Bank Fees and Household Financial Well-Being*

Michaela Pagel† Sharada Sridhar‡ Emily Williams§

March 3, 2026

In this study, we examine policy changes from large U.S. banks between 2017 and 2022, which eliminated non-sufficient funds (NSF) fees and relaxed overdraft policies. Using individual transaction-level data, we find that the elimination of NSF fees, not surprisingly, resulted in immediate reductions in NSF charges across the income distribution. However, relaxing overdraft policies resulted in reductions in overdraft fees only for wealthier households, along the dimensions of income and liquidity, and only those enjoyed subsequent declines in late fees, interest payments, account maintenance fees, and the use of alternative financial services, such as payday loans. Our results thus suggest that the policy changes were not substantial enough to significantly reduce the financial stress of the more vulnerable households. As our setting features multiple treatments and variation in treatment intensities, we theoretically motivate and empirically implement a new stacked event study estimator closely related to de Chaisemartin and D'Haultfoeuille (2024) to address the biases arising from staggered DID specifications.

Keywords: Bank Fees, Financial Well-Being, Staggered DiD Specifications
JEL codes: G5, D14

*We thank participants at the Advances in Economic Research using Bank Transactions Data at the 2025 Barcelona School of Economics Summer Forum, Boulder Colorado Consumer Finance Conference, Junior Household Finance Virtual Seminar Series, Virtual Household Finance Seminar, Rice University Finance Seminar Series, and Georgia Tech Scheller Academic Seminar. We also thank Emily Gallagher for valuable thoughts and comments as well as Anton Tarasenko and Xinyue Zhao for valuable research assistance.

†Washington University, Olin Business School, NBER & CEPR
‡Georgia Tech, Scheller College of Business
§Harvard Business School

1 Introduction

It is crucial to gain a deeper understanding of how to help households that experience financial stress. While existing work has focused on how changes to income (Agarwal and Qian, 2014; Johnson et al., 2006; Baker, 2018) and credit supply (Gross and Souleles, 2002; Agarwal et al., 2018) affect household spending, there has been less focus on the effects of liquidity shocks when households face hardship. Many financial products charge fees, such as overdraft, late, maintenance, or non-sufficient fund (NSF) fees, specifically at the point of very low liquidity. Furthermore, consumer regulators view overdraft and NSF charges[1] as regressive since they disproportionately affect lower-income households. Banks, on the other hand, justify these fees as a way to manage risk, enable broader access to deposit accounts, and help stressed households with unsecured credit in times of need. Because bank fees charged on deposit accounts have been a serious point of contention (Federal Reserve Board, 2001, 2009; CFPB, 2017, 2024), this paper studies how permanent and temporary relief in the fees charged affect household well-being in the short-, medium-, and longer-term and across the income distribution. We show that even very straightforward policy changes in the provision of consumer financial services can fail to impact the vulnerable households they would have been expected to benefit.

Facing scrutiny from regulators and consumer advocates, a number of large banks restructured their depositor fees from 2017 and 2022. Banks made varying adjustments to different fees. NSF fees saw only extensive margin changes, as nearly all banks eliminated this category of charge. Overdraft policies, on the other hand, changed along the intensive margin, affecting multiple dimensions. These include increasing the minimum overdraft amount before a fee is applied, implementing a 24-hour grace period before charging a fee, and reducing the number of overdrafts that incur a fee. At the core of these reforms lies the assumption that reducing or eliminating certain fees should directly improve financial well-being, reduce liquidity stress, and mitigate financial insecurity; primarily for lower-income more vulnerable

[1] Overdraft fees occur when a customer's bank allows a transaction to proceed despite insufficient funds, often charging a fee for this service. NSF fees, by contrast, are charged when a transaction cannot be processed due to insufficient funds.

households.[2]

The actual impact of these fee reductions on household financial outcomes, however, remains unclear. On the one hand, extending overdraft grace periods could provide much-needed relief to financially stressed households, improving liquidity and reducing reliance on high-cost alternative financial services, such as payday loans or pawnshops. On the other hand, critics argue that these policy changes may not go far enough to address the root causes of financial instability or that their impact may be diluted by other factors, such as ingrained financial habits, access to alternative credit sources, or behavioral responses to changes in fee structures.

We utilize a random subsample, 35 thousand individuals, of transactions-level data from a large account aggregator, covering approximately 50 million individuals, which enables us to observe transaction amounts and descriptions. We can then identify overdrafts, NSF fees, and other indicators of financial well-being, such as late fees, cash-advance fees, credit card debt, and payday loans. We then aggregate the account-level data to an MSA-week level and construct 2019 MSA-level bank market shares using Summary of Deposits (SOD) data from the Federal Deposit Insurance Corporation (FDIC). This provides a metric to determine which MSAs were affected by each policy change and the extent to which they were impacted.

The main empirical challenge is the staggered nature of the twenty-two bank policy changes we study. Since MSAs were treated at different times and to varying degrees, we utilize recent advancements in staggered difference-in-differences (DID) methodologies (Baker et al., 2022; Gardner, 2022; Callaway and Sant'Anna, 2021; Abraham and Sun, 2021), specifically the stacked estimation technique. Since MSAs can be treated many times, we innovate beyond the standard control groups of never-treated and not-yet-treated MSAs. Using never-treated MSAs as controls raises concerns about parallel trends, as these areas lack any top-20 bank branches, making them fundamentally different from treated MSAs. Conversely, using not-yet-treated MSAs as controls limits statistical power—after the first policy change in 2017, most MSAs are treated to some degree, leaving only two or three policy shocks available

[2]These motivations originate directly from the banks themselves. See, for example, press releases from Bank of America, TD Bank, and Citibank.

for analysis. Given that treatment and control groups are based on our bank-share metric rather than a direct institutional measure, further reducing the number of policy shocks weakens our statistical power. Instead, we theoretically motivate and empirically implement a novel event study estimator designed to address biases introduced by staggered treatment adoption in DID settings. Building on the work of de Chaisemartin and D'Haultfoeuille (2024); de Chaisemartin and D'Haultfoeuille (2023a,b), our estimator allows for multiple treatments and adjusts for variations in the intensity of policy changes across different regions.

Our methodology measures the average effect of an additional treatment, conditional on the past treatments. We implement our design within the stacked framework of Baker et al. (2022) to execute our strategy. In the standard approach, each stack contains first-treated units of a given event, with clean controls. The model is then analyzed with unit-by-stack fixed effects, ensuring the captured estimate is unbiased. With our approach, each stack consists of units treated by a given event, with a control group that matches the pre-treatment intensity distribution of the treated group in the stack. As the treated and control groups have been pre-treated in similar ways and all fixed effects are interacted with the stack indicator, creating this ex-ante match mitigates concerns of early versus late comparisons. Note, the standard stacked estimation nests into our framework.

Using the MSA-week panel as well as the treatment events and bank shares, we first simulate a binomial outcome variable, as, our main real-world outcome variable, NSF fees, contains a considerable share of zeros. In turn, we modify the outcome variable to be affected by each bank's policy events proportional to each MSA's bank shares. We then use a number of different event-study estimators to try and recover the true effects. We show that the staggered design yields severely biased estimates and that the stacked designs with never- or not-yet-treated controls yield very noisy estimates that are biased towards zero. In contrast, using all events with pre-treatment exposures being matched between treatment and control yields estimates that are close to the true coefficients.

We then turn to the actual data. We first examine the immediate impact of banks' policy changes on NSF and overdraft fees. In the aggregate data, in the

same quarter as the policy intervention, we observe a significant drop in NSF fees. Within 18 months, fees decline by 10.3% relative to the average MSA-quarter, which is within estimation error of the average treatment effect (i.e., the average bank share of treated MSAs) of 6.8% and thus translates to -100% decline in NSF fees. This result is in line with banks adjusting NSF fees on the extensive margin, eliminating them completely. Next, we analyze both the number of overdrafts and overdraft fees, finding a reduction in both. However, this effect is delayed by two quarters, with a smaller magnitude of only 4.89% relative to the average MSA-quarter.

We then analyze the distributional effects of bank fee changes and find that the benefits are much larger for wealthier MSAs than for less-wealthy ones. For each individual, we estimate the average monthly income based on observed "payroll" transactions and compute the median across MSAs. We then divide MSAs into "high-income" and "low-income" groups based on this median. Only richer MSAs experienced a significant decline in overdraft fees. This disparity underscores the original policy tension: intensive margin fee changes primarily benefit wealthier households, whereas extensive margin adjustments affect all consumers. Addressing concerns about MSA-specific confounding effects, we replicate our findings when conducting the analysis separately for higher- and lower-income households within each MSA.

Next, we explore spillover effects on household balance sheets. We find a significant reduction in interest payments and late fees, with declines of 5% to 6% relative to the mean MSA-quarter. We also observe a decrease in account maintenance charges. Since these fees can often be waived if a minimum balance is maintained[3], the policy changes likely improved liquidity, allowing households to meet minimum balance requirements more easily. Additionally, we find a significant effect on payday, FinTech, or other high-interest loans. However, again, all of these effects are concentrated in higher-income MSAs whereas lower-income MSAs do not appear to be benefiting in a measurable way. Again, we find similar results when comparing poor and rich households within MSAs.

Incorporating liquidity as a second dimension reveals that high income, high liq-

[3]See `https://www.bankrate.com/banking/checking/how-to-avoid-monthly-fees/`.

uidity households consistently experience the largest gains, while low income, low liquidity MSAs see no effect. That said, poorer geographies with greater liquidity experience positive spillover outcomes, though smaller than their higher-income counterparts.

The rest of the paper is organized as follows. Section 2 reviews the related literature and Section 3 describes the data used in our analysis. Section 4 details our empirical strategy, including the construction of our novel control groups, which are theoretically motivated in Subsection 4.2 and backed by simulation results in Subsection 4.3. In turn, Section 5 presents our findings. Finally, Section 6 concludes.

2 Literature Review

Research on NSF and overdraft fees is present in the fields of behavioral economics, consumer finance, and banking regulation. These fees are scrutinized due to their disproportionate effect on low-income and financially vulnerable individuals, as well as their role in enhancing bank profitability.

A significant strand of research focuses on how overdraft and NSF fees impact consumer behavior and welfare. Many studies highlight that banks often fail to adequately disclose fees, leading to consumer confusion and financial hardship. Researchers have argued that banks might encourage customers to overdraw their accounts, assuming that they will either repay quickly or pay hefty fees for overdrafts. This raises concerns about whether overdraft fees are genuinely in the best interest of consumers or are simply a predatory loan revenue source for banks (Avery et al., 2008; Campbell, 2011).

The decision-making process surrounding overdrafts often involves behavioral biases. Studies have shown that individuals may not fully understand the costs of overdrafts, often relying on mental shortcuts (heuristics) that lead to suboptimal choices, and are inattentive (Stango and Zinman, 2009; Medina, 2021; Carlin et al., 2023; Olafsson and Pagel, 2025; Carvalho et al., 2024). Consumers tend to underestimate the frequency with which they incur fees, leading to poor financial planning. Research also delves into the emotional and psychological impact of frequent over-

draft fees on consumers. These fees are linked to feelings of financial insecurity, embarrassment, and stress, which can exacerbate existing financial difficulties (Hogarth and Hilgert, 2002).

A number of papers focus on the regulatory environment surrounding overdraft fees. In the US, for instance, the Federal Reserve has introduced regulations to limit the scope of overdrafts and the ways in which banks can charge fees for them (Tufano, 2009). Research has examined the effectiveness of such regulations in reducing consumer harm while maintaining bank profitability (Tufano, 2009; DiMaggio et al., 2024). Several other studies explore the effects of the CARD Act (2009) and Dodd-Frank Act (2010), which introduced more transparency and consumer protections related to overdraft fees. Researchers have found mixed results, with some evidence that consumers benefit from such reforms but also concerns about unintended consequences, such as the reduction in access to certain banking services (Linden, 2010). An alternative fee structure, such as offering more affordable short-term credit or incorporating a low-cost overdraft protection option, could better serve consumers without sacrificing bank profitability (Ben-David, 2012).

Finally, studies have looked at the pricing strategies banks use to structure their overdraft fees, such as differentiating between types of transactions or creating tiered fee structures. These practices have raised concerns about price discrimination, where banks charge higher fees to consumers who are least able to afford them (Baker and McIndoe-Calder, 2016). Finkelstein (2008) shows that in markets with limited banking competition, fees may be higher, and consumers may have fewer alternatives. Conversely, increased competition in financial markets may drive down overdraft fees, though such effects are slow to materialize.

3 Data

3.1 Transactions-Level Bank Account Data

Description In this paper, we use bank account and credit card transaction data from a large U.S. account aggregator. Partnering with thousands of financial institu-

tions and FinTech firms, they link multiple financial accounts to a single individual while stripping private individual information. While the overall sample includes 50 million individuals, we use a random subsample of about 35 thousand people for an unbalanced panel between January 2018 and May 2023. Each observation consists of an individual identifier, a bank account identifier, a transaction amount, the day of the transaction, a description of the transaction, the zip code of the transaction location (if it is at a point-of-sale terminal), and a broad category name (i.e., "Groceries," "Rent," etc). We use the most frequently observed transaction zip code for an individual across all their accounts to identify their home market.

To answer our question, we must be able to identify transactions that are overdrafts and non-sufficient fund fees. We do so by manually searching within the description variable for keywords, such as "overdraft," "insufficient," "od fee," or "od item," to name a few. In turn, we restrict the relevant outgoing transactions to those below 100 USD. Given our interest in the spillover effects of bank policy changes on other indicators of financial well-being, we also do description searches for late fees, interest charges, maintenance fees, early wage access, payday loans, and FinTech loans. Credit card interest charges are categorized by the data provider as such.

For our heterogeneous results, we also identify payroll transactions, as well as, relative measures of liquidity. Income is defined as an inflow into the account with a description containing the words "payroll" or "income." Given we do not observe the starting balance of accounts, we can only measure relative liquidity levels. We calculate this using a running sum of all inflows and outflows for each individual-week pair. Then, for each individual, across all weeks, we categorize the percentile, in bins of 5%, of the level of the running sum.

