[8] T. Washington and A. Horowitz-Ghazi, *Buy now, pay dearly?*, *NPR* (May, 2022).

[9] T. S. Bernard, *Consumers and Companies Are Buying In on Paying Later*, *The New York Times* (Sept., 2021).

[10] J. Boxell, *Analysis | How Old-Style Buy Now, Pay Later Became Trendy 'BNPL'*, *Washington Post* (Sept., 2021).

[11] P. Krishna, *Eat Now, Pay Later: Going Into Debt for Food*, *The New York Times* (Aug., 2022).

[12] D. P. Morgan, M. R. Strain, and I. Seblani, *How Payday Credit Access Affects Overdrafts and Other Outcomes*, *Journal of Money, Credit and Banking* **44** (2012), no. 2-3 519–531. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1538-4616.2011.00499.x.

[13] J. Zinman, *Restricting consumer credit access: Household survey evidence on effects around the Oregon rate cap*, *Journal of Banking & Finance* **34** (Mar., 2010) 546–556.

[14] A. Morse, *Payday lenders: Heroes or villains?*, *Journal of Financial Economics* **102** (Oct., 2011) 28–44.

[15] D. Karlan and J. Zinman, *Expanding Credit Access: Using Randomized Supply Decisions to Estimate the Impacts*, *The Review of Financial Studies* **23** (2010), no. 1 433–464. Publisher: [Oxford University Press, Society for Financial Studies].

[16] R. V. Galperin and A. Weaver, *Payday Lending Regulation and the Demand for Alternative Financial Services*, Oct., 2014.

[17] J. Gathergood, B. Guttman-Kenney, and S. Hunt, *How Do Payday Loans Affect Borrowers? Evidence from the U.K. Market*, *The Review of Financial Studies* **32** (Feb., 2019) 496–523.

[18] P. M. Skiba and J. Tobacman, *Do Payday Loans Cause Bankruptcy?*, *Journal of Law and Economics* **62** (2019), no. 3 485–519. Num Pages: 35.

[19] D. Campbell, F. Asís Martínez-Jerez, and P. Tufano, *Bouncing out of the banking system: An empirical analysis of involuntary bank account closures*, *Journal of Banking & Finance* **36** (Apr., 2012) 1224–1235.

[20] S. Carrell and J. Zinman, *In Harm's Way? Payday Loan Access and Military Personnel Performance*, *The Review of Financial Studies* **27** (Sept., 2014) 2805–2840.

[21] N. Bhutta, *Payday loans and consumer financial health*, *Journal of Banking & Finance* **47** (Oct., 2014) 230–242.

[22] B. T. Melzer, *The Real Costs of Credit Access: Evidence from the Payday Lending Market**, *The Quarterly Journal of Economics* **126** (Feb., 2011) 517–555.

[23] P. Heidhues and B. Kőszegi, *Exploiting Naïvete about Self-Control in the Credit Market*, *American Economic Review* **100** (Dec., 2010) 2279–2303.

[24] S. Meier and C. Sprenger, *Present-Biased Preferences and Credit Card Borrowing*, *American Economic Journal: Applied Economics* **2** (Jan., 2010) 193–210.

[25] T. Kuchler and M. Pagel, *Sticking to your plan: The role of present bias for credit card paydown*, *Journal of Financial Economics* **139** (Feb., 2021) 359–388.

[26] J. Alcazar and T. Bradford, *The Appeal and Proliferation of Buy Now, Pay Later: Consumer and Merchant Perspectives*, .

[27] B. J. Keys and J. Wang, *Minimum payments and debt paydown in consumer credit cards*, *Journal of Financial Economics* **131** (Mar., 2019) 528–548.

[28] M. Di Maggio, E. Williams, and J. Katz, *Buy Now, Pay Later Credit: User Characteristics and Effects on Spending Patterns*, Sept., 2022.

