WORKING PAPER SERIES
NO. 2026-18 NOVEMBER 2026

Federal Reserve Bank *of* Atlanta

Free School Meals, Credit Access, and Household Debt

Orgul Ozturk, University of South Carolina
M. Melinda Pitts, Federal Reserve Bank of Atlanta

Summary:
Under the Community Eligibility Provision (CEP), schools serve free meals to every student without an application or a household means test, thus removing a recurring expense for many households. We examine the effect of CEP on the financial health of adults in those households, combining the staggered rollout of CEP with quarterly credit records from the Federal Reserve Bank of New York/Equifax Consumer Credit Panel. Our results suggest that CEP draws households further into credit markets. Access to credit steps up when CEP arrives and stays higher, most durably in revolving credit, the kind that lets a household choose when to spend. The Equifax Risk Score barely moves. Severe delinquency rises modestly, but not among people observed continuously in the credit records before and after onset. The gains are largest where the fewest students were already eligible for free or reduced-price meals before CEP. Existing evidence shows that borrowing rises when nutrition assistance is withdrawn; our results suggest that it also rises when assistance is extended.

JEL classification: G51, I38, H53, I28

Key words: household finance, school lunches, consumer credit

https://doi.org/10.29338/wp2026-18

Federal Reserve Bank of Atlanta Policy Hub papers are available on the Atlanta Fed's website at www.atlantafed.org. Click "Publications" and then "Policy Hub." To receive e-mail notifications about new papers, use atlantafed.org/subscribe.

Orgul Ozturk (odozturk@moore.sc.edu) is with the Darla Moore School of Business of the University of South Carolina, 1014 Greene Street, Columbia, SC 29208.

M. Melinda Pitts (melinda.pitts@atl.frb.org) is research center director of the Center for Human Capital Studies in the Research Department of the Federal Reserve Bank of Atlanta.

The authors thank Benjamin Miller for excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Atlanta or the Federal Reserve System.



Federal Reserve Bank *of* Atlanta

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed's website at www.atlantafed.org. Click "Publications" and then "Working Papers." To receive e-mail notifications about new papers, use atlantafed.org/subscribe.

1. Introduction

Many households in the United States are financially fragile, with little ability to absorb an unexpected expense or a loss of income. In the Federal Reserve's Survey of Household Economics and Decision Making for 2025, 37 percent of adults would not have covered an unexpected $400 expense entirely with cash or its equivalent, and 12 percent could not have paid it at all (Board of Governors of the Federal Reserve System, 2026). Household debt stands at $18.8 trillion, about 3 percent below its 2008 peak once prices are taken into account, and delinquency, after falling during the pandemic, has returned to its pre-pandemic level, with 4.7 percent of outstanding debt delinquent in mid-2026, the same share as at the end of 2019. Our sample covers the recovery from the financial crisis. Household debt fell from $12.1 trillion at the start of 2010 to $11.2 trillion in mid-2013 and rose to $13.9 trillion by mid-2019, while the share of debt in some stage of delinquency fell from 11.8 percent to 4.4 percent.[1] Government programs provide substantial assistance that may alleviate some financial concerns, but there is limited empirical evidence on how in-kind assistance programs affect broader measures of household financial well-being. In this paper, we examine how an in-kind nutritional assistance program, specifically free school meals, affects household financial health. We contribute to the literature on how government assistance programs, such as the Earned Income Tax Credit (EITC) (Goodman-Bacon and McGranahan, 2008), the Supplemental Nutrition Assistance Program (SNAP) (Hoynes and Schanzenbach, 2009), or the safety net as a whole, including unemployment insurance (Bitler and Hoynes, 2016), among others, affect household well-being.

[1] Federal Reserve Bank of New York, Quarterly Report on Household Debt and Credit, 2019:Q4 and 2026:Q2. Real values use the CPI-U, deflating the 2008 third-quarter peak from September 2008 to June 2026.

That literature has focused largely on poverty and income, and we study a different outcome, households' use of credit and the financial health it reflects.