Once the fee transactions have been identified, we collapse the individual-week data to the MSA-week level. For aggregation, all finance fees, loans, and payroll variables are summed and the percentile of relative liquidity are averaged.

Summary Statistics In this section, we quantitatively describe the transactions data being used. First, we want to understand the trends of non-sufficient fund fees and overdraft fees. Figure 1 shows the reduction of both these fees over time. A

sharp decline in both categories of fees in March 2020 stands out, as indicated by the vertical dashed line. This is the month of the first round of stimulus checks that were part of the CARES Act in response to COVID-19. A reduction of bank fees is expected with this injection of liquidity. The first policy change that occurred in our sample was in September 2020. After March 2020, there is a clear downward trend in non-sufficient fund fees until the end of the sample. Overdraft fees also have a downward trend, though less pronounced.

Next, we consider summary statistics of the transactions data. In Table 2, we report means and percentiles for spending, income and liquidity variables, as well as, NSF fees, overdraft fees, late fees, credit card interest payments, and other unsecured short-term loans. We report summary statistics for the individual-level transactions data as well as aggregated to the MSA-quarter level which corresponds to the dynamic regression coefficients. Notice, on an individual level, the distribution of the percentile of relative liquidity is by construction – the fifth percentile is 5%, the twenty-fifth percentile is 25% and so on. On the MSA quarter level, this distribution narrows, as we take the mean for every MSA quarter. As we collapse across all individuals, rather than conditioning on individuals who ever paid an overdraft fee or ever incurred an insufficient fund fee, even the aggregated bank fee variables contain a significant share of zeros. The average spending and income transactions, are in line with representative survey data from the Bureau of Labor Statistics Consumer Expenditure Survey[4] and the American Community Survey from the U.S. Census[5].

3.2 Policy Changes

The Consumer Finance Protection Bureau (CFPB) has long followed bank overdraft and non-sufficient fund fee policies, providing data summaries and research reports. As banks began to change their fee policies, favoring depositors, the CFPB documented these changes for the top 15 financial institutions in the U.S. Figure 2 is a snapshot of their compilation. For each bank in the table, they report fee revenue

[4]See `https://www.bls.gov/news.release/cesan.nr0.htm`.
[5]See `https://www.census.gov/programs-surveys/acs.html`.

in 2021, and various types of policy changes. Two patterns can be noticed. First, every large bank completely eliminated their non-sufficient fund fee. Second, besides Citibank, Capital One, and USAA Federal Savings Bank, overdraft fee changes were made on the intensive margin rather than the extensive margin. That is, instead of their complete removal, many financial institutions are now allowing for a 24-hour grace window, increasing the minimum amount withdrawn before an overdraft fee was charged, or reducing the number of overdrafts an individual could be charged for.

In order for us to study the impact of these changes in policies, we need the dates of the policy changes. Because overdraft and non-sufficient fund fees have always been controversial, banks were eager to put out news releases about their relaxation of these charges to their customers. Utilizing bank press releases and news articles reporting on the policy changes, we are able to pinpoint when they went into effect. Table 1 reports these dates. For most policy changes, we are able to find the exact day they went into effect. For a few, we can only identify the month or the quarter. In total, we have 22 policy changes across 15 large banks in the U.S.

3.3 Bank Market Shares

Our transactions-level data does not identify the financial institution of each observed account. This is key to determine treatment, as policy changes are bank-specific. To proxy for the banks individuals have an account with, we construct 2021 MSA-level bank market shares using bank branch data from the summary of deposits data provided by the Federal Deposit Insurance Corporation (FDIC)[6]. For example, an MSA in Greensboro-High Point, NC, has 10 branches of Woodforest Bank, 4 branches of PNC Bank, 10 branches of Bank of America, 18 branches of Wells Fargo, and a total of 144 bank branches across all banks. The corresponding bank shares for Woodforest, PNC, Bank of America, and Wells Fargo are .07, .03, .07, and .13. We use these shares as intensity of treatment, as further described in Section 4.1.

[6]See `https://www7.fdic.gov/sod/dynaDownload.asp?barItem=6`

4 Empirical Strategy

4.1 Set-up

We need to be cognizant of two empirical challenges before setting up our model.

First, as just described, while the policy changes are at the bank level, the transaction data does not link individual accounts and financial institutions. To overcome this obstacle, we define an MSA to be treated by, or "exposed" to, a bank's policy change if the bank has branches in that MSA in 2019. For example, as detailed in Section 3.3, Wells Fargo's exposure in the Greensboro-High Point, NC, MSA is .13. We utilize both extensive and intensive margin measures of exposure. The extensive measure is 0 or 1 – did a bank with a branch in an MSA have a policy change or not? The intensive margin considers every incremental policy change. For example, again consider the MSA in Greensboro-High Point, NC. Before our sample begins in 2018, Wells Fargo made an overdraft policy change in 2017, and Bank of America made two policy changes in 2014. This means in January 2020, this MSA has a cumulative pre-treatment intensive margin exposure of 1*.13 + 2*.07 = .27, but an incremental intensive margin exposure of 0. When Wells Fargo makes an additional policy change in September 2020, this cumulative pre-treatment intensive margin exposure increases to .27 + 1*.13 = .40, with the incremental margin exposure just .13. While we use the extensive margin exposure as treatment for all our empirical specifications, for models that use all bank policy changes, we will use the increasing intensive margin exposure to construct control groups.

Second, as policy changes occur in different weeks over the three year period, we must address the issue of biased estimates in staggered DID designs, as discussed in the growing literature (Baker et al., 2022; Gardner, 2022; Callaway and Sant'Anna, 2021; Abraham and Sun, 2021). Following Baker et al. (2022) and Cengiz et al. (2019), we consider a stacked framework.

The new DID methodologies were developed to mitigate time-varying confounds (Baker et al., 2022; Gardner, 2022; Callaway and Sant'Anna, 2021; Abraham and Sun, 2021). This occurs as early-treated groups in control groups for later treatments can

bias estimates, even if the standard assumptions of parallel trends and no anticipation hold. The solution of many of these strategies is to use the "never treated" or the "last to be treated" groups as controls.

The standard stacked estimation design, which considers only the first treatment and never-treated controls, is limited in scope for our setting for two reasons. As banks strategically open and close branches, exposed MSAs are not randomly treated; this implies never-treated MSAs are a poor control group. Second, and this is true for the approach using not-yet-treated controls as well, limitation to the first treatment of each unite significantly reduces estimation power.

We thus develop a modified stacked differences-in-differences (DID) model that addresses both the empirical hurdles and temporal sources of bias.

4.2 Theoretical Discussion of New Event Study Estimator

Following Wing et al. (2024), let $i = 1, \ldots, N$ index MSAs, and $t = 1, \ldots, T$, index week-of-years. For each MSA i, let $A_i = \{A_{i_1}, \ldots, A_{i_\ell}\}$ be the ordered set of calendar periods when i is treated. The sub-subscript denotes the order of treatment, and for each $j \in \{1, \ldots, \ell\}$, $1 \leq i_j \leq T$. If an MSA is never treated, assume $A_i = \{\infty\}$.

We next want to represent causal relationships using potential outcomes. For each MSA i, for each $z \in A_i, z \neq A_{i_1}$, let z_{-d} represent the treatment date d treatments before z. $Y_{it}(z_{-1})$ represents the outcome MSA i would experience in calendar period t under the hypothetical scenario where i is never treated at time z, but *was* treated up to and including time z_{-1}; that is, the counterfactual state where $\{A_{i_1}, \ldots, z_{-1}\} \subset A_i$, but $z \notin A_i$. Similarly, let $Y_{it}(z)$ represents the outcome i would experience at t if treated up to including z; that is $\{A_{i_1}, \ldots, z_{-1}, z\} \subset A_i$.

In our modified DID setting, we also consider a pre-z treatment condition, $\Omega_z(\Theta_i, \theta_{i_{-1}})$, conditional on the degree of policy exposure for both the treatment and the control group. As discussed in Section 3.3, we measure if an MSA is treated based on its share of bank branches for the institution making a policy change. We define $\Theta_i = \{\theta_{i_1}, \ldots, \theta_{i_\ell}\}$ to be the set of incremental bank branch share exposures, corresponding to each element in ordered set A_i. For example, consider i, the MSA of

Greensboro-High Point, NC as discussed in Section 3.3. Its first event A_{i_1} in our sample is September 2020, when Wells Fargo made an policy change, with its respective bank share .13. This means $\theta_{i_1} = .13$. We also define $\theta_{i_{-1}}$, which is the cumulative exposure before our sample starts. For example, for the MS m of Greensboro-High Point, NC this would be its share of Wells Fargo, .13 which made a policy change in 2017, plus two times its share of Bank of America branches, $2 * .07 = .14$, which made policy changes in 2014 and 2017, resulting in $\theta_{i_{-1}} = .27$. MSAs that are never treated have $\Theta_i = \{0\}$ and $\theta_{i_{-1}} = 0$.

The causal effect of treatment up to and including z relative to treatment up to and including z_{-1} is $\beta_{it}(z) = Y_{it}(z) - Y_{it}(z_{-1})$. The subscripts of β_{it} means the treatment effect can differ across MSAs and time periods. In our modified DID setting, we are interested in the average incremental treatment effect on the treated (AITT) in a particular week. This distinction is important, as we avoid the "bad" or biased effect (Wing et al., 2024, pp 12) by estimating the effect of every *additional* treatment, instead of measuring the effect from no treatment to treatment. We denote this as

$$AITT(z,t) = \mathbb{E}(\beta_{it}(z)|z \in A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})).$$

The realized outcome at any event time $z \in \cup_{i=1,\ldots,N}, A_i$ is defined as the outcome at time z_{-1} plus the marginal treatment outcome at time z, if treatment at time z actually occurs.

$$Y_{it} = Y_{it}(z_{-1}) + \beta_{i_t}(z) * \mathbb{1}\{z \in A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})\} \tag{1}$$

In order to recover causal relationships in our DID setting we need two main assumptions: a non-anticipation assumption, and a common trends assumption.

> **Assumption 1** *No Anticipation: The average causal effect of being treated at time z is zero for all calendar periods between z_{-1} and z. That is, for period $z_{-1} < t < z$, if pre-treatment condition Ω_z holds,*
>
> $$\mathbb{E}(Y_{it}(z) - Y_{it}(z_{-1})|z \in A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})) = 0.$$

Assumption 2 *Common Trends: In the absence of treatment exposure at time z, the average change across post-z treatment time periods would be the same in the treated group, $\forall i \in \{1, \ldots, N\}$ such that $z \in A_i$, and the comparison group, $\forall i \in \{1, \ldots, N\}$ such that $z \notin A_i$ and a pre-treatment condition Ω_z holds. For periods $t > z$,*

$$E(Y_{it}(z_{-1}) - Y_{it-1}(z_{-1}) \mid z \in A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})) = E(Y_{it}(z_{-1}) - Y_{it-1}(z_{-1}) \mid z \notin A_i, \Omega_z(\Theta_i, \theta_{i_{-1}}))$$

As in the standard DID framework, these two assumptions are not testable, as they depend on unobserved counterfactual data.

We can now show how our estimation strategies, under the no anticipation and common trends assumptions, will recover the average incremental treatment effect $AITT(z,t)$,

$$\begin{aligned}
\Delta_{IT}^{z,z+k} &= \mathbb{E}(Y_{i,z+k} - Y_{i,z_{-1}} | z \in A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})) - \mathbb{E}(Y_{i,z+k} - Y_{i,z_{-1}} | z \notin A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})) \\
&= \mathbb{E}(Y_{i,z+k}(z) - Y_{i,z_{-1}}(z) | z \in A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})) \; - \\
&\quad \mathbb{E}(Y_{i,z+k}(z_{-1}) - Y_{i,z_{-1}}(z_{-1}) | z \notin A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})) \\
&= \mathbb{E}(Y_{i,z+k}(z) - Y_{i,z_{-1}}(z_{-1}) | z \in A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})) \; - \\
&\quad \mathbb{E}(Y_{i,z+k}(z_{-1}) - Y_{i,z_{-1}}(z_{-1}) | z \notin A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})) \\
&= \mathbb{E}(\beta_{i,z+k} | z \in A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})) \; + \\
&\quad \Big[\mathbb{E}(Y_{i,z+k}(z_{-1}) - Y_{i,z_{-1}}(z_{-1}) | z \in A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})) \; - \\
&\quad \mathbb{E}(Y_{i,z+k}(z_{-1}) - Y_{i,z_{-1}}(z_{-1}) | z \notin A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})) \Big] \\
&= E(\beta_{i,z+k} | z \in A_i, \Omega_z(\Theta_i, \theta_{i_{-1}})) \\
&= AITT(z, z+k)
\end{aligned}$$

The second equality substitutes the potential outcomes, and the third equality holds by the first assumption. Substituting the treated outcomes expressed as the untreated outcome plus the treatment effect (1) we obtain the fourth equality. The fifth equality comes from re-writing $Y_{i,z+k}(z)$ in terms of $\beta_{i,z+k}$ and $Y_{i,z+k}(z_{-1})$. And finally, repeatedly applying the common trends assumption brings us to our desired

conclusion.

Next, we detail our baseline specification, and run simulations on this model, using different control group construction innovations. We find that the best performing control group matching the ex-ante distribution of incremental bank policy changes of the treated, utilizing all policy change events and rigorously addressing early versus late treatment comparisons.

4.3 Baseline Specification

Every week a bank policy change occurs is a treatment event, with its own treatment-specific "stack" s. Each stack is its own DID event study. We then append all the individual stacks s into a single dataset and measure β, the weighted average of each stack's treatment effect. We estimate dynamic regressions, where an event time k, β_k measures the difference between treated and control observations at k relative to the period before the policy change, $k = -1$. For each stack, we use data within a specified window, K relative periods before and K relative periods after treatment.