[29] E. deHaan, J. Kim, B. Lourie, and C. Zhu, *Buy Now Pay (Pain?) Later*, Sept., 2022.

[30] B. Guttman-Kenney, C. Firth, and J. Gathergood, *Buy Now, Pay Later (BNPL)...On Your Credit Card*, Mar., 2022. arXiv:2201.01758 [econ, q-fin].

[31] A. Cooban, *Buy now, pay later is a huge hit with shoppers. Just how dangerous is it?*, Dec., 2021.

[32] A. Sng and C. Tan, *Buy Now Pay Later in Singapore: Regulatory Gaps and Reform*, Apr., 2022.

[33] C. Reagan, *Retailers bid farewell to layaway, as shoppers embrace buy now, pay later options*, Sept., 2021. Section: Retail.

[34] M. Black and R. Saks Frankel, *What Is The Average Credit Card Interest Rate?*, Apr., 2022. Section: Credit cards.

[35] K. Hoevelmann, *How Payday Loans Work*, *St. Louis Federal Reserve* (July, 2019).

[36] M. Lux and B. Epps, *Grow Now, Regulate Later? Regulation urgently needed to support transparency and sustainable growth for Buy-Now, Pay-Later*, *Mossavar-Rahmani Center for Business and Government Working Papers* **M-RCBG Associate Working Paper No. 182** (2022).

[37] G. Reich, *30+ Buy Now, Pay Later Trends & Statistics for Banks in 2022*, May, 2022.

[38] T. Paul, *Affirm review: When should you use 'buy now, pay later' provider?*, July, 2021. Section: Select: Reviews.

[39] M. Hanbury, *This is what the average Walmart shopper looks like*, 2020.

[40] D. Laibson, *Golden Eggs and Hyperbolic Discounting*, *The Quarterly Journal of Economics* **112** (1997), no. 2 443–477. Publisher: Oxford University Press.

[41] M. Leyes, *DFPI Report Shows Changes in Consumer Lending, Decrease in PACE Program*, Oct., 2021.

[42] J. L. Liu, *Beyond Neighborhood Food Environments: Distance Traveled to Food Establishments in 5 US Cities, 2009–2011*, *Preventing Chronic Disease* **12** (2015).

[43] U. C. Bureau, *US Census Bureau Quarterly Retail E-Commerce Sales 3rd Quarter 2021*, .

[44] R. J. Lee, I. N. Sener, P. L. Mokhtarian, and S. L. Handy, *Relationships between the online and in-store shopping frequency of Davis, California residents*, *Transportation Research Part A: Policy and Practice* **100** (June, 2017) 40–52.

[45] J. Chen and J. Roth, *Logs with zeros? Some problems and solutions*, *Papers* (July, 2023). Number: 2212.06080 Publisher: arXiv.org.

[46] J. White, *How Utilization Rate Affects Credit Scores - Experian*, Jan., 2018.

[47] D. Lott and a. p. r. e. i. t. R. P. R. F. a. t. A. Fed, *Now in the U.S.: Buy Now, Pay Later*, Mar., 2021.

[48] R. Shevlin, *Buy Now, Pay Later: The "New" Payments Trend Generating $100 Billion In Sales*, Sept., 2021. Section: Fintech.

[49] P. Gerrans, D. G. Baur, and S. Lavagna-Slater, *Fintech and responsibility: Buy-now-pay-later arrangements*, *Australian Journal of Management* (July, 2021) 03128962211032448. Publisher: SAGE Publications Ltd.

[50] R. K. Crump, V. J. Hotz, G. W. Imbens, and O. A. Mitnik, *Dealing with limited overlap in estimation of average treatment effects*, *Biometrika* **96** (Mar., 2009) 187–199.

[51] T. Lau, *A Bid to Counter Big Money in Politics Is Gaining Steam | Brennan Center for Justice*, 2019.

[52] B. Jones, *Most Americans want to limit campaign spending, say big donors have greater political influence*, May, 2018.