Specifically, the focus of this research is the National School Lunch Program, which is the largest form of nutritional assistance for school-aged children. This program expanded dramatically beginning in the fall of 2011 with the Community Eligibility Provision (CEP). Traditionally, the free or reduced-price school lunch was means-tested at the student level and required individual families to submit applications. Under CEP, the eligibility is instead determined at the school or district level and is based on the share of students directly certified through other assistance programs. As a result, CEP expands access to in-kind nutritional resources not only for low-income families but also for households who may not qualify for other forms of income assistance, such as SNAP, and for those who may rely on aid from private or non-profit organizations.

Although school meals are provided directly to children during the school day, the economic unit affected by the program is the household. By lowering the effective price of food, CEP may free up resources within the household budget. These freed resources can be reallocated toward non-food consumption, debt repayment or savings. They may also change the timing of spending, making it easier to spread purchases across the month, to carry costs through the summer months when school meals stop, or to buy in bulk, and to take on and repay the credit that makes such purchases possible. Thus, even though the program targets children's nutrition, it is reasonable to expect spillover effects on adult financial outcomes.

This household-budget mechanism is consistent with a broader literature showing that in-kind public programs can affect financial well-being. Evidence from the Oregon Health Insurance Experiment shows that Medicaid coverage reduces out-of-pocket medical

expenditures, medical debt, and financial strain (Finkelstein et al., 2012). Studies of the Affordable Care Act Medicaid expansions similarly find reductions in medical collections, delinquency, unpaid bills, and adverse public records (Hu et al., 2018; Brevoort et al., 2020; Miller et al., 2021). Related work on housing mobility assistance finds effects on credit-market outcomes through changes in neighborhood conditions (Miller and Soo, 2021). Recent evidence from SNAP further shows that benefit access can affect household borrowing and credit outcomes. Dodini, Larrimore, and Tranfaglia (2024) find that the imposition of SNAP work requirements increases credit seeking, new account openings, balances, and past-due balances, while Homonoff, Lee, and Meckel (2025) show that households denied SNAP because of administrative barriers experience increased debt and delinquency and lower credit scores. Together, this evidence suggests that public programs can affect household finances not only through direct changes in income, but also by altering the resources available for borrowing, repayment, and other financial decisions.

Consistent with this mechanism, CEP adoption significantly reduces household grocery purchases (Handbury and Moshary, 2021) and demand for community food resources (Ruffini, Öztürk, and Pekgün, 2025).[2] These spillovers suggest that school meals function as a meaningful in-kind transfer at the household level. Yet little is known about how this transfer affects broader measures of financial stability and, particularly, access to credit. By freeing resources otherwise devoted to necessities, even relatively small in-kind transfers may alter household budget constraints in ways that affect both credit use and access. Importantly, greater access to credit

[2]These spillover effects are in addition to more direct effects of school meals, including improved nutrition and health (Bhattacharya et al., 2006; Gleason and Suitor, 2001; Gundersen et al., 2012), reduced food insecurity (Fletcher and Frisvold, 2017; Gundersen et al., 2012), increased meal consumption (Schwartz and Rothbart, 2020; Ruffini, 2022), and improved academic behavior and outcomes (Frisvold, 2015; Fuller and Comperatore, 2020; Gordanier et al., 2020; Gordon and Ruffini, 2021; Kho and Hunter, 2026; Ruffini, 2022).

need not translate into higher credit scores or lower debt; it may instead expand households' capacity to finance other consumption or investment.

This distinction provides an important contrast with the recent SNAP evidence. Existing studies by Dodini, Larrimore, and Tranfaglia (2024) and Homonoff, Lee, and Meckel (2025) primarily identify the consequences of losing or being denied benefits, where increased borrowing may reflect households turning to credit to replace lost resources. Much less is known about the other side of this relationship, namely whether gaining access to an in-kind transfer can relax household budget constraints sufficiently to expand credit-market participation and/or alter borrowing behavior. CEP provides a useful setting for examining this margin because it provides a relatively small, recurring in-kind transfer that reduces an essential household expenditure without directly increasing cash income.