Our model is as follows

$$y_{s,i,t} = \sum_{k \in [-K,K], k \neq -1} \beta_k \mathbb{1}_{s,i}\{t = k\} + X_{s,i,t} + \theta_{s \times i} + \phi_{s \times m \times y} + \epsilon_{s,i,t} \tag{2}$$

The outcome variable for a given stack s, MSA i, in week t, is $y_{s,i,t}$. The treatment effects β_k is measured via dummies $\mathbb{1}_{s,i}\{t = k\}$, which take on the value of 1 if MSA i is treated in stack s, and week t is within k relative periods pre/post policy change. We control for stimulus payments into bank accounts with $X_{s,i,t}$. We also include a rich set of fixed effects, where $\theta_{s \times i}$ is MSA-by-stack, and $\phi_{s \times m \times y}$ is month-by-year-by-stack. Respectively, this allows us to control, within a stack, for MSA attributes that are time-invariant, and time trends common to all MSAs, Our standard errors are clustered on a state-by-stack level.

4.3.1 Simulation Results

We use simulations to demonstrate the viability of our new estimator. Taking the MSA-week panel of our previous exercise including all the policy changes and treatment intensities determined by the bank shares, we simulate a binomial outcome variable with a zero probability of 50%. This corresponds to our NSF fee outcome, which contains a significant share of zeros in the real data. In turn, we modify the outcome variable to be affected by the events proportionally to their bank shares: whenever there is an event of a certain bank, every MSA's outcome in the first quarter after that change is reduced by its respective bank share times 0.2 if the outcome is one in that week and unaffected if the outcome is zero. In the second quarter the bank share is weighted by 0.4, in the third by 0.6, in the fourth by 0.8, and by the full bank share in all quarters afterwards. This method will imply that the outcome's slope is changed after each event. In turn, the true coefficients we would expect to estimate are given by 0.1, 0.2, 0.3, 0.4, and 0.5 for quarters 1 to 4 and beyond that, given there are 50% zeroes in the outcome variable. We test four stacked DiD methodologies, with two additional approaches and empirical estimations detailed in the Supplemental Appendix.

Method 1: First Policy Change as Treatment, Not-Yet-Treated MSAs as Control This model follows the standard approach in the stacked DID literature, where only the first treatment to each MSA is studied, and the not-yet-treated MSAs make up the control group. This design conservatively ensures the DID estimates are not biased. As this methodology considers only the first policy change to each MSA, we use nine events occurring in different weeks between 2021 and 2022. The control group consists of MSAs that are first treated after the treatment in question. We can represent the pre-treatment condition as

$$\Omega_z(\Theta_i, \theta_{i_{-1}}) = \{i, n | z \in A_i, z = A_{i_1}, \theta_{i-1} = 0, z \notin A_n, z < A_{n_1}, \theta_{n_{-1}} = 0\}.$$

Since these are geographical areas that will be treated in the future, they make up a realistic control group for our treated MSAs. However, as we cannot use all events,

statistical power limited.

Method 2: Treatment of First Policy Change, Control of Never-Treated We consider only the first treatment to each MSA, with never-treated MSAs as the control group. We can think of the pre-treatment condition to be

$$\Omega_z(\Theta_m, \theta_{m_{-1}}) = \{m, n | z \in A_m, z = A_{m_1}, \theta_{m_{-1}} = 0, z \notin A_n, A_n = \{\infty\}, \theta_{n_{-1}} = 0\}$$

Given our direct method of treatment, reducing the number of treatment events limits our power, and as banks do not randomly select bank branch locations, the control group MSAs are very different from the treated MSAs. In both the simulation exercise and empirical implementation, we find very noisy estimates. Analogous to Method 1, this can be attributed to limited statistical power and mismatched controls.

Method 3: All Policy Changes as Treatment, Controls Matched on Average Pre-Treatment Exposure In this method, we consider all policy changes. We construct a group of control MSAs such that each has a cumulative pre-treatment exposure equal to the *average* cumulative pre-treatment exposure of the treated group. By matching the ex-ante degree of policy exposure, we mitigate early versus late comparisons, as the treated and control MSAs are, on average, similarly exposed prior to treatment.

Let N be the number of MSAs treated at time z; that is $N = |\{i | z \in A_i\}|$. The mean cumulative pre-treatment exposure of the soon-to-be-treated group is

$$\overline{\theta_z} = \frac{1}{N}\left[\sum_{i|z\in A_i}\sum_{i_j \in A_i | i_j < z} \theta_{i_j} + \sum_{i|z\in A_i} \theta_{i_{-1}}\right].$$

Correspondingly, the pre-treatment condition is

$$\Omega_z(\Theta_i, \theta_{i_{-1}}) = \{i, n | z \in A_i, z \notin A_n, \theta_{n_{-1}} + \textstyle\sum_{n_j \in A_n | n_j < z} \theta_{n_j} = \overline{\theta_z}\}.$$

When constructing these control groups empirically, we allow for small slack in the equality condition to ensure the size of the control group is at least 10% the size of the treated group. Our results are robust to increasing and decreasing this slack.

Method 4: All Policy Changes as Treatment, Controls Matched on Distribution of Pre-Treatment Exposure In this final method, we again consider all policy changes and match the control and treatment groups based on pre-treatment exposure levels. However, we now construct a control sample that matches the coarsened distribution of incremental pre-treatment exposure for the past d events. Specifically, we minimize the differences between the treated and control groups on six key summary statistics for each past event–mean, minimum, 25th, 50th, 75th percentiles, and maximum.

Let $\boldsymbol{f_z^T}$ be the matrix of these attributes for any set of MSAs that are treated at time z, where each row represents a prior event, and each column a distribution statistic. Analogously, let $\boldsymbol{f_z^C}$ be the matrix containing these attributes for any set of MSAs not treated at time z. The pre-treatment condition is

$$\Omega_z(\Theta_i, \theta_{i_{-1}}) = \{i, n | z \in A_i, z \notin A_n, \arg\min_{i,n} |\boldsymbol{f_z^T} - \boldsymbol{f_z^C}|\},$$

with minimization performed row by row, for each of the d prior events. As before, some treated MSAs will be excluded in stack z if their removal yields a control group with a better distributional match. We implement this coarsened distribution match using the `cem` command in Stata.

Because units are treated multiple times in our setting and each treatment changes the evolution of the outcome of interest over time, matching the treatment and control group based only on the cumulative exposure might be insufficient. Instead it might be relevant to match on the history of treatment events. This is achieved by this method.

Comparison of Simulation Results for all Methods In Figure 3, it is clear that the stacked DiD methods using a matching technique of the history of previous

treatment intensity exposure, as discussed in Section 4.1, yield estimates that are closer to the truth. The staggered method is biased considerably. Given the noise in the outcome variable, the not-yet-treated and never-treated controls that restrict to the first event are not well estimated. Instead, using all events and matching the mean pre-treatment exposure of the treatment and control groups is yielding tighter estimates that are less biased. Matching the distribution of pre-treatment exposure further reduces the bias, and, finally, matching the distribution of the history of pre-treatment exposure performs best.

The evaluation of our methodologies naturally leads us to use Method 4.3.1 for all of our main empirical specifications.

4.4 Heterogeneous Specification

Beyond our main model, we want to study the distributional effect for each of our outcome variables. Our control and treatment groups are constructed as per Method 4.3.1 in Section **??**. For each variable d that may produce heterogeneous effects, we split all MSAs into two categories: above and below the median value of d, M_d, calculated across all MSA-week pairs. For static splits, we compare the within-MSA median, $M_{i,d}$ against M_d. Using the same notation as in Section **??** our model is

$$y_{s,i,t} = \sum_{k \in [-K,K], k \neq -1} \beta_k \mathbb{1}_{s,i}\{t = k\} + \sum_{k \in [-K,K], k \neq -1} \beta_{k,d} \mathbb{1}_{s,i}\{t = k\} \times \mathbb{1}_{s,i}\{M_{i,d} \leq M_d\} \tag{3}$$
$$+ X_{s,i,t} + \theta_{s \times i} + \phi_{s \times m \times y} + \epsilon_{s,i,t}$$

To mitigate MSA specific confounding effects, we further split our MSA-week data into two categories: users above and below the median value of d, M_d, calculated across all unique users. For example, when splitting by income, each MSA–week pair yields two observations: one for users with monthly income below the overall median, and one for users with income above it. The specification is analogous to Equation (3).

For dynamic splits, we compare the value of $d_{s,i,t}$ in a given stack-MSA-week pair to M_d. Again, using the same notation as in Section **??** our model is

$$y_{s,i,t} = \sum_{k\in[-K,K],k\neq -1} \beta_k \mathbb{1}_{s,i}\{t=k\} + \sum_{k\in[-K,K],k\neq -1} \beta_{k,d} \mathbb{1}_{s,i}\{t=k\} \times \mathbb{1}_{s,i}\{d_{s,i,t} \leq M_d\} \tag{4}$$
$$+ X_{s,i,t} + \theta_{s\times i} + \phi_{s\times m\times y} + \epsilon_{s,i,t}$$

For both split-types, we have two sets of coefficients for each outcome $y_{s,i,t}$. The first set is $\{\beta_k\}$ which measures the treatment effect for MSAs above the median M_d, and the second set is $\{\beta_{k,d}\}$ which measures the additive treatment effect, on top of $\{\beta_k\}$, for MSAs below the median M_d. Controls, fixed effects, and standard error calculations exactly follow Section **??**.

In our next session we discuss our empirical results for our baseline and heterogeneous specifications, following Equations 2, 3, and 4.

5 Main Results

In this section, we present the causal effects of changes to non-sufficient funds and overdraft fee policies on household financial well-being.

5.1 Effects on Overdraft and Non-Sufficient Fund Fees

In Figure A6, we study the main effect of the banks' policy changes. The top panel reports the effects of non-sufficient funds (NSF) fees. Starting from the quarter of the bank policy intervention, there is a significant drop in fees. Within a year and a quarter of a bank policy change, the NSF fees in a MSA-week reduce by \$2.18. Relative to an average MSA, this is a 10.3% reduction, calculated as calculated as (\$2.18 per week $\times$ 12 weeks per quarter)/(\$253 in average NSF fees per MSA-quarter). This is within estimation error of the average bank share of a treated MSA of 6.8% and thus corresponds to a -100% decline. Given the policy changes

regarding non-sufficient fund fees were unilaterally at the extensive margin, seeing the immediate drop was expected.

The bottom panel of Figure A6 report the effect of bank policy changes on overdraft fees. While we observe a reduction in fees paid, it is not significant until the fourth quarter. As seen in Table 1, for an average MSA, the total overdraft fee payments equal 2.5 times the total NSF fee payments. However, in response to bank policies, the effect is roughly the same in magnitude for both fees. One explanation for this muted result could be the *intensive margin* nature of the adjustments to overdrafts. Only three banks completely removed this type of fee. The others either allowed for a 24-hour grace window, increased the minimum amount withdrawn before an overdraft fee was charged, or reduced the number of overdrafts an individual could be charged for. Given this, observing only mild effects can be explained by the bank policy changes not being significant enough to help households in need. To further study who is helped, we explore the heterogeneous effects on overdraft fees in Sections 5.3 and 5.4.

5.2 Spillover Effects

After observing a decline in overdrafts and NSF fees, we again apply Equation 2 to four other financial well-being indicators: late fees, interest rate charges (reflecting rolled-over credit card debt), account maintenance fees, and having any high-interest loan (payday, EWA, or other FinTech loans). In Figure 5 we see a significant drop of 50% to 60% ,relative to an average MSA, in all variables, indicating these policies improve household financial well-being. It is important to note that even though 2018 to 2023 was a time period of both decreasing and increasing interest rates, our month-by-stack fixed effects should control for these changes.

Now that we have established both first-order and spillover effects, we want to look to the distributional impact. This will allow us to determine if the effect of the policy was uniform across important dimensions, or if there was disproportionate impact.

5.3 Heterogeneity: Income

The first dimension we study is income. For each individual, we identify their income transactions, as described in Section 3.1. We find that the average payroll per MSA-quarter is $8,638 or just under $3,000. Using a static split, we estimate the treatment effects using Equation 3.

In Figure 6, we plot the results for above median income MSAs and the additional interaction effect for below median income MSAs. Given the straightforward, and expected, decreases to NSF fees for all users, we focus on the effect on overdrafts. We find that high income areas enjoy decreases to overdraft fees, whereas low income areas do not. Looking at spillover effects, Figure 7, shows that poor areas enjoy none of the spillover effects we observed in the baseline figure with the benefits going to the richer areas.

Given it is possible there are MSA-level confounding results, we re-run this specification but splitting within MSA income levels. As seen in Figure 6, the effects are robust to this within-MSA split. Figure 8 shows a similar scissor effect as the across MSA results – however, the differences between poor and rich households is less muted and concentrated in the short term. We also observe significant differences in the pre-period between poor and rich individuals within MSA, indicating the outcome *level* are different for different households. This is not surprising, given household income levels are not randomly assigned.

This is our first piece of evidence that while intensive margin changes to overdraft policies are promoted as a way to alleviate financial burdens for struggling households, they tend to disproportionately benefit wealthier individuals. Moreover, all spillover benefits are reaped by the wealthier households.

5.4 Heterogeneity: Liquidity

The next dimension we consider is liquidity. Detailed in Section 3.1, we can reliably only measure the relative percentile of liquidity for each person in our sample. Averaging across individuals in a MSA quarter, the liquidity range is narrow, from the 50th to the 60th percentile. Given we normalize our data in this manner on an

individual level, we allow for dynamic splits for our model following Equation 4.

Figure 9 shows the reduction in overdrafts is significant but not significantly different between groups. However, also seen in Figure 9, the spillover effects in late fees, account maintenance fees and high interest loans only benefit those with high balances. Where low liquidity geographies do see financial gains are from reduced interest payments, but again this is half the magnitude relative to more liquid households.

These results again show while there are modest gains for the more financially vulnerable due to these bank policies, once again the more advantaged benefit the most. A first-order concern is that income and liquidity are highly correlated, leading to these very similar distributional results. To address this, we dynamically split the sample along liquidity, but statically split along income. If the two dimensions were perfectly correlated, this should result in two identical graphs, which is not what we find.

Examining income or liquidity in isolation masks important distributional effects. When income is considered alone, richer households benefit more from the overdraft fee reductions; when liquidity is considered alone, both low and high liquidity groups fare well. However, Figure 10 shows that the rich with low liquidity reap the benefits while the poor with low liquidity experience a modest increase in overdraft fees, once again highlighting the weakness of the intensive margin lever.