[53] D. Montanaro, *Poll: Despite Record Turnout, 80 Million Americans Didn't Vote. Here's Why*, *NPR* (Dec., 2020).

[54] B. Ferman and C. Pinto, *Inference in Differences-in-Differences with Few Treated Groups and Heteroskedasticity*, *The Review of Economics and Statistics* **101** (July, 2019) 452–467.

[55] L. Bouton, J. Cage, E. Dewitte, and V. Pons, *Small Campaign Donors*, SSRN Scholarly Paper 3978318, Social Science Research Network, Rochester, NY, Dec., 2021.

[56] M. Gentzkow, J. M. Shapiro, and M. Sinkinson, *The Effect of Newspaper Entry and Exit on Electoral Politics*, *American Economic Review* **101** (Dec., 2011) 2980–3018.

[57] J. S. Hastings, T. J. Kane, D. O. Staiger, and J. M. Weinstein, *The effect of randomized school admissions on voter participation*, *Journal of Public Economics* **91** (June, 2007) 915–937.

[58] R. Shachar and B. Nalebuff, *Follow the Leader: Theory and Evidence on Political Participation*, *American Economic Review* **89** (June, 1999) 525–547.

[59] A. J. Hoffer, *Are Voters Rational on the Margin?: A Spatial Analysis of Voter Turnout in U.S. Presidential Elections*, Apr., 2014.

[60] J. Yu, *Does State Online Voter Registration Increase Voter Turnout?**, *Social Science Quarterly* **100** (2019), no. 3 620–634. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ssqu.12598.

[61] R. E. Hogan, *Campaign Spending and Voter Participation in State Legislative Elections*, *Social Science Quarterly* **94** (2013), no. 3 840–864. Publisher: Wiley.

[62] K. Milligan, E. Moretti, and P. Oreopoulos, *Does education improve citizenship? Evidence from the United States and the United Kingdom*, *Journal of Public Economics* **88** (Aug., 2004) 1667–1695.

[63] O. Falck, R. Gold, and S. Heblich, *E-lections: Voting Behavior and the Internet*, *American Economic Review* **104** (July, 2014) 2238–2265.

[64] L. R. Heller, J. Miller, and E. F. Stephenson, *Voter ID Laws and Voter Turnout*, *Atlantic Economic Journal* **47** (June, 2019) 147–157.

[65] B. J. McCabe and J. A. Heerwig, *Diversifying the Donor Pool: How Did Seattle's Democracy Voucher Program Reshape Participation in Municipal Campaign Finance?*, *Election Law Journal: Rules, Politics, and Policy* **18** (Sept., 2019) 323–341. Publisher: Mary Ann Liebert, Inc., publishers.

[66] C. Yorgason, *Campaign finance vouchers do not reduce donor inequality*, . Publisher: SocArXiv.

[67] A. Griffith and T. Noonen, *The effects of public campaign funding: Evidence from Seattle's Democracy Voucher program*, *Journal of Public Economics* **211** (July, 2022) 104676.

[68] S. Kliff, *Seattle's radical plan to fight big money in politics*, Nov., 2018.

[69] C. of Seattle, *About the Program - DemocracyVoucher | seattle.gov*, 2021.

[70] M. Mohrman, *State population steadily increases, tops 7.7 million residents in 2021*, *State of Washington Office of Financial Management* (2021) 10.

[71] Z. L. Hajnal, *Opinion | Why Does No One Vote in Local Elections?*, *The New York Times* (Oct., 2018).

[72] A. F. Haughwout, R. P. Inman, and J. V. Henderson, *Should Suburbs Help Their Central City? [with Comment]*, *Brookings-Wharton Papers on Urban Affairs* (2002) 45–94. Publisher: Brookings Institution Press.