To study these potential effects of CEP, we combine information on the timing and expansion of the program with detailed measures of individual financial health from the Federal Reserve Bank of New York/Equifax Consumer Credit Panel (CCP). The CCP provides quarterly credit bureau data on debt balances, delinquencies, loan composition, credit access, and the Equifax Risk Score. We focus on adults aged 25–54, the group most likely to have school-aged children. Using a dose-response difference-in-differences framework, we estimate the impact of potential CEP exposure on credit-market participation, debt and delinquency across auto, mortgage, and revolving credit, as well as bankruptcy and foreclosure.

We find that exposure to schoolwide free meals draws households further into credit markets. The probability of holding any credit account rises by about one percentage point, and the number of accounts per person rises by about 2 percent. Access steps up when CEP arrives, including among people observed continuously before and after onset, and the number of

accounts builds over the following year and then levels off. The gains are most durable in revolving credit, the kind households can draw on and repay as needs change, and revolving access stays higher, among them as well. The lasting gain in the number of accounts does not appear among people observed continuously before and after onset, and may come from people new to the credit panel. Because four in five people already hold an account, the probability of holding any account moves by less, and mainly at onset. Severe delinquency rises modestly, but not among people observed continuously before and after onset, and the rise may come from people new to the credit panel. Because eligibility for free and reduced-price meals is set by family income, at 130 and 185 percent of the poverty line, the share of a school's students already eligible measures how many families qualified on income before CEP. Our results suggest the gains are largest where that share was lowest, among families whose incomes sat just above the cutoffs and who paid for school meals before CEP.

The credit records show the credit households hold, not whether the change came from households seeking more credit or lenders offering more. Because the Equifax Risk Score does not improve, the expansion does not appear to reflect lenders responding to households becoming better risks. The Equifax Risk Score falls by about 0.9 score points, a seventh of one percent of its mean, consistent with the recent opening of new accounts, one of the factors that lower credit scores. A loosening of lending standards would not have to move the Equifax Risk Score, but it could explain our results only if it fell on CEP tracts when they adopted, since the state-by-quarter effects absorb anything common to a state. Our results therefore suggest that the expansion reflects households seeking more credit, although the credit records cannot confirm it.

By documenting spillover effects of a non-means-tested, in-kind assistance program on credit-market participation, this paper adds a new dimension to the literature on the financial

consequences of the safety net. Existing work shows that losing or being denied benefits can increase borrowing and financial distress. We show that expanding in-kind assistance can also increase credit-market participation and debt, while barely moving the Equifax Risk Score, a measure of creditworthiness. These findings highlight an understudied margin of household response to public transfers. Assistance may alter not only financial distress and consumption smoothing, but also households' access to credit and the set of financial choices available to them.

2. Conceptual Framework and Institutional Setting

2.1 *Conceptual Framework*

To illustrate the potential mechanisms linking schoolwide free meals to financial outcomes, consider a simple household model in which families maximize utility over food and non-food consumption. Households obtain food through private grocery purchases and government-provided meals. Although government assistance is monetarily free, recipients may incur non-monetary costs such as stigma or administrative burdens (Byrne, Just, and Barrett, 2023). Adoption of a schoolwide free meal program increases government-provided food among students who were previously not certified to receive free meals and lowers the effective price of food for these households. Moreover, the universal nature of CEP may reduce stigma among previously eligible students.