We see similar nuanced results for our spillover effects. Figure 11 shows all high income areas reduce interest payments, but the addition of high liquidity decreases payments even further. Low income, high liquidity areas also see this reduction, but at half the magnitude of their counterparts. Low income, low liquidity households see no change. Finally, only high income, high liquidity households reduce their take-up of high interest loans.

While our initial liquidity results indicated positive findings for all liquidity types, for both intensive and extensive margin changes, our two-way split shows us things are not that clear. The poorest of the poor, who are disadvantaged on the dimensions of liquidity and income, do not experience beneficial changes. And while lower income households with some amount of liquidity have benefits, they are small relative to

richer households.

6 Conclusion

This paper examines the impact of recent changes in depositor fee policies by large U.S. banks on household financial well-being. We find that while the elimination of NSF fees led to immediate cost reductions, the effects of relaxed overdraft fee policies were more gradual and less pronounced. Despite reductions in late payments, interest charges, and account maintenance fees, there is limited evidence of significant changes in households' reliance on high-interest loans, such as payday loans. We estimate these effects using a novel control group construction, matching previously treated bank market shares to the coarsened distribution of the previously treated bank share of the newly treated observations.

Additionally, we observe a distributional impact: wealthier households benefit more from the intensive margin changes to overdraft policies, while poorer households experience zero or relatively smaller spillover effects from both policy changes. These results are robust to splitting either across or within MSA. Interacting with a second dimension of liquidity, we find high income, high liquidity households are always the winner, and low income, low liquidity households have no change or fare worse. However, when considering secondary effects, low income households with higher liquidity do have positive outcomes. Taken together, these findings suggest that the intensive margin bank policy changes did little to significantly improve the financial standing of the financially vulnerable; instead, they benefited the rich.

Overall, this paper contributes to the literature on financial regulation by providing a nuanced view of the impacts of fee reductions, challenging the assumption that such reforms inherently lead to improved household financial health. It highlights the need for more comprehensive banking policies to effectively address underlying financial vulnerabilities.

References

Abraham, S. and L. Sun (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects.

Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2018). Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow? *The Quarterly Journal of Economics*.

Agarwal, S. and W. Qian (2014). Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore. *American Economic Review 104*(12), 4205–4230.

Avery, R. B., K. P. Canner, and G. B. Cook (2008). The impact of overdraft protection on consumer welfare. *Journal of Financial Economics 90*(3), 504–522.

Baker, A. C., D. F. Larcker, and C. C. Wang (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics 144*(2), 370–395.

Baker, A. E. and B. McIndoe-Calder (2016). Price discrimination in overdraft fees: Evidence from bank data. *Review of Financial Studies 29*(7), 1864–1891.

Baker, S. R. (2018). Debt and the response to household income shocks: Validation and application of linked financial account data. *Journal of Political Economy 126*(4), 1504–1557.

Ben-David, I. (2012). Alternative credit models: Lower-cost overdraft protection and consumer choice. *Journal of Banking and Finance 36*(10), 2774–2788.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*.

Campbell, J. Y. (2011). Household finance: The role of overdraft fees and credit cards. *Journal of Finance 66*(4), 1255–1285.

Carlin, B., A. Olafsson, and M. Pagel (2023). Mobile apps and financial decision making. *Review of Finance 27*(3), 977–996.

Carvalho, L., A. Olafsson, and D. Silverman (2024). Misfortune and mistake: The financial conditions and decision-making ability of high-cost loan borrowers. *Journal of Political Economy 132*(9), 3173–3213.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019, 05). The effect of minimum wages on low-wage jobs*. *The Quarterly Journal of Economics 134*(3), 1405–1454.

CFPB (2017). Cfpb sues tcf national bank for tricking consumers into costly overdraft service.

CFPB (2024). Cfpb takes action to stop banks from harvesting overdraft fees without consumers' consent.

de Chaisemartin, C. and X. D'Haultfoeuille (2023a). Two-way fixed effects and differences-in-differences estimators with several treatments. *Journal of Econometrics 236*(2), 105480.

de Chaisemartin, C. and X. D'Haultfoeuille (2023b). Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey. *Econometrics Journal 26*(3), C1–C30.

de Chaisemartin, C. and X. D'Haultfoeuille (2024). Difference-in-differences estimators of intertemporal treatment effects.

DiMaggio, M., A. Ma, and E. Williams (2024). In the red: Overdrafts, payday lending, and the underbanked. *Journal of Finance*.

Federal Reserve Board (2001, December). Attachment 1 to the federal reserve board press release, december 11, 2001. Technical report, Federal Reserve Board, Washington, D.C. Accessed: 2025-01-31.

Federal Reserve Board (2009, November). Federal reserve board press release on bank capital regulation, november 12, 2009. https://www.federalreserve.gov/newsevents/pressreleases/bcreg20091112a.htm. Accessed: 2025-01-31.

Finkelstein, A. (2008). The impact of market concentration on bank fees: Evidence from overdraft charges. *Review of Economics and Statistics 90*(4), 829–842.

Gardner, J. (2022, July). Two-stage differences in differences. *arXiv:2207.05943*.

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

Hogarth, J. R. and M. J. Hilgert (2002). Financial literacy, education, and inclusion: The role of overdraft protection. *Journal of Consumer Affairs 36*(2), 226–246.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review 96*(5), 1589–1610.

Linden, J. (2010). Consumer protection in retail banking: Overdraft regulation and its impact. *Journal of Consumer Policy 33*(1), 19–38.

Medina, P. C. (2021). Side effects of nudging: Evidence from a randomized intervention in the credit card market. *The Review of Financial Studies 34*(5), 2580–2607.

Olafsson, A. and M. Pagel (2025). The ostrich in us: Selective attention to personal finances. *Review of Economics and Statistics*, 1–26.

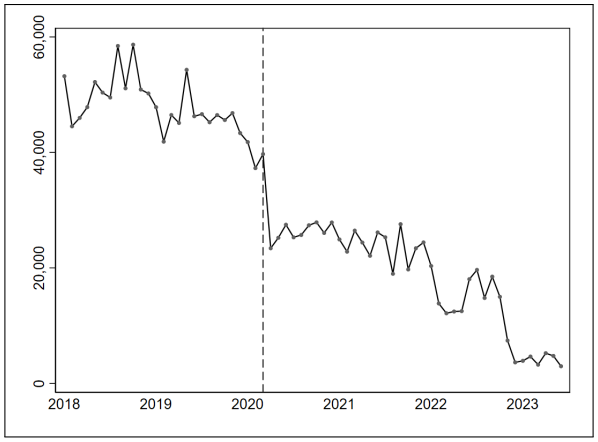
Rambachan, A. and J. Roth (2023). A more credible approach to parallel trends. *Review of Economic Studies*.

Stango, V. and J. Zinman (2009). Exponential growth bias and household finance. *Review of Financial Studies 22*(7), 2807–2849.

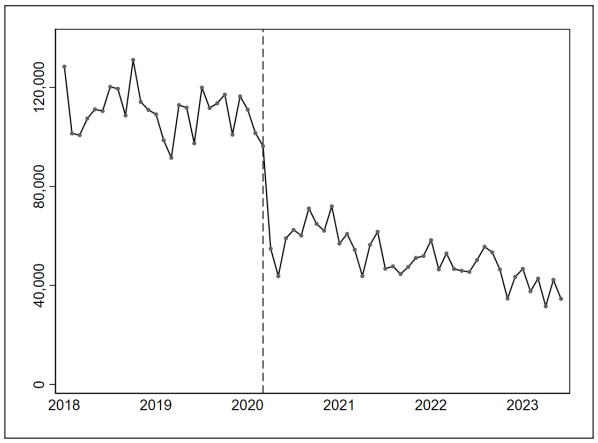
Tufano, P. (2009). Regulating overdraft protection: A case study of the u.s. experience. *Journal of Financial Economics 93*(3), 243–263.

Wing, C., M. Yozwiak, A. Hollingsworth, S. Freedman, and K. Simon (2024). Designing difference-in-difference studies with staggered treatment adoption: Key concepts and practical guidelines. *Annual Review of Public Health 45*, 485–505. Accessed: 2025-05-05.

Figures



(a) Non-sufficient Fund Fees ($)



(b) Overdraft Fees ($)

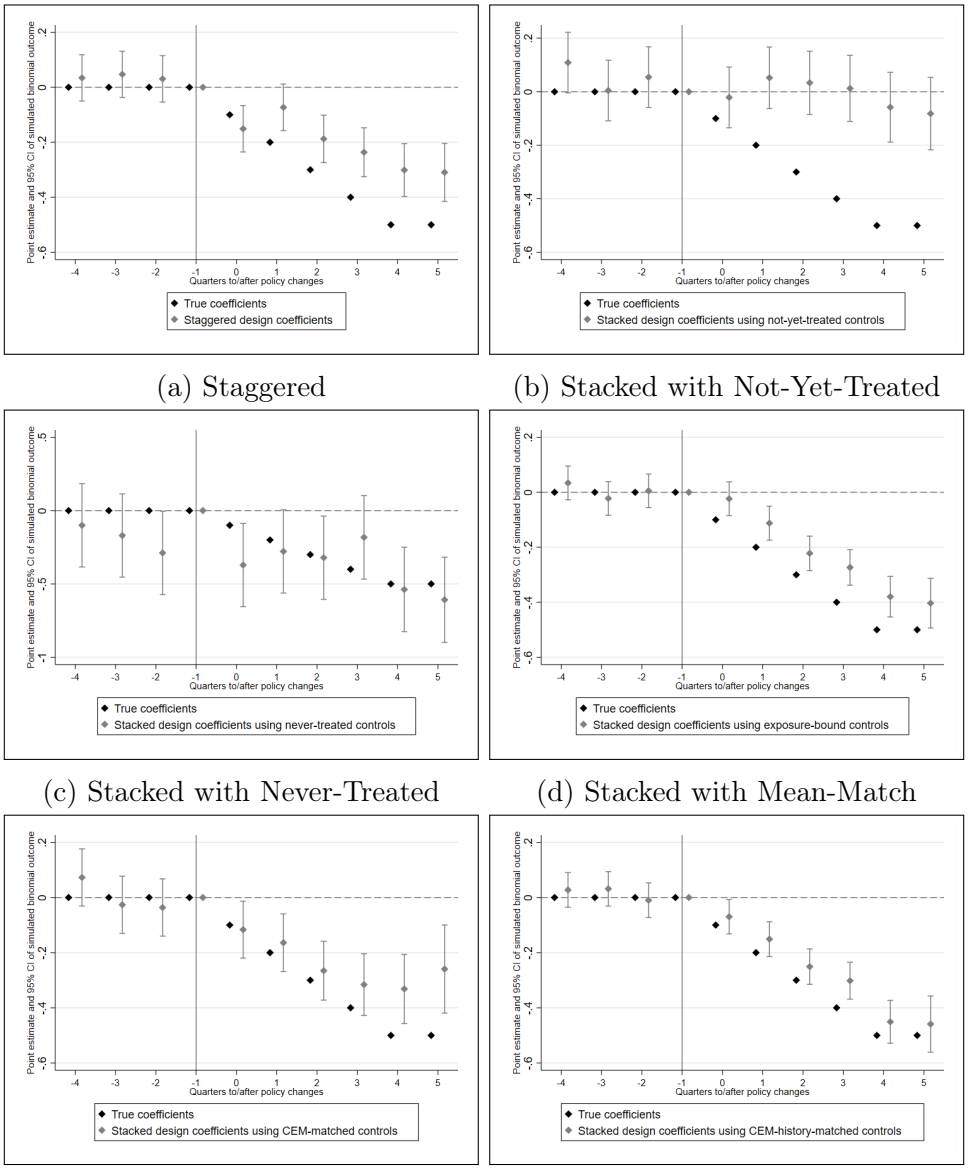
Figure 1: Bank Fees by Month-Year

Notes: This figure shows the total amount of bank fees by month between January 2018 and May 2023. Figure (a) plots the amount of non-sufficient fund fees over time, and Figure (b) plots the amount of overdraft fees over time. The dotted line at March 2023 indicates when the first stimulus payments were distributed via the CARES act in response to COVID. As can be seen, this liquidity reduced bank fees. The first bank policy change in our sample occurred 6 months later, in September 2020. We study 22 bank policy changes made between 2018 and 2023.

Bank	Overdraft/NSF Revenue Reported for 2021	No overdraft fees on any transactions[1]	No NSF fees	No overdraft fees on debit card purchases[2]	No overdraft fees on ATM withdrawals	No extended/ sustained overdraft fees	Size of overdraft and/or NSF fee[3]	Daily limit on number of overdraft/NSF fees[4]	Cushion before overdraft fee is charged[5]	Extended grace period
Truist Bank	$415 million		✓			✓	$36	3 ($108)	$5	--
U.S. Bank N.A.	$338 million		✓			✓	$36	3 ($108)	$50	next day
Regions Bank	$300 million		✓			✓	$36	3 ($108)	$5	next day
PNC Bank, N.A.	$269 million		✓			✓	$36	1 ($36)	$5	next day
USAA Federal Savings Bank	$197 million		✓	✓	✓	✓	$26	1 ($29)	$50	--
Huntington National Bank	$152 million		✓			✓	$15	3 ($45)	$50	next day
Citizens Bank, N.A.	$145 million		✓			✓	$35	5 ($175)	$5	next day
Woodforest National Bank	$145 million		✓			✓	$32	3 ($96)	$1	--
KeyBank N.A.	$116 million		✓				$20	3 ($60)	$20	--
First National Bank Texas DBA First Convenience Bank	$108 million						$34	2 ($68)	$25	next day

Figure 2: CFPB Bank Fee Policy Changes Snapshot

Notes: The CFPB compiled a list of overdraft and insufficient fund fee policy changes for the top 20 banks in the U.S.. This is a partial snapshot of the table. The complete table can be found on https://files.consumerfinance.gov/f/documents/cfpb_overdraft-table_2024-04.pdf. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020.