[73] M. J. Malbin, P. W. Brusoe, and B. Glavin, *Small Donors, Big Democracy: New York City's Matching Funds as a Model for the Nation and States*, *Election Law Journal: Rules, Politics, and Policy* **11** (Mar., 2012) 3–20.

[74] T. G. Conley and C. R. Taber, *Inference with "Difference in Differences" with a Small Number of Policy Changes*, *The Review of Economics and Statistics* **93** (Feb., 2011) 113–125.

[75] S. G. Donald and K. Lang, *Inference with Difference-in-Differences and Other Panel Data*, *Review of Economics and Statistics* **89** (May, 2007) 221–233.

[76] J. G. MacKinnon and M. D. Webb, *Wild Bootstrap Inference for Wildly Different Cluster Sizes*, *Journal of Applied Econometrics* **32** (2017), no. 2 233–254. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jae.2508.

[77] D. Roodman, M. Nielsen, J. G. MacKinnon, and M. D. Webb, *Fast and wild: Bootstrap inference in Stata using boottest*, *The Stata Journal* **19** (Mar., 2019) 4–60. Publisher: SAGE Publications.

[78] K. H. Brodersen, F. Gallusser, J. Koehler, N. Remy, and S. L. Scott, *Inferring causal impact using Bayesian structural time-series models*, *The Annals of Applied Statistics* **9** (Mar., 2015) 247–274. Publisher: Institute of Mathematical Statistics.

[79] A. Abadie and J. Gardeazabal, *The Economic Costs of Conflict: A Case Study of the Basque Country*, *American Economic Review* **93** (Mar., 2003) 113–132.

[80] T. C. Buchmueller, J. DiNardo, and R. G. Valletta, *The Effect of an Employer Health Insurance Mandate on Health Insurance Coverage and the Demand for Labor: Evidence from Hawaii*, *American Economic Journal: Economic Policy* **3** (Nov., 2011) 25–51.

[81] S. Cunningham and M. Shah, *Decriminalizing Indoor Prostitution: Implications for Sexual Violence and Public Health*, *The Review of Economic Studies* **85** (July, 2018) 1683–1715.

[82] E. Andreyeva and B. Ukert, *The Impact of Medical Marijuana Laws and Dispensaries on Self-Reported Health*, *Forum for Health Economics & Policy* **22** (Oct., 2019). Publisher: De Gruyter Section: Forum for Health Economics & Policy.

[83] D. F. Ullman, *The Effect of Medical Marijuana on Sickness Absence*, *Health Economics* **26** (2017), no. 10 1322–1327. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/hec.3390.

[84] M. Baggio, A. Chong, and S. Kwon, *Marijuana and alcohol: Evidence using border analysis and retail sales data*, *Canadian Journal of Economics/Revue canadienne d'économique* **53** (2020), no. 2 563–591. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/caje.12437.

[85] N. W. Chan, J. Burkhardt, and M. Flyr, *The Effects of Recreational Marijuana Legalization and Dispensing on Opioid Mortality*, *Economic Inquiry* **58** (2020), no. 2 589–606. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ecin.12819.

[86] T. D. Gundersen, N. Jørgensen, A.-M. Andersson, A. K. Bang, L. Nordkap, N. E. Skakkebæk, L. Priskorn, A. Juul, and T. K. Jensen, *Association Between Use of Marijuana and Male Reproductive Hormones and Semen Quality: A Study Among 1,215 Healthy Young Men*, *American Journal of Epidemiology* **182** (Sept., 2015) 473–481. Publisher: Oxford Academic.

[87] M. Bari, N. Battista, V. Pirazzi, and M. Maccarrone, *The manifold actions of endocannabinoids on female and male reproductive events*, *Frontiers in Bioscience (Landmark Edition)* **16** (Jan., 2011) 498–516.

[88] A. J. Sun and M. L. Eisenberg, *Association Between Marijuana Use and Sexual Frequency in the United States: A Population-Based Study*, *The Journal of Sexual Medicine* **14** (Nov., 2017) 1342–1347.