Before the introduction of CEP, households above the free-meal threshold (130 percent of the poverty line) that were paying reduced or full price for school meals faced a recurring food cost for every child. In the 2014–15 school year, among schools that charged for meals, a full-price lunch averaged $2.42 and a full-price breakfast $1.43, so a child paying full price for both every school day cost a household about $690 over a 180-day school year, about $440 of it for

lunch (Fox and Gearan, 2019).[3] The CEP program eliminates this cost. With increased receipt of government-provided nutrition, households may substitute away from grocery purchases if food sources are imperfect substitutes with positive cross-price elasticities. Marcus and Yewell (2022) find an average reduction in household grocery spending between 5 and 19 percent, depending on the zip code level of exposure to CEP. The reduction in grocery spending effectively relaxes the household budget constraint; if both food and non-food goods are normal goods, the in-kind transfer expands the effective budget set. Households may respond by (1) increasing non-food consumption, (2) reducing borrowing, (3) paying down existing debt, or (4) taking on new credit to finance consumption or durable goods that were previously difficult to afford. If households were previously borrowing to smooth consumption, particularly for food expenditures, greater government-provided nutrition may reduce reliance on revolving credit and lower delinquent balances. Alternatively, by freeing resources otherwise committed to recurring food expenditures, CEP may increase households' capacity to take on and service new debt. Thus, a relaxed budget constraint can generate either less reliance on credit or greater access to and use of credit, depending on how households reallocate the resources released by the transfer.

This mechanism is consistent with a broader literature on public assistance and household consumption smoothing. Gundersen and Ziliak (2003) find that Food Stamps reduce income and food-consumption volatility among households with a high probability of program participation. More recent evidence indicates that the timing of SNAP disbursement affects the smoothness of household expenditures, underscoring the importance of predictable resources for financially constrained households (Cotti, Gordanier, and Ozturk, 2021). Although the form of assistance

[3] Few paying students ate school breakfast before CEP, 5 percent against the 35 percent who ate school lunch, but a child who eats a free breakfast at school also spares the household the cost of breakfast at home, so the saving does not depend on the family having bought school breakfast before.

differs across programs, the underlying mechanism is similar. When a public program covers an essential expenditure, it relaxes the household budget constraint and may alter borrowing, repayment, and other financial decisions. Evidence from SNAP further shows that these adjustments can occur through credit markets, as households increase credit seeking and borrowing when benefits are lost or restricted (Dodini, Larrimore, and Tranfaglia, 2024; Homonoff, Lee, and Meckel, 2025). Our setting examines the other side of this relationship, whether an expansion of in-kind resources changes not only financial distress but also households' access to and participation in credit markets.

How households use such resources also depends on how they arrive. Stimulus delivered as a flow of reduced withholding raised spending at about half the rate of a one-time payment (Sahm, Shapiro, and Slemrod, 2012), and many households used the 2011 payroll tax cut, also a flow, to smooth the repayment of debt rather than to spend (Sahm, Shapiro, and Slemrod, 2015). CEP likewise delivers its resources as a flow, a cost removed every school day. This logic extends beyond nutrition programs to housing, utility, health, and child-care assistance, each of which can reduce the resources households must devote to recurring necessities. CEP provides a useful setting in which to examine this mechanism because it expands in-kind support to households that may not receive other means-tested benefits.

The impact of CEP on household finances is expected to vary with how many students were already eligible for free or reduced-price meals before CEP. Where nearly every student was already certified, through SNAP or Medicaid or by application, CEP changes little for most families, and the marginal cash-flow benefit is small. Where few students were eligible, CEP makes meals free for many families who had been paying the full or reduced price, or providing their own lunches, and the gain in the free-meal share, and the dollars it releases, is greatest. The

cash-flow channel therefore predicts that the financial response is largest where the share of students already eligible was lowest.

Broader equilibrium effects may also occur. Because CEP is not means-tested at the household level and does not impose a direct income cliff, it may alter labor-supply and employment decisions. The resources released by free school meals may also facilitate access to transportation. Reliable automobile access can expand the set of jobs that workers can feasibly reach and retain, particularly where employment is spatially dispersed or work schedules are irregular. Prior work finds that automobile access is associated with improved employment outcomes among welfare recipients and other low-income workers (Raphael and Rice, 2002; Ong, 2002; Gurley and Bruce, 2005). At the same time, automobiles require financing, insurance, maintenance, and fuel, so greater transportation access may increase demand for auto credit and may not unambiguously improve measured financial health. Thus, an increase in debt following CEP need not by itself indicate deteriorating financial well-being; it may also reflect expanded access to financed durable goods among previously constrained households. CEP could therefore affect both the extensive margin of credit access and the risk associated with additional borrowing. The predicted effects on financial outcomes are therefore mixed. CEP may simultaneously reduce financial distress while increasing participation in certain credit markets. At the same time, actual behavior may deviate from this simple model. Benefits may not be fully salient to households, take-up may be incomplete since meals are take-it-or-leave-it, or assistance from family and friends may be crowded out. Therefore, empirical analysis is necessary to determine how households adjust their financial behavior in response to CEP.