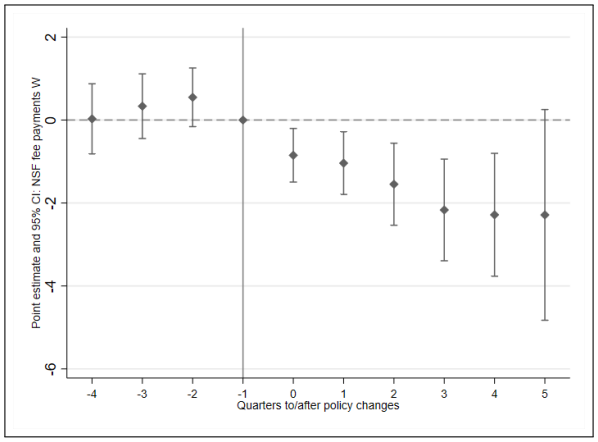
(a) Staggered (b) Stacked with Not-Yet-Treated

(c) Stacked with Never-Treated (d) Stacked with Mean-Match

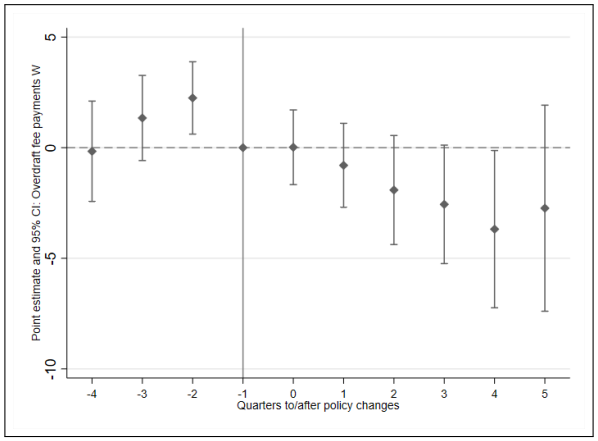
(e) Stacked with CEM-Match (f) Stacked with CEM-History-Match

Figure 3: Simulation results for the different DiD event study methods

Notes: We take the MSA-week panel of our previous exercise including all the policy changes and treatment intensities determined by the bank shares. We then simulate a 0/1 outcome variable with a probability to obtain a 0 or 1 given by 50%. In turn, we modify the outcome variable to be affected by the events proportionally to their bank shares: whenever there is an event of a certain bank, every MSA's outcome in the first quarter after that change is reduced by its respective bank share times 0.2 (if the outcome is one in that week and unaffected if the outcome is zero). In the second quarter the bank share is weighted by 0.4, in the third by 0.6, in the fourth by 0.8, and by the full bank share in all quarters afterwards. In turn, the true coefficients we would expect to estimate are given by 0.1, 0.2, 0.3, 0.4, and 0.5 for quarters 1 to 4 and beyond that. We then compare a staggered event study estimator, a stacked estimator with never-treated and not-yet-treated controls that restrict to the first event for each MSA. And finally estimators using all events and matching the mean pre-treatment exposure, the distribution of pre-treatment exposure, and the distribution of the history of pre-treatment exposure. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020.



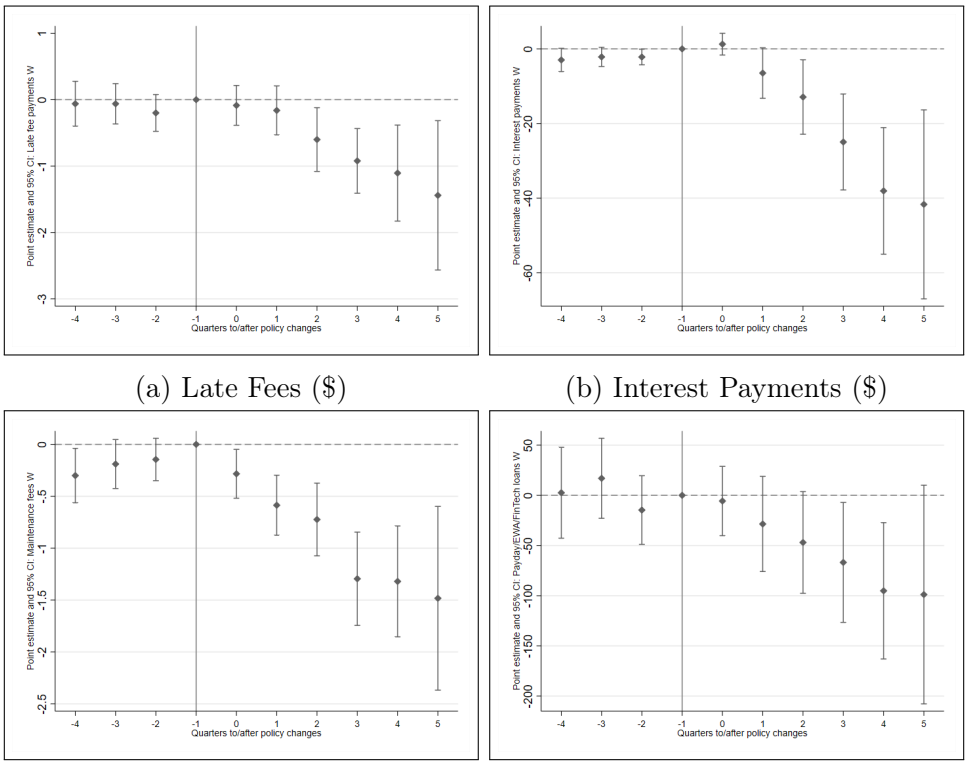
(a) Non-sufficient fund fees ($)



(b) Overdraft Fees ($)

Figure 4: Effects of Bank Policy Changes on Depositor Fees

Notes: This figure shows the average incremental effect a change in bank fee policy has on both overdraft and non-sufficient fund fees. The effects are estimated using Equation 2, i.e., the stacked estimation design using all events and for each treatment MSA matching the distribution of the history of pre-treatment exposure to form the control group, controlling for MSA times stack and month-by-year times stack fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state-by-stack level.



(a) Late Fees ($)

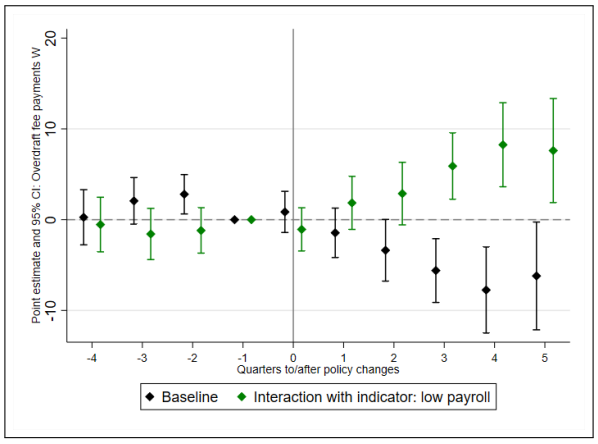
(b) Interest Payments ($)

(c) Account Maintenance Fees ($)

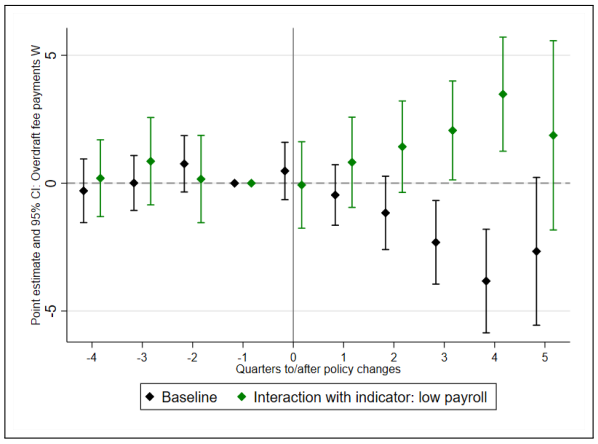
(d) High Interest Loans ($)

Figure 5: Spillover Effects of Bank Policy Changes

Notes: This figure shows the average incremental spillover effects a change in bank fee policy (targeting overdraft and non-sufficient fund fees) has on late fees, interest charges, and any high interest loans. The effects are estimated using Equation 2, i.e., the stacked estimation design using all events and for each treatment MSA matching the distribution of the history of pre-treatment exposure to form the control group, controlling for MSA times stack and month-by-year times stack fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state-by-stack level.



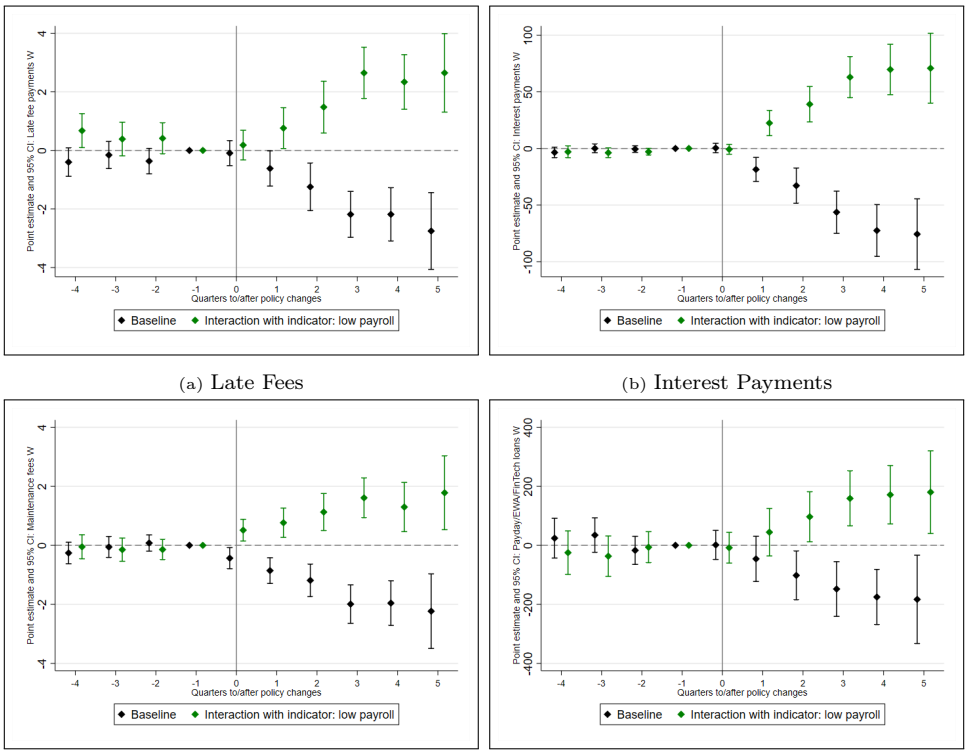
(a) OD Fees, By Income Across MSAs



(b) OD Fees, By Income Within MSA

Figure 6: Effects of Bank Policy Changes on Depositor Fees: By Income

Notes: This figure shows the average incremental effect a change in bank fee policy has on both overdraft and non-sufficient fund fees, by income group. The results for **above median income** MSAs are in **black** and the results for **below median income** MSAs are in **green**. The effects are estimated using Equation 2, i.e., the stacked estimation design using all events and for each treatment MSA matching the distribution of the history of pre-treatment exposure to form the control group, controlling for MSA times stack and month-by-year times stack fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state-by-stack level.



(a) Late Fees

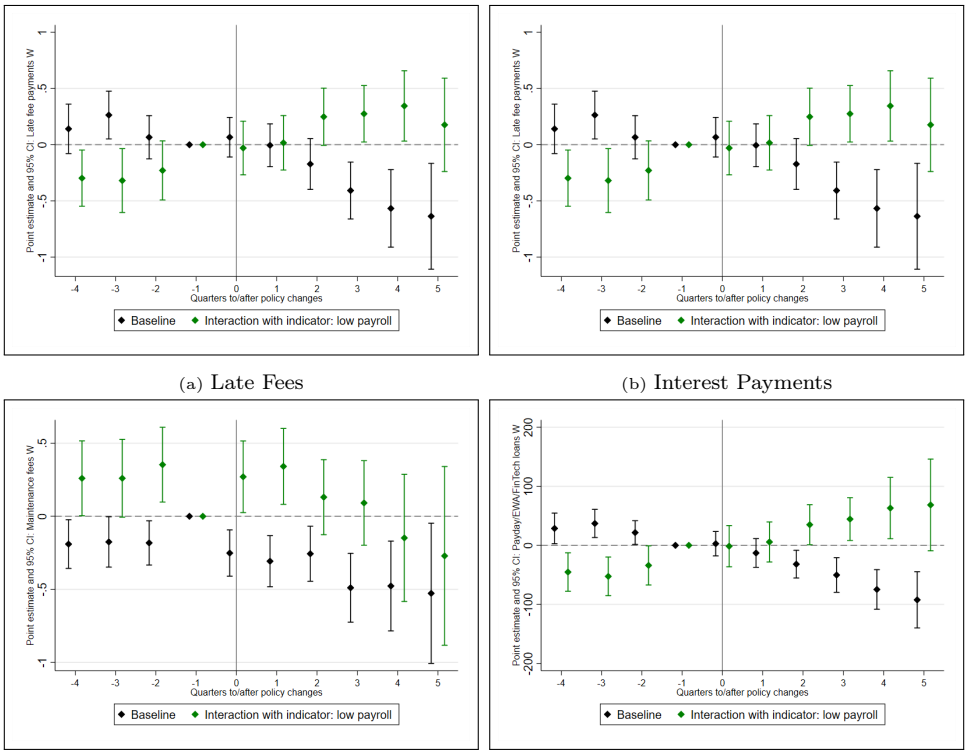
(b) Interest Payments

(c) Account Maintenance Fees

(d) High Interest Loans

Figure 7: Spillover Effects of Bank Policy Changes: By Income Across MSAs

Notes: This figure shows the average incremental spillover effects a change in bank fee policy (targeting overdraft and non-sufficient fund fees) has on late fees, interest payments, account maintenance fees and high interest loans, by income group. The results for **above median income** MSAs are in **black** and the results for **below median income** MSAs are in **green**. The effects are estimated using Equation 2, i.e., the stacked estimation design using all events and for each treatment MSA matching the distribution of the history of pre-treatment exposure to form the control group, controlling for MSA times stack and month-by-year times stack fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state-by-stack level.