[89] J. Guo, I.-J. Chung, K. G. Hill, J. D. Hawkins, R. F. Catalano, and R. D. Abbott, *Developmental relationships between adolescent substance use and risky sexual behavior in young adulthood*, *The Journal of Adolescent Health: Official Publication of the Society for Adolescent Medicine* **31** (Oct., 2002) 354–362.

[90] C. de Chaisemartin and X. D'Haultfœuille, *Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects*, *American Economic Review* **110** (Sept., 2020) 2964–2996.

[91] M. Leyton, *Cannabis legalization: Did we make a mistake? Update 2019*, *Journal of Psychiatry & Neuroscience : JPN* **44** (Sept., 2019) 291–293.

[92] H. Wen, J. M. Hockenberry, and J. R. Cummings, *The effect of medical marijuana laws on adolescent and adult use of marijuana, alcohol, and other substances*, *Journal of Health Economics* **42** (July, 2015) 64–80.

[93] M. Baggio, A. Chong, and D. Simon, *Sex, marijuana and baby booms*, *Journal of Health Economics* **70** (Mar., 2020) 102283.

[94] D. S. Hasin, A. L. Sarvet, M. Cerdá, K. M. Keyes, M. Stohl, S. Galea, and M. M. Wall, *US Adult Illicit Cannabis Use, Cannabis Use Disorder, and Medical*

Marijuana Laws: 1991-1992 to 2012-2013, *JAMA psychiatry* **74** (June, 2017) 579–588.

[95] R. L. Pacula, D. Powell, P. Heaton, and E. L. Sevigny, *Assessing the Effects of Medical Marijuana Laws on Marijuana Use: The Devil is in the Details*, *Journal of Policy Analysis and Management* **34** (2015), no. 1 7–31. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/pam.21804.

[96] S. S. Martins, C. M. Mauro, J. Santaella-Tenorio, J. H. Kim, M. Cerda, K. M. Keyes, D. S. Hasin, S. Galea, and M. Wall, *State-level medical marijuana laws, marijuana use and perceived availability of marijuana among the general U.S. population*, *Drug and Alcohol Dependence* **169** (Dec., 2016) 26–32.

[97] A. C. Bradford and W. D. Bradford, *Medical Marijuana Laws Reduce Prescription Medication Use In Medicare Part D*, *Health Affairs* **35** (July, 2016) 1230–1236. Publisher: Health Affairs.

[98] C. M. Mauro, P. Newswanger, J. Santaella-Tenorio, P. M. Mauro, H. Carliner, and S. S. Martins, *Impact of Medical Marijuana Laws on State-Level Marijuana Use by Age and Gender, 2004–2013*, *Prevention Science* **20** (Feb., 2019) 205–214.

[99] A. R. Williams, J. Santaella-Tenorio, C. M. Mauro, F. R. Levin, and S. S. Martins, *Loose regulation of medical marijuana programs associated with higher rates of adult marijuana use but not cannabis use disorder*, *Addiction* **112** (2017), no. 11 1985–1991. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/add.13904.

[100] M. Cerdá, C. Mauro, A. Hamilton, N. S. Levy, J. Santaella-Tenorio, D. Hasin, M. M. Wall, K. M. Keyes, and S. S. Martins, *Association Between Recreational Marijuana Legalization in the United States and Changes in Marijuana Use and Cannabis Use Disorder From 2008 to 2016*, *JAMA Psychiatry* **77** (Feb., 2020) 165–171.

[101] M. S. Subbaraman and W. C. Kerr, *Subgroup trends in alcohol and cannabis co-use and related harms during the rollout of recreational cannabis legalization in Washington state*, *International Journal of Drug Policy* **75** (Jan., 2020) 102508.