2.2 *Policy Reform: School Meals and the Community Eligibility Provision (CEP)*

The school meals program is one of the oldest and largest nutritional assistance programs for families with school-aged children. In the past decade, the nature of the school meals program fundamentally changed due to reforms that allowed schools to offer free meals to all students, regardless of family income. The largest of these reforms, the Community Eligibility Provision (CEP), began as a pilot in the 2011–12 school year and expanded nationwide in 2014–2015.[4]

Under the traditional National School Lunch and School Breakfast Programs, students qualified for free or reduced-price meals through household-level applications or direct certification. Under CEP, eligibility is determined at the school, group-of-schools, or district level. A school, group, or district may participate when at least 40 percent of enrolled students are directly certified through SNAP, TANF, the Food Distribution Program on Indian Reservations, or another categorical eligibility pathway. Participating schools receive federal reimbursement for meals at a rate equal to 1.6 times their directly certified share, capped at 100 percent. Thus, when the directly certified share is at least 62.5 percent, all meals are reimbursed at the federal free-meal rate; below that threshold, the school or district bears the cost of meals not fully reimbursed. Districts may also pool students across multiple schools when determining eligibility and reimbursement. CEP adoption is therefore more attractive where the directly certified share is high, but it remains a local decision shaped by expected reimbursement, administrative costs, and district resources.

CEP changes the incidence of school-meal assistance in two ways. First, it eliminates the application requirement for students in participating schools, reducing administrative burdens

[4] CEP became available to eligible schools in Illinois, Kentucky, and Michigan in the 2011–12 school year; the District of Columbia, New York, Ohio, and West Virginia in 2012–13; Florida, Georgia, Maryland, and Massachusetts in 2013–14; and in all remaining states in the 2014–15 school year.

and potentially stigma associated with individual certification. Second, it extends free meals to students who would otherwise have paid full or reduced prices, including children in households that are not eligible for SNAP or other cash-equivalent assistance. The financial implications of CEP should therefore be concentrated in locations where participation produces a meaningful increase in the share of students receiving free meals, rather than in locations where most students were already directly certified before adoption.

The stated goals of CEP were to expand access to school meals and simplify program administration. Consistent with these goals, participation among eligible schools increased over time. By the 2019–20 school year, 14.9 million children attended a school offering free meals to all students through CEP (Food Research & Action Center, 2020), close to 30 percent of the 50.8 million students enrolled in public schools (National Center for Education Statistics, 2023, Table 203.10), and adoption varied widely across states and was more common in higher-poverty schools (Food Research & Action Center, 2020). Because participation is voluntary, adoption may also reflect local conditions associated with financial need and district capacity. Our empirical approach therefore uses variation in the timing and intensity of CEP exposure while accounting for persistent differences across individuals and locations. The institutional features of CEP imply that its financial effects depend both on whether schools adopt the program and on the extent to which adoption increases the value of free meals for households in the surrounding area.

3. Data

3.1 Access to Free School Meals

To measure local exposure to schoolwide free meals, we combine school-level information on participation in the Community Eligibility Provision (CEP) with geographic information linking

schools to residential areas. The starting point for our CEP exposure measure is a panel dataset of public schools containing information on CEP participation and school characteristics for academic years 2009–2010 through 2018–2019. We use this data to measure, for each tract and quarter, the share of the tract's school-age children likely to attend a school covered by CEP, which Section 3.4 splits into coverage at onset and coverage gained afterwards.[5]

The school-level dataset includes indicators for overall and CEP participation, school grade level (elementary, middle, and high school), and baseline measures of student disadvantage such as the share of students eligible for free or reduced-price lunch in 2010. Charter schools are excluded from the construction of CEP exposure because they frequently operate outside traditional attendance zones and therefore may not reflect the schools most likely attended by students residing in a given neighborhood. Furthermore, New York City is excluded from the analysis for similar reasons.