(a) Late Fees

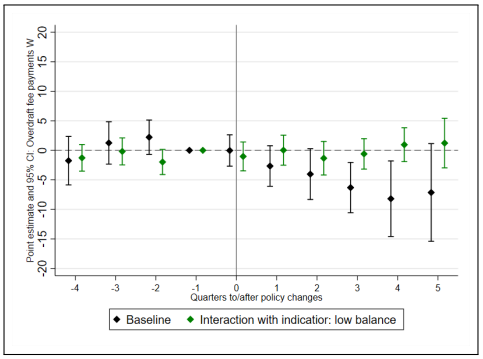
(b) Interest Payments

(c) Account Maintenance Fees

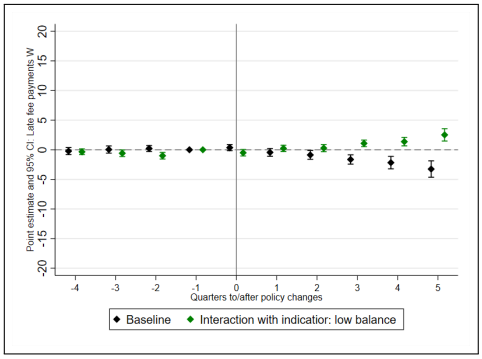
(d) High Interest Loans

Figure 8: Spillover Effects of Bank Policy Changes: By Income Within MSA

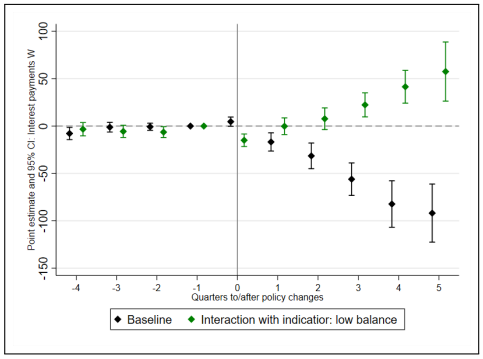
Notes: This figure shows the average incremental spillover effects a change in bank fee policy (targeting overdraft and non-sufficient fund fees) has on late fees, interest payments, account maintenance fees and high interest loans, by income group. Splitting within an MSA, the **above median income** households are in **black** and the results for **below median income** households are in **green**. The effects are estimated using Equation 2, i.e., the stacked estimation design using all events and for each treatment MSA matching the distribution of the history of pre-treatment exposure to form the control group, controlling for MSA times stack and month-by-year times stack fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state-by-stack level.



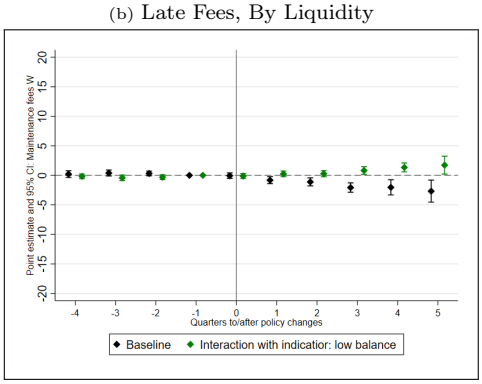
(a) OD Fees, By Liquidity



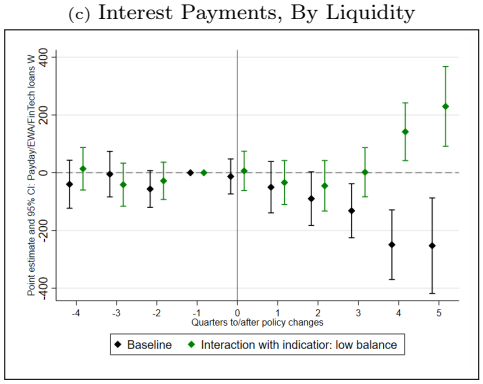
(b) Late Fees, By Liquidity



(c) Interest Payments, By Liquidity



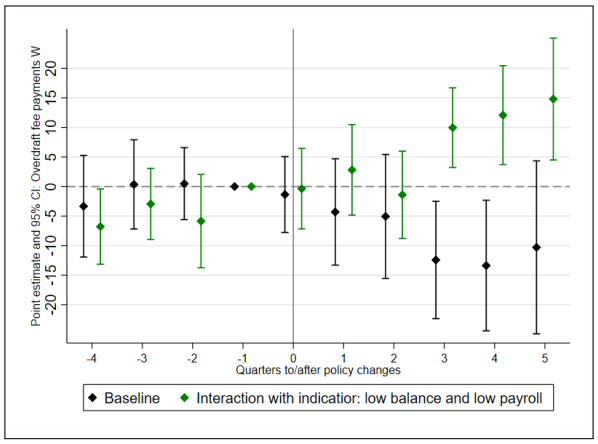
(d) Account Maintenance Fees, By Liquidity



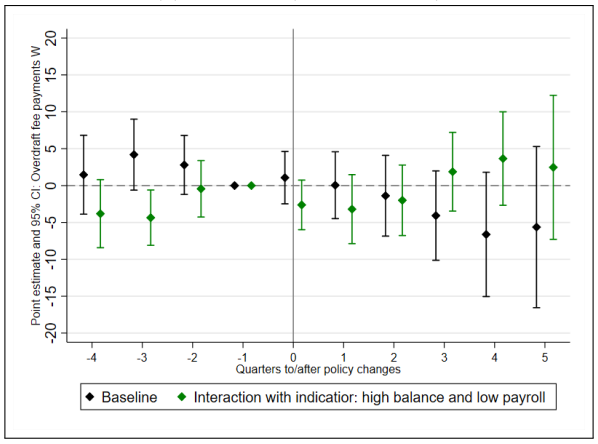
(e) High Interest Loans, By Liquidity

Figure 9: Effects of Bank Policy Changes: By Liquidity

Notes: This figure shows the average incremental effects a change in bank fee policy (targeting overdraft and non-sufficient fund fees) has on overdraft fees, late fees, interest payments, account maintenance fees and high interest loans, by liquidity group. The results for **above median liquidity** MSAs are in **black** and the results for **below median liquidity** MSAs are in **green**. The effects are estimated using Equation 2, i.e., the stacked estimation design using all events and for each treatment MSA matching the distribution of the history of pre-treatment exposure to form the control group, controlling for MSA times stack and month-by-year times stack fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state-by-stack level.



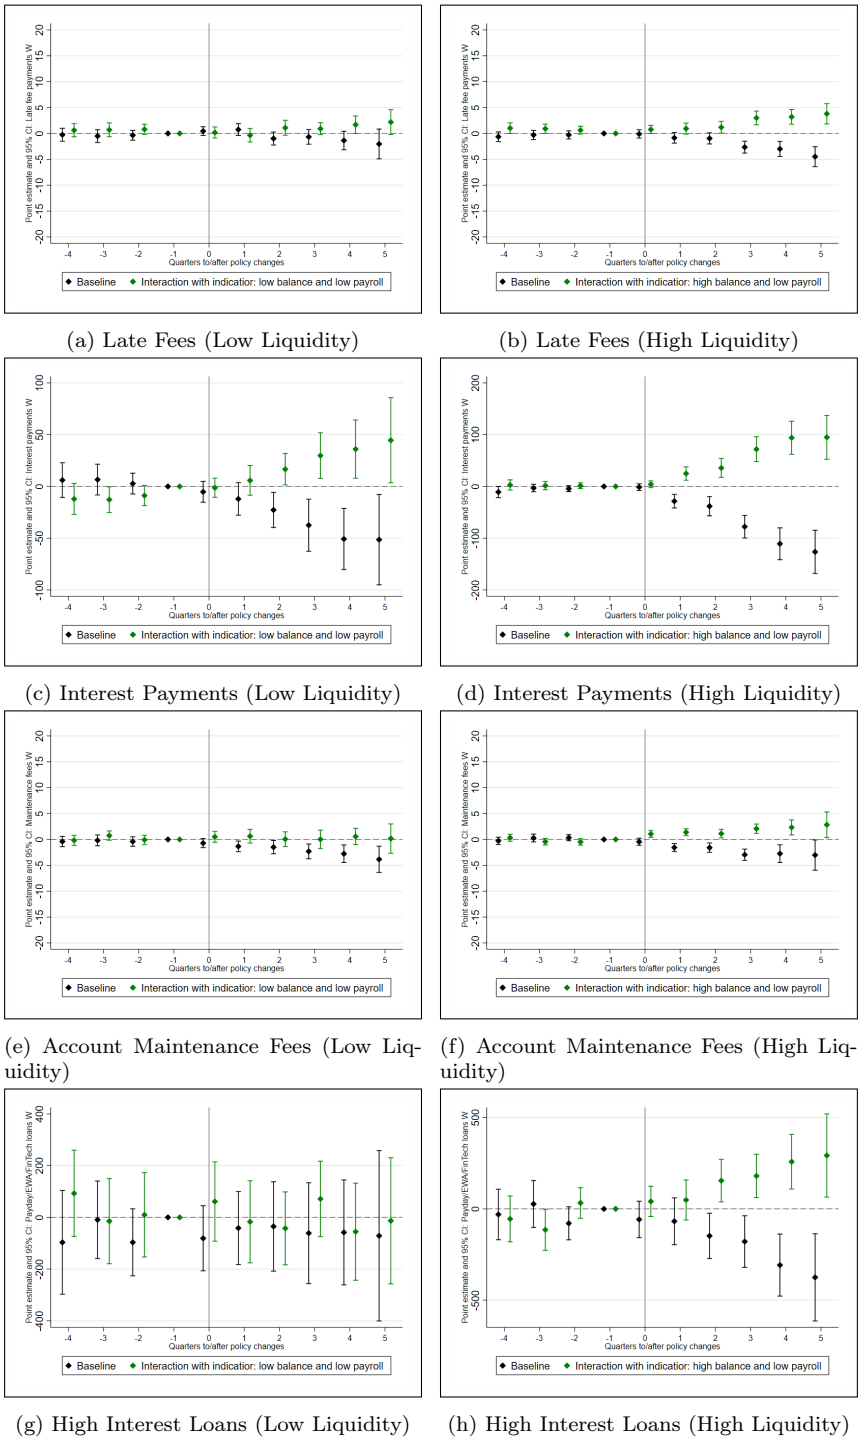
(a) OD Fees, (Low Liquidity)



(b) OD Fees, (High Liquidity)

Figure 10: Effects of Bank Policy Changes on Depositor Fees: By Liquidity and Income

Notes: This figure shows the average incremental effect a change in bank fee policy has on overdraftfees, by liquidity and income. The first column is for the low-liquidity sample, and the second column is for the high-liquidity sample. The results for **above median income** MSAs are in **black** and the results for **below median income** MSAs are in **green**. The effects are estimated using Equation 2, i.e., the stacked estimation design using all events and for each treatment MSA matching the distribution of the history of pre-treatment exposure to form the control group, controlling for MSA times stack and month-by-year times stack fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state-by-stack level.



(a) Late Fees (Low Liquidity)

(b) Late Fees (High Liquidity)

(c) Interest Payments (Low Liquidity)

(d) Interest Payments (High Liquidity)

(e) Account Maintenance Fees (Low Liquidity)

(f) Account Maintenance Fees (High Liquidity)

(g) High Interest Loans (Low Liquidity)

(h) High Interest Loans (High Liquidity)

Figure 11: Spillover Effects of Bank Policy Changes: By Liquidity and Income

Notes: This figure shows the average incremental spillover effects a change in bank fee policy (targeting overdraft and non-sufficient fund fees) has on late fees, interest payments, account maintenance fees and high interest loans, by liquidity and income. The first column is for the low-liquidity sample, and the second column is for the high-liquidity sample. The results for **above median income** MSAs are in **black** and the results for **below median income** MSAs are in **green**. The effects are estimated using Equation 2, i.e., the stacked estimation design using all events and for each treatment MSA matching the distribution of the history of pre-treatment exposure to form the control group, controlling for MSA times stack and month-by-year times stack fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state-by-stack level.

Tables

Bank	Date	Policy Change
Wells Fargo	Nov 2017	"Rewinds" overdraft fee if sufficient funds arrive by next morning
	Sep 2020	Introduced Clear Access Banking - account without overdraft fees; no service fee for ages 13-24
	Q1 2022	Eliminated NSF fees, transfer fees between linked accounts; introduced 24-hour overdraft grace period and small dollar loans
Chase	August 2021	No overdraft fees for amounts <$50; next-day grace period for sufficient funds
Bank of America	2010	Eliminated overdraft fees with debit/ATM fees
	Mar 2014	Introduced SafeBalance account - no overdraft fees, $5 monthly fee, no paper checks
	Sep 2021	Allowed linkage of up to 5 accounts to avoid overdraft fees
	Feb 2022	Eliminated NSF fees
	May 2022	Reduced overdraft fees from $35 to $10; eliminated transfer fees
TD Bank	Approx. Feb 2022	Eliminated NSF fees; no overdraft fee for amounts <$50; next-day grace period
US Bank	Jan 3, 2022	Eliminated NSF fees
	Jun 2022	Increased no-fee overdraft threshold from $5 to $50; 24-hour grace period
Regions Bank	Mar 17, 2022	Eliminated overdraft transfer fees
	Jun 16, 2022	Eliminated NSF fees; maximum three overdraft fees per day
PNC Bank	Apr 13, 2021	Introduced 24-hour grace period for overdraft fees
	Aug 2022	Eliminated NSF fees
Citibank	Summer 2022	Eliminated overdraft fees, returned item fees, NSF fees, and overdraft transfer fees
Capital One	Feb 2022	Eliminated overdraft and NSF fees
USAA Bank	Dec 8, 2022	Eliminated NSF and overdraft fees
Huntington	Jul 2022	Reduced overdraft fee from $36 to $15
Citizens Financial	Oct 7, 2021	Introduced 24-hour overdraft reversal
	Sep 26, 2022	Eliminated NSF fees
Woodforest Bank	May-Sep 2022	Eliminated NSF fees
Key Bank	Sep 16, 2022	No fee for overdrafts $20 or less; eliminated NSF fees; reduced overdraft fees and frequency
Fifth Third Bank	Jun 23, 2022	Eliminated NSF fees

Table 1: Timeline of Bank Overdraft and NSF Policy Changes (2010-2022)

Notes: To determine the time of the policy change, we cross-checked each policy change reported by the CFPB with official bank press releases and newspaper articles. The most granular dates are to the day and the most broad ones are within quarter. Corresponding to our transaction data date range, we are able to study 22 of 25 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020.