[102] M. Gnofam, A. A. Allshouse, E. H. Stickrath, and T. D. Metz, *Impact of Marijuana Legalization on Prevalence of Maternal Marijuana Use and Perinatal Outcomes*, *American Journal of Perinatology* **37** (Jan., 2020) 59–65. Publisher: Thieme Medical Publishers.

[103] E. A. Stormshak, A. S. Caruthers, J. M. Gau, and C. Winter, *The impact of recreational marijuana legalization on rates of use and behavior: A 10-year comparison of two cohorts from high school to young adulthood*, *Psychology of*

Addictive Behaviors **33** (2019) 595–602. Place: US Publisher: American Psychological Association.

[104] A. Meinhofer, A. E. Witman, J. M. Hinde, and K. Simon, *Marijuana liberalization policies and perinatal health*, *Journal of Health Economics* **80** (Dec., 2021) 102537.

[105] D. M. Anderson and D. I. Rees, *The Public Health Effects of Legalizing Marijuana*, *Journal of Economic Literature* **61** (Mar., 2023) 86–143.

[106] J. Currie and H. Schwandt, *Short- and long-term effects of unemployment on fertility*, *Proceedings of the National Academy of Sciences* **111** (Oct., 2014) 14734–14739.

[107] K. Mulligan, *Contraception Use, Abortions, and Births: The Effect of Insurance Mandates*, *Demography* **52** (2015), no. 4 1195–1217. Publisher: [Population Association of America, Springer].

[108] S. Fischer, H. Royer, and C. White, *The impacts of reduced access to abortion and family planning services on abortions, births, and contraceptive purchases*, *Journal of Public Economics* **167** (Nov., 2018) 43–68.

[109] B. Branan, *Does it make financial sense to keep your medical marijuana card after Jan. 1?*, *The Sacramento Bee* (Dec., 2017).

[110] A. Daniller, *Two-thirds of Americans support marijuana legalization*, Nov., 2019.

[111] N. McCarthy, *The Arguments For And Against Marijuana Legalization In The U.S. [Infographic]*, June, 2019. Section: Industry.

[112] G. Galvin, *U.S. Births Continue to Fall, Fertility Rate Hits Record Low*, *US News & World Report* (May, 2020).

[113] CDC, *Fertility Rates by State*, tech. rep., Centers for Disease Control and Prevention, Jan., 2022.

[114] L. Cameron, J. Seager, and M. Shah, *Crimes Against Morality: Unintended Consequences of Criminalizing Sex Work**, *The Quarterly Journal of Economics* **136** (Feb., 2021) 427–469.

[115] CDC, *STD Facts - Gonorrhea*, *Centers for Disease Control and Prevention* (Mar., 2022).

[116] V. Koppa, *The Effect of Alcohol Access on Sexually Transmitted Diseases: Evidence From the Minimum Legal Drinking Age*, *American Journal of Health Economics* **4** (May, 2018) 164–184. Publisher: The University of Chicago Press.

[117] C. Carpenter, *Youth alcohol use and risky sexual behavior: evidence from underage drunk driving laws*, *Journal of Health Economics* **24** (May, 2005) 613–628.

[118] B. B. Gorzalka, M. N. Hill, and S. C. H. Chang, *Male-female differences in the effects of cannabinoids on sexual behavior and gonadal hormone function*, *Hormones and Behavior* **58** (June, 2010) 91–99.

[119] J. B. Kingree, R. Braithwaite, and T. Woodring, *Unprotected sex as a function of alcohol and marijuana use among adolescent detainees*, *Journal of Adolescent Health* **27** (Sept., 2000) 179–185.

[120] A. D. Bryan, S. J. Schmiege, and R. E. Magnan, *Marijuana use and risky sexual behavior among high-risk adolescents: Trajectories, risk factors, and event-level relationships*, *Developmental Psychology* **48** (2012), no. 5 1429–1442. Place: US Publisher: American Psychological Association.