Because CEP participation is determined at the school (or school district) level while our empirical analysis is conducted at the residential level, we construct a measure of CEP exposure by linking schools to nearby Census tracts. We begin with a dataset containing the geographic centroids of all U.S. Census tracts. For each tract and year, we identify the nearest elementary, middle, and high school using the geographic distance between the tract centroid and school coordinates. Distances are Euclidean. To ensure that tract-school matches reflect plausible enrollment patterns, matches are restricted to schools located within the same state and within a fixed geographic radius of the tract centroid.

[5] These data extend and update the school-level CEP participation measures used in Ruffini (2022) and Gordon and Ruffini (2021).

In addition to geographic proximity, we further restrict tract-school matches using school district boundaries. Specifically, we merge the school dataset with a crosswalk linking school districts to counties. A tract is matched only to schools operated by districts that serve the tract's county. This step helps ensure that the assigned schools are part of the local public school system and prevents tracts from being linked to nearby schools that do not serve residents of the county. Once the nearest elementary, middle, and high schools are identified for each tract-year, the school-level CEP data are merged to the tract-level dataset. This allows us to determine whether the schools most likely attended by students in a tract offer schoolwide free meals in a given year.

Because the analysis ultimately aggregates CEP exposure across residential areas, we combine tract-level school matches with population data describing the number of school-aged children living in each tract. Population counts for ages 5–9, 10–14, and 15–19 are obtained from the National Cancer Institute's SEER population database and correspond to elementary, middle, and high school age groups, respectively. These population counts are used as weights when aggregating CEP exposure measures. Using these weights, we construct tract-level measures of CEP exposure that reflect the population-weighted share of nearby schools participating in CEP. Because CEP adoption expanded gradually over time and varied across locations, this approach produces meaningful geographic and temporal variation in access to free school meals. This exposure variable is the basis for the CEP measures used in the analysis. Districts must elect CEP by June 30 for the following school year (7 C.F.R. 245.9(f)(4)(i)), so coverage changes, with few exceptions, only at the start of a school year, and a tract's onset is the first quarter of the school year in which its coverage turns positive.[6]

[6] USDA extended the election deadline to August 31 for 2014–15 and 2015–16, which still falls in our onset quarter, and allowed some districts to elect CEP during the school year (U.S. Department of Agriculture, Food and Nutrition

3.2 Measure of Financial Well-Being

To create a comprehensive measure of household financial well-being we use the Federal Reserve Bank of New York/Equifax Consumer Credit Panel (CCP). The CCP is a nationally representative 5 percent random sample of the U.S. population with an Equifax credit report collected from 1999 to the present. The consumer records are anonymous, and the identity of the consumers is not provided. The sample is structured in a quarterly panel, so that once selected for inclusion an individual stays in the panel until they no longer have a credit history. New individuals are added to the panel each quarter to maintain the 5 percent sample, thus repopulating the panel to account for both attrition and population growth. Each quarter, the CCP reports all of the information available on a credit report including the Equifax Risk Score, account balances and delinquencies across a wide range of categories, birth year and geographic location down to the census tract level.

We begin with every individual in the CCP from 2009 to 2019, 15,561,633 people, and apply the restrictions below in the order listed. We drop quarters after a person's death and people with no recorded age. Because 2010 census tracts first appear in the CCP in 2014 Q2, we keep people present in that quarter or who enter the CCP later, and use each person's own address history to place earlier quarters in 2010 tracts. People who moved from a 2000 tract that was later split cannot be placed this way, so they enter the sample from the first quarter their tract can be identified.[7] We drop anyone with a bankruptcy on file in the year before they enter the sample, as people in bankruptcy generally lack access to new credit, and anyone living in a

Service, 2015, 2016). In 2014–15 these mid-year adopters were 414 of the 14,214 CEP schools (Center on Budget and Policy Priorities, 2015), so for about 3 percent of schools the onset may be timed a quarter or two off.