	Mean	Standard deviation	Median	5th percentile	25th percentile	75th percentile	95th percentile
Individual-level (no zeros):							
All retail spending	75.18	893.5	21.43	2.600	9.730	51.95	220.4
Online spending	126.6	1361.7	26.80	2.720	10.79	80	400
Restaurant spending	26.52	120.5	14.50	3.210	8.100	27.89	72
Grocery spending	38.47	153.7	19.46	1.610	7.440	42.98	129.1
Essential spending	71.70	1029.5	20	2.440	8.880	46	200
Income	1138.2	2668.6	487.0	0.01000	4.500	1549.8	3699.0
Relative Liquidity, Percentile	54.12	28.24	55	10	30	80	95
NSF fees	27.11	19.71	27.37	4.050	10	36	67.12
Overdraft fees	37.45	18.55	36	7	35	36	100
Interest payments	51.65	70.28	27.87	0.520	7.650	73.04	174.3
Late fees	28.90	7.765	27	15	25	35	39
Maintenance fees	8.506	56.51	3	1.500	3	6	24.95
Payday/EWA/FinTech loans	1123.6	3599.8	197.0	15.10	75	1181.1	3708.6
Transactions per user per quarter	1848.2	11228.1	372	99	241	548	1080
Observations (total)	87,525,521						
Aggregated to MSA-quarter level (with zeros):							
NSF fees	253.3	4507.6	0	0	0	0	72.91
Overdraft fees	655.0	11230.8	0	0	0	35	360
Interest payments	2809.8	49175.6	0	0	0	241.7	1343.0
Late fees	112.5	1993.9	0	0	0	0	64
Maintenance fees	185.5	3561.6	0	0	0	6	83
Payday/EWA/FinTech loans	9435.4	154768.3	0	0	0	0	2914.0
Average Income	13893.3	11802.0	12230.6	0.0700	7338.7	17687.8	32854.0
Share of Essential Spending	0.247	0.138	0.228	0.0604	0.164	0.305	0.494
Relative Liquidity, Percentile	54.24	4.891	54.49	48.65	51.58	56.34	59.62
Number of users per quarter	14571.4	8353.2	14580	1468	7365	21759	27546

Table 2: Summary Statistics

Notes: This table reports the distributions of various spending, income, and financial fee payments. The top panel statistics use the individual-level transactions data while the bottom panel sums up the financial fee payments to the MSA-quarter level. The categorizations of spending, income, as well as interest payments come from the account aggregator's classifications. The financial fee payments are obtained through manual searches in transaction descriptions. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020.

Supplemental Appendix

Robustness: Additional Estimator Methodologies

For robustness, we consider two additional methodologies, the first with with a never-treated MSAs control group, and the second where the treated and control groups are matched on the ex-ante distribution of cumulative treatment exposure.

Method 5: All Policy Changes as Treatment, Controls Matched on Propensity Score of Pre-Treatment Exposure We once again consider all policy changes. This time, the control group is constructed by matching propensity scores for cumulative ex-ante policy exposure, using the nearest x neighbors. This approach strives to avoid biases from early versus late comparisons by closely controlling for the distribution of prior exposure. Our pre-treatment condition is then,

$$\Omega_z(\Theta_i, \theta_{i_{-1}}) = \{i, n | \forall m \text{ s.t. } z \in A_i, \exists x \text{ number of } n \text{ s .t } z \notin A_n, \theta_{i_{-1}} + \textstyle\sum_{i_j \in A_i | m_j < z} \theta_{i_j} = \theta_{n_{-1}} + \sum_{n_j \in A_n | n_j < z} \theta_{n_j}\}.$$

Not all MSAs that treated at time z will end up in stack z. The reason is that it is possible that x nearest neighbors cannot be found in the control group. We implement this matching procedure using the `kmatch` command in Stata.

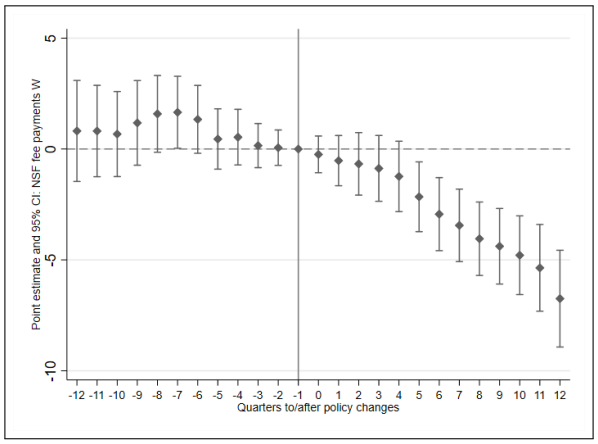
Method 6: Treatment of All Policy Changes, Control of Ex-Ante Equal Distribution of Policy Exposure Similar to Section 6, as we are matching the control and treatment groups on the coursened distribution of pre-treatment exposure levels. However, instead of considering the incremental treatment intensity distribution for the past d events, we just match on the distribution of the cumulative sum of all past treatments. Again, we minimize the differences of the control and treated distributions the mean, min 25%, 50%, 75%, max, for each past event.

Let f_z^T be the vector containing these attributes for any set of MSAs, with entry representing a distribution statistic. Analogously, let f_z^C be the matrix vector these attributes for any set of MSAs not treated at time z. Our pre-treatment condition is

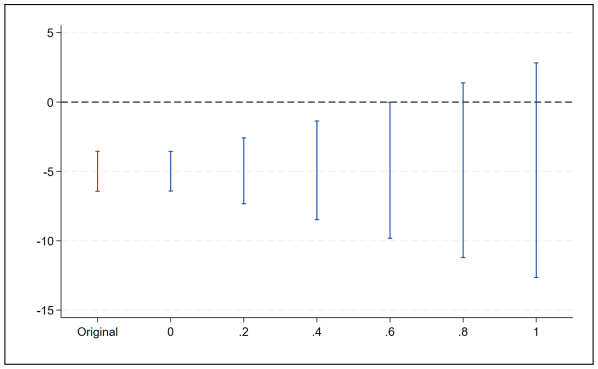
$$\Omega_z(\Theta_m, \theta_{m_{-1}}) = \{m, n | z \in A_m, z \notin A_n, \arg\min_{m,n} |f_z^T - f_z^C|\}.$$

Our simulation and empirical implementation reveal significant estimates, with minimal pre-trends.

Figures: Empirical Implementation of Estimator Methodologies



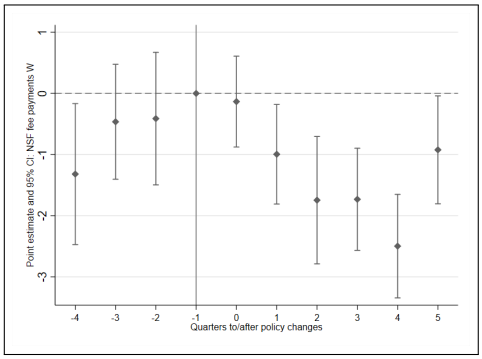
(a) Non-sufficient Fees ($)



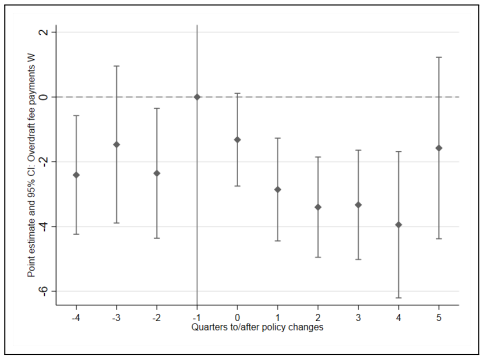
(b) Honest DID, Parallel Trends

Appendix Figure A1: Non-Sufficient Fund Fees Robustness Checks

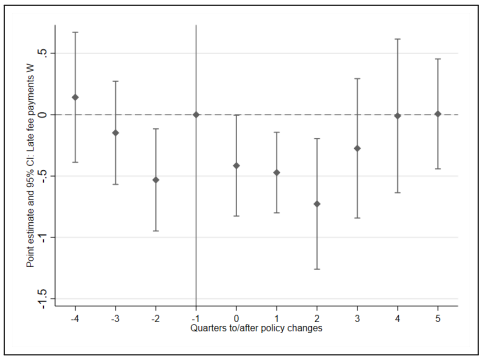
Notes: This figure provides robustness for the NSF finding, specifically considering the pre-trend. The top panel re-runs Equation 2, but with relative months instead of relative quarters, i.e., the stacked estimation design using all events and for each treatment MSA matching the distribution of the history of pre-treatment exposure to form the control group, controlling for MSA times stack and month-by-year times stack fixed effects. This more granular graph shows a lack of pre-trend, the effect of the policy change occurs within a quarter, just as in Panel (a) of Figure A6. The bottom panel utilizes the Honest DID framework of Rambachan and Roth (2023). The results remain significant long as the confounding factors which might be producing non-parallel trends post-treatment is at most 60% of the worst pre-treatment violation of parallel trends. In both panels, we study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals.



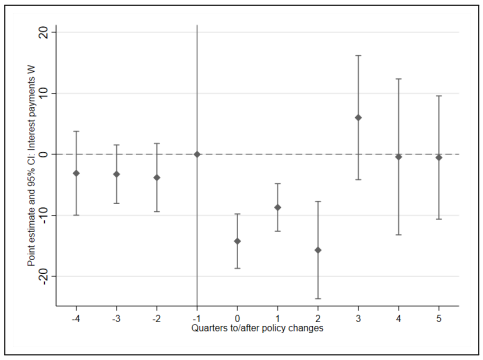
(a) Non-sufficient fund fees ($)



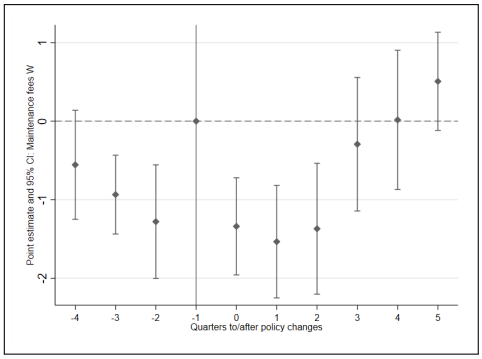
(b) Overdraft Fees ($)



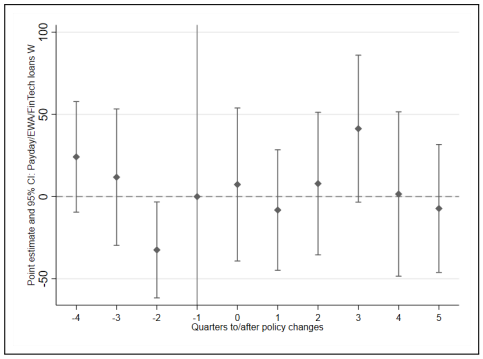
(c) Late Fees ($)



(d) Interest Payments ($)



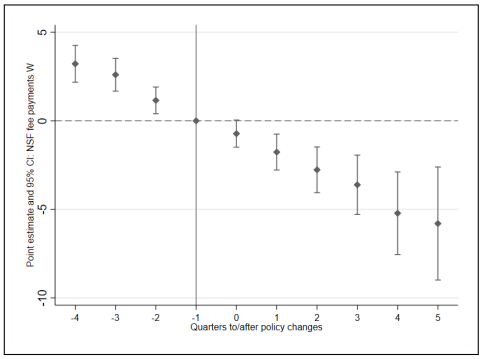
(e) Account Maintenance Fees ($)



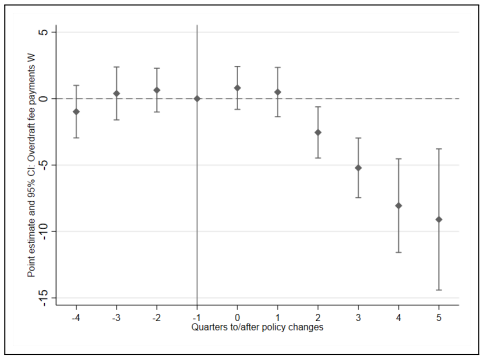
(f) High Interest Loans ($)

Appendix Figure A2: DiD Variation: Staggered

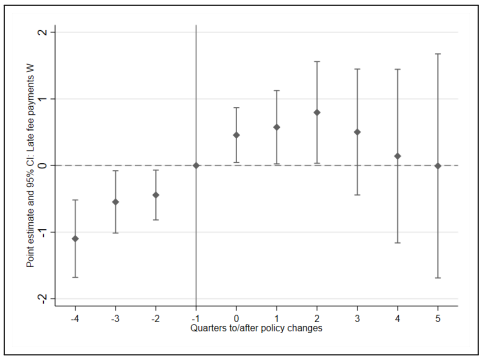
Notes: This figure shows the average incremental spillover effects a change in bank fee has policy overdraft, non-sufficient fund fees, late fees, interest charges, and any high interest loans. The effects are estimated using a staggered differences-in-differences model, controlling for MSA times stack and month-by-year time fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state level.