[121] J. J. Palamar, P. Acosta, D. C. Ompad, and S. R. Friedman, *A Qualitative Investigation Comparing Psychosocial and Physical Sexual Experiences Related to Alcohol and Marijuana Use among Adults*, *Archives of sexual behavior* **47** (Apr., 2018) 757–770.

[122] E. Wiebe and A. Just, *How Cannabis Alters Sexual Experience: A Survey of Men and Women*, *The Journal of Sexual Medicine* **16** (Nov., 2019) 1758–1762.

[123] B. K. Lynn, J. D. López, C. Miller, J. Thompson, and E. C. Campian, *The Relationship between Marijuana Use Prior to Sex and Sexual Function in Women*, *Sexual Medicine* **7** (June, 2019) 192–197.

[124] K. F. Boehnke, S. Gangopadhyay, D. J. Clauw, and R. L. Haffajee, *Qualifying Conditions Of Medical Cannabis License Holders In The United States*, *Health Affairs* **38** (Feb., 2019) 295–302. Publisher: Health Affairs.

[125] S. S. du Plessis, A. Agarwal, and A. Syriac, *Marijuana, phytocannabinoids, the endocannabinoid system, and male fertility*, *Journal of Assisted Reproduction and Genetics* **32** (Nov., 2015) 1575–1588.

[126] D. M. Anderson and D. I. Rees, *The Legalization of Recreational Marijuana: How Likely is the Worst-Case Scenario?*, *Journal of Policy Analysis and Management* **33** (2014), no. 1 221–232.

[127] A. F. Harlow, A. K. Wesselink, E. E. Hatch, K. J. Rothman, and L. A. Wise, *Male Preconception Marijuana Use and Spontaneous Abortion: A Prospective Cohort Study*, *Epidemiology (Cambridge, Mass.)* **32** (Mar., 2021) 239–247.

[128] W. H. Organization, *Indicator Metadata Registry Details*, *World Health Organization* (2022).

[129] N. Austin and S. Harper, *Constructing a longitudinal database of targeted regulation of abortion providers laws*, *Health Services Research* **54** (2019), no. 5 1084–1089. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1475-6773.13185.

[130] K. Neels, Z. Theunynck, and J. Wood, *Economic recession and first births in Europe: recession-induced postponement and recuperation of fertility in 14 European countries between 1970 and 2005*, *International Journal of Public Health* **58** (Feb., 2013) 43–55.

[131] J. Williams, R. Liccardo Pacula, F. J. Chaloupka, and H. Wechsler, *Alcohol and marijuana use among college students: economic complements or substitutes?*, *Health Economics* **13** (2004), no. 9 825–843. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/hec.859.

[132] L. Cameron and J. Williams, *Cannabis, Alcohol and Cigarettes: Substitutes or Complements?*, *Economic Record* **77** (2001), no. 236 19–34. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1475-4932.00002.

[133] T. Wong, A. Singh, J. Mann, L. Hansen, and S. McMahon, *Gender Differences in Bacterial STIs in Canada*, *BMC Women's Health* **4** (Aug., 2004) S26.

[134] M. S. Kearney and P. B. Levine, *Investigating recent trends in the U.S. teen birth rate*, *Journal of Health Economics* **41** (May, 2015) 15–29.

[135] C. Chuard and P. Chuard-Keller, *Baby bonus in Switzerland: Effects on fertility, newborn health, and birth-scheduling*, *Health Economics* **30** (Sept., 2021) 2092–2123.

[136] Y. Lu and D. J. G. Slusky, *The Impact of Women's Health Clinic Closures on Fertility*, *American Journal of Health Economics* **5** (July, 2019) 334–359. Publisher: The University of Chicago Press.

[137] J. Roth, P. H. Sant'Anna, A. Bilinski, and J. Poe, *What's trending in difference-in-differences? A synthesis of the recent econometrics literature*, *Journal of Econometrics* (Apr., 2023) S0304407623001318.

[138] W. C. Kerr, Y. Ye, M. S. Subbaraman, E. Williams, and T. K. Greenfield, *Changes in Marijuana Use Across the 2012 Washington State Recreational Legalization: Is Retrospective Assessment of Use Before Legalization More Accurate?*, *Journal of Studies on Alcohol and Drugs* **79** (May, 2018) 495–502.

[139] M. Cerdá, M. Wall, T. Feng, K. M. Keyes, A. Sarvet, J. Schulenberg, P. M. O'Malley, R. L. Pacula, S. Galea, and D. S. Hasin, *Association of State Recreational Marijuana Laws With Adolescent Marijuana Use*, *JAMA Pediatrics* **171** (Feb., 2017) 142–149.

[140] D. M. Anderson, B. Hansen, D. I. Rees, and J. J. Sabia, *Association of Marijuana Laws With Teen Marijuana Use: New Estimates From the Youth Risk Behavior Surveys*, *JAMA Pediatrics* **173** (Sept., 2019) 879–881.

[141] E. K. Choo, M. Benz, N. Zaller, O. Warren, K. L. Rising, and K. J. McConnell, *The Impact of State Medical Marijuana Legislation on Adolescent Marijuana Use*, *Journal of Adolescent Health* **55** (Aug., 2014) 160–166.

[142] S. Harper, E. C. Strumpf, and J. S. Kaufman, *Do Medical Marijuana Laws Increase Marijuana Use? Replication Study and Extension*, *Annals of Epidemiology* **22** (Mar., 2012) 207–212.

[143] D. S. Hasin, M. Wall, K. M. Keyes, M. Cerdá, J. Schulenberg, P. M. O'Malley, S. Galea, R. Pacula, and T. Feng, *Medical marijuana laws and adolescent marijuana use in the USA from 1991 to 2014: results from annual, repeated cross-sectional surveys*, *The Lancet Psychiatry* **2** (July, 2015) 601–608.

[144] S. D. Lynne-Landsman, M. D. Livingston, and A. C. Wagenaar, *Effects of State Medical Marijuana Laws on Adolescent Marijuana Use*, *American Journal of Public Health* **103** (Aug., 2013) 1500–1506. Publisher: American Public Health Association.

[145] K. M. Keyes, M. Wall, M. Cerdá, J. Schulenberg, P. M. O'Malley, S. Galea, T. Feng, and D. S. Hasin, *How does state marijuana policy affect US youth? Medical marijuana laws, marijuana use and perceived harmfulness: 1991–2014*, *Addiction* **111** (2016), no. 12 2187–2195. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/add.13523.

[146] N. I. on Drug, *Sex and Gender Differences in Substance Use*, *National Institute on Drug Abuse* (2022).

[147] J. Copeland, W. Swift, and V. Rees, *Clinical profile of participants in a brief intervention program for cannabis use disorder*, *Journal of Substance Abuse Treatment* **20** (Jan., 2001) 45–52.

[148] H. Carliner, P. M. Mauro, Q. L. Brown, D. Shmulewitz, R. Rahim-Juwel, A. L. Sarvet, M. M. Wall, S. S. Martins, G. Carliner, and D. S. Hasin, *The widening gender gap in marijuana use prevalence in the U.S. during a period of economic change, 2002–2014*, *Drug and Alcohol Dependence* **170** (Jan., 2017) 51–58.

[149] S. Kausar and S. Bewley, *Pregnancy after the Age of 40*, *Women's Health* **2** (Nov., 2006) 839–845. Publisher: SAGE Publications Ltd STM.

[150] D. Bruce, T. J. Grove, E. Foster, and M. Shattell, *Gender Differences in Medical Cannabis Use: Symptoms Treated, Physician Support for Use, and Prescription*

Medication Discontinuation, *Journal of Women's Health* **30** (June, 2021) 857–863. Publisher: Mary Ann Liebert, Inc., publishers.