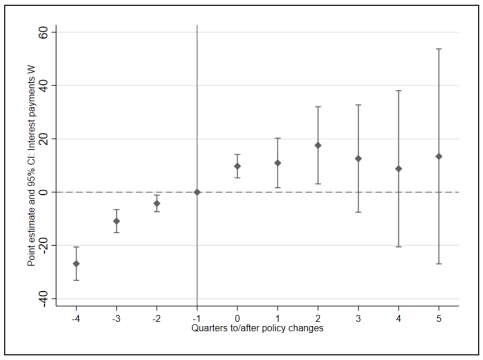
(a) Non-sufficient fund fees ($)



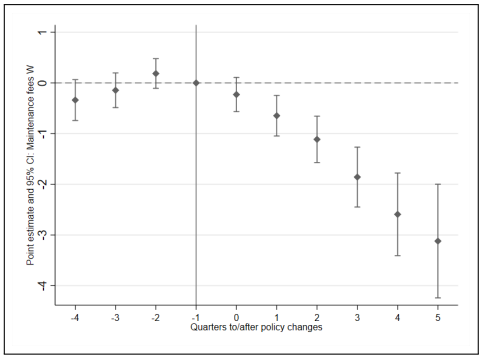
(b) Overdraft Fees ($)



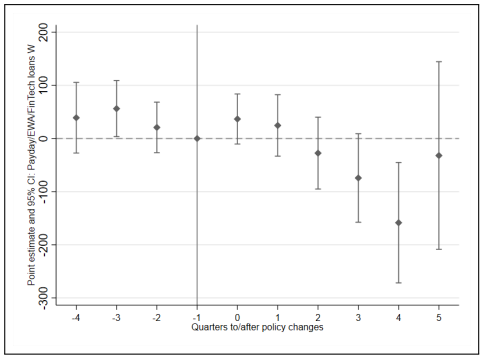
(c) Late Fees ($)



(d) Interest Payments ($)



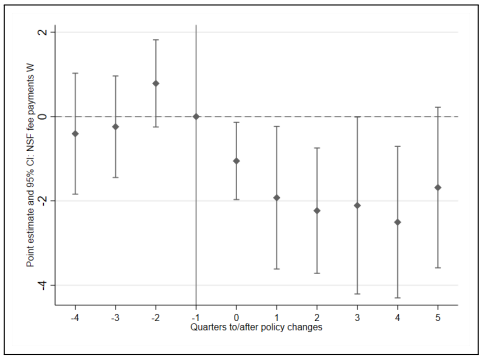
(e) Account Maintenance Fees ($)



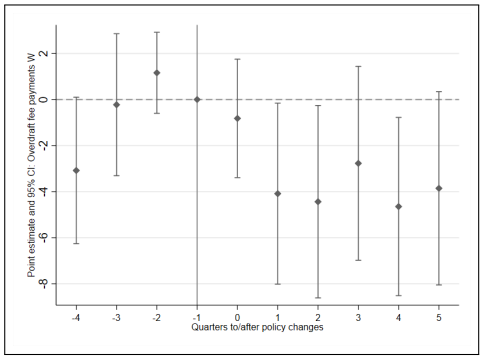
(f) High Interest Loans ($)

Appendix Figure A3: DiD Variation: Stacked with Never Treated Controls

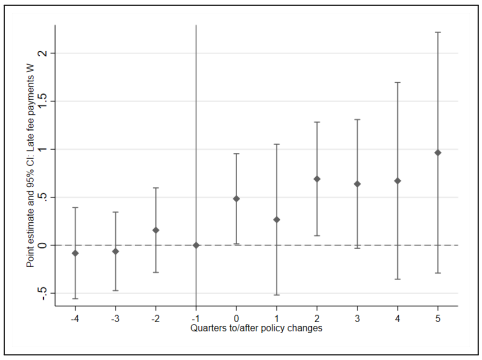
Notes: This figure shows the average incremental spillover effects a change in bank fee has policy overdraft, non-sufficient fund fees, late fees, interest charges, and any high interest loans. The effects are estimated using a stacked differences-in-differences model, using never-treated MSAs as our control. We also control for MSA times stack and month-by-year times stack fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state-by-stack level.



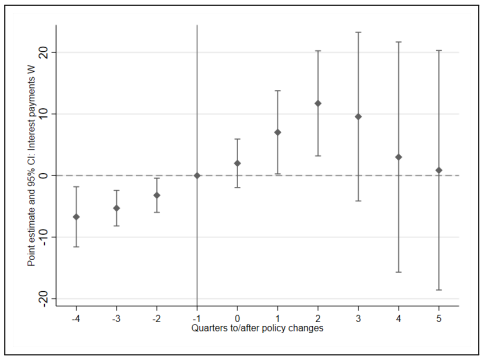
(a) Non-sufficient fund fees ($)



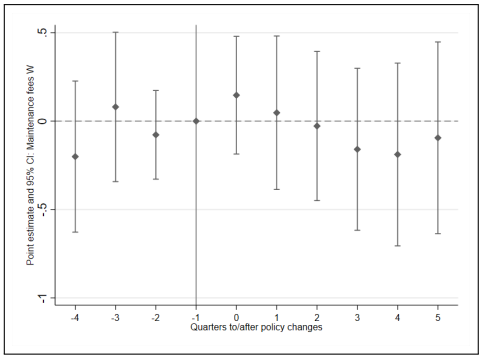
(b) Overdraft Fees ($)



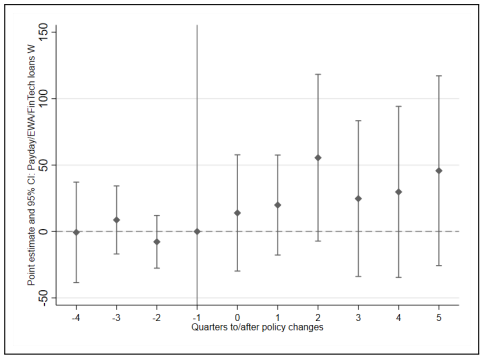
(c) Late Fees ($)



(d) Interest Payments ($)



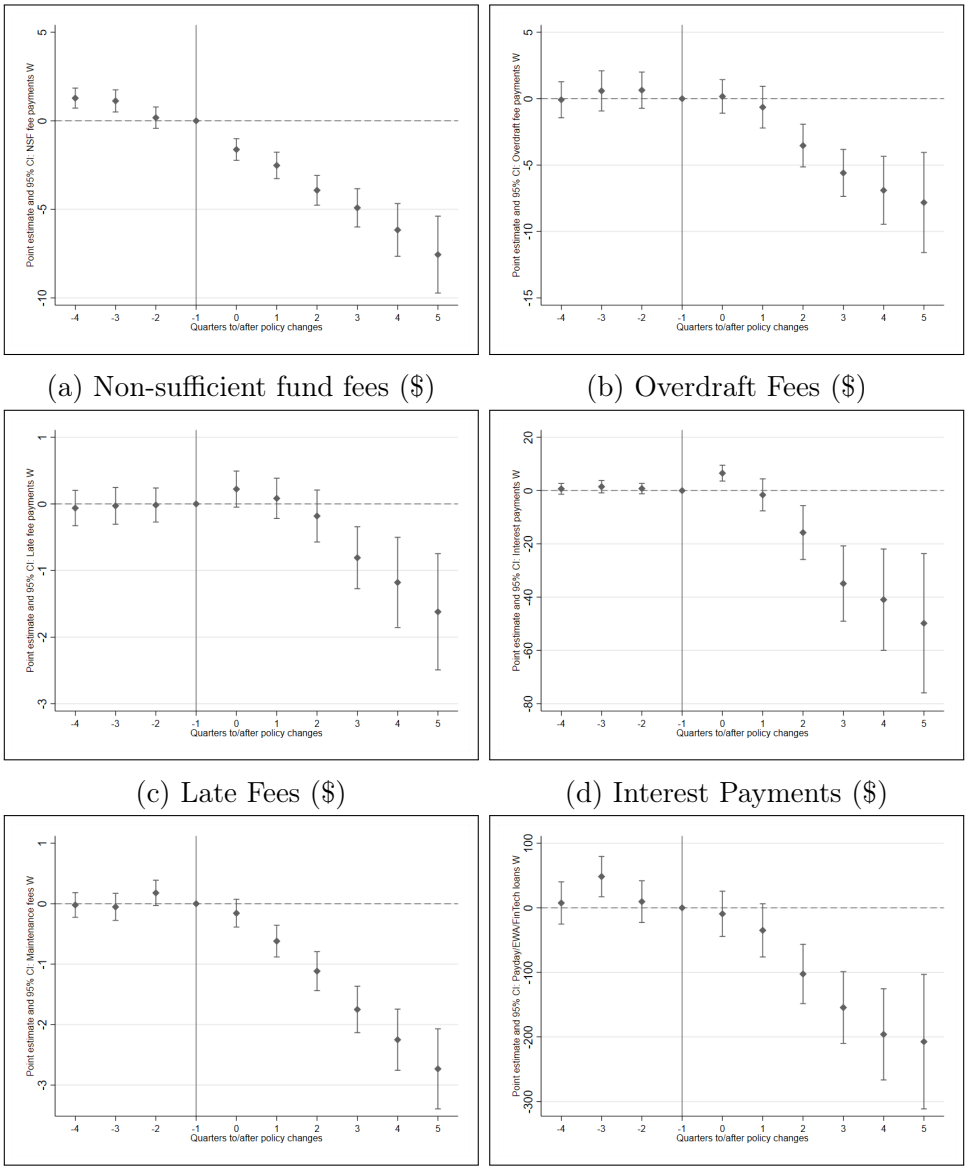
(e) Account Maintenance Fees ($)



(f) High Interest Loans ($)

Appendix Figure A4: DiD Variation: Stacked with Not-Yet-Treated Controls

Notes: This figure shows the average incremental spillover effects a change in bank fee has policy overdraft, non-sufficient fund fees, late fees, interest charges, and any high interest loans. The effects are estimated using a stacked differences-in-differences model, using not-yet-treated MSAs as our control. We also control for MSA times stack and month-by-year times stack fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state-by-stack level.



(a) Non-sufficient fund fees ($)

(b) Overdraft Fees ($)

(c) Late Fees ($)

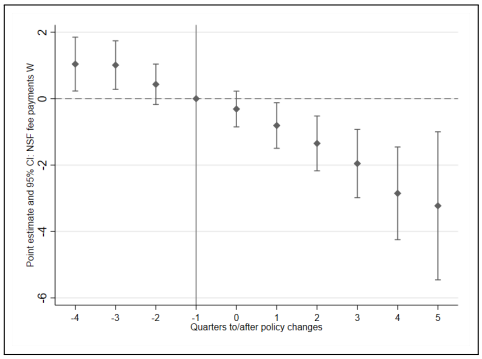
(d) Interest Payments ($)

(e) Account Maintenance Fees ($)

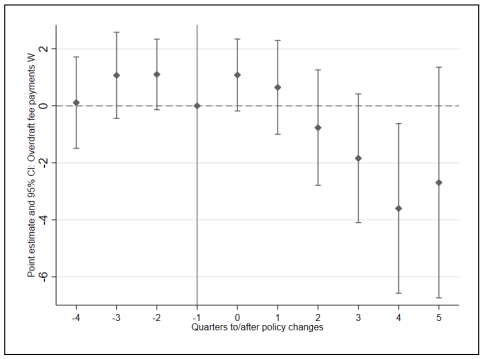
(f) High Interest Loans ($)

Appendix Figure A5: DiD Variation: Stacked with Mean-Match Controls

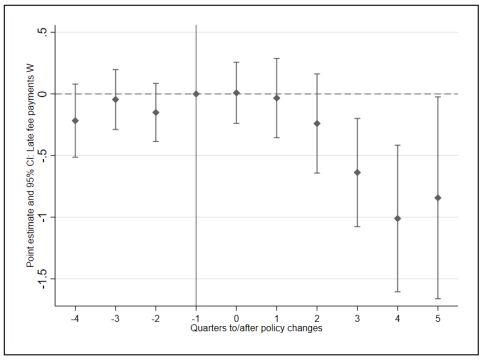
Notes: This figure shows the average incremental spillover effects a change in bank fee has policy overdraft, non-sufficient fund fees, late fees, interest charges, and any high interest loans. The effects are estimated using a stacked differences-in-differences model, with control group MSAs such that each has a cumulative pre-treatment exposure equal to the average cumulative pre-treatment exposure of the treated group. We also control for MSA times stack and month-by-year times stack fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state-by-stack level.



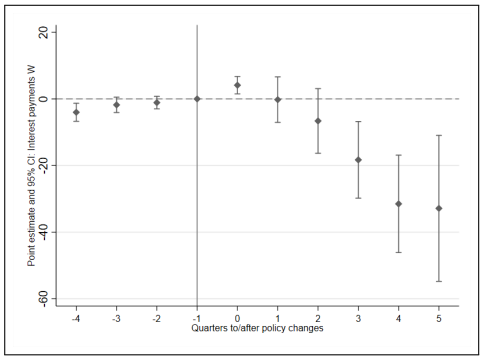
(a) Non-sufficient fund fees ($)



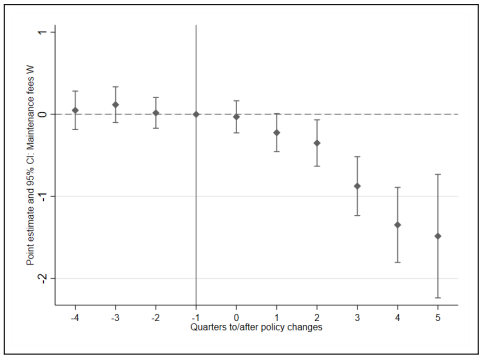
(b) Overdraft Fees ($)



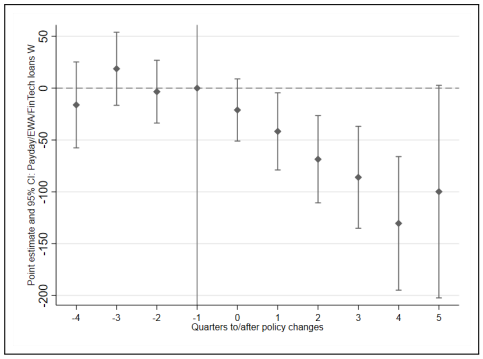
(c) Late Fees ($)



(d) Interest Payments ($)



(e) Account Maintenance Fees ($)



(f) High Interest Loans ($)

Appendix Figure A6: DiD Variation: Stacked with CEM Controls

Notes: This figure shows the average incremental spillover effects a change in bank fee has policy overdraft, non-sufficient fund fees, late fees, interest charges, and any high interest loans. The effects are estimated using a stacked differences-in-differences model. We construct a control sample that matches the coarsened distribution of incremental pre-treatment exposure for the past 1 event. We also control for MSA times stack and month-by-year times stack fixed effects. We study 22 bank policy changes made between 2018 and 2023 with the first one occurring in September 2020. The bars represent 95% confidence intervals which the standard errors being clustered at the state-by-stack level.