[7] We drop 1,080 individuals, or 0.007 percent of the full sample, for whom we do not have enough locatable quarters to fully perform pre-entry filtering.

CEP-covered tract in any of the four quarters before they enter, so that no one enters the sample already covered. We then keep the quarters in which a person is aged 25 to 54, the ages most likely to have a child in school. Using five-year American Community Survey data for 2013 to 2017, we keep person-quarters in census tracts where at least 50 percent of households with children have a child aged 6 to 17. This removes tracts where most families have only children too young for school, whom CEP would not reach. Within the tracts kept, the share averages 62 percent (Table 1). Each person must then remain in the sample for at least four consecutive quarters. Finally, we drop quarters after a person's first bankruptcy filing, since from then on their use of credit is constrained by the bankruptcy proceedings rather than their own choices.

Key dependent variables used in the analysis are (1) the Equifax Risk Score, (2) access to any credit as well as specifically auto, revolving and mortgage loans, (3) indicators of severe delinquency overall and specifically in auto, revolving and mortgage accounts and (4) entering into bankruptcy or foreclosure. Severe delinquency (SD) means an account is 90 days or more past due. The Equifax Risk Score is similar to a FICO score and reflects the probability of being severely delinquent in the next 24 months. It ranges from 280 to 850, with higher numbers being less risky. While the specifics of the calculation of the Equifax Risk Score are unknown, credit scores in general are influenced by payment and credit history, what types of credit accounts are open, the recent opening of new accounts and what share of available credit is utilized.[8]

Access to credit is measured both as an indicator of whether the individual has any credit accounts and the number of accounts overall and for each of the three categories of auto, revolving and mortgage accounts. Severe delinquency is captured in multiple ways, including an indicator of whether the individual has any accounts in severe delinquency, as well as measures

[8] https://www.equifax.com/personal/education/credit/score/articles/-/learn/how-is-credit-score-calculated/

that are calculated using both dollars and the number of accounts in delinquency. Specifically, the measures include the number and share of accounts in severe delinquency and the total amount of debt, in 2014 dollars, and share of debt in severe delinquency overall and separately for the three categories of auto, revolving and mortgage accounts. As noted in Argys et al. (2020), higher income individuals have, in general, higher levels of debt and, thus, the potential to have higher levels of delinquency. Thus, an increase in delinquent debt could represent an increase in financial fragility or the ability to take on more debt, or both. It could also be the case that someone paid off a debt and closed the account, resulting in the delinquent debt accounting for a larger share of debt. For these reasons, we consider these variables as a group rather than focusing on a single measure.

Tract level characteristics, including household income decile, median household income, and the share of households with children that have a child aged 6 to 17 within the tract, were obtained from the American Community Survey 2013–2017 five-year estimates and held fixed across the panel (Manson et al., 2019).

3.3 Sampling Framework

After the noted restrictions were put in place, we were left with a sample of 150,470,569 person-quarters for 5,815,590 unique individuals. The sample was first split into the four waves of CEP rollout. We refer to the states first eligible in 2014 as the main pool and the three pilot waves together as the pilot pool. The first wave includes data through the school year 2011–2012 for the first three pilot states of Illinois, Kentucky, and Michigan. The second wave includes data through the 2012–2013 school year for the second-year pilot states of the District of Columbia, New York, Ohio, and West Virginia. The third wave included data through the 2013–2014 school year for the third-year pilot states of Florida, Georgia, Maryland, and Massachusetts. The

main sample includes data through 2019 Q2 for all of the remaining states that became eligible for the CEP program in the 2014–15 school year. For purposes of estimation, we created a sub-sample that includes everyone that resided in a treated tract and a sample of never treated individuals that are selected to resemble the broadly-defined treated group, on the tract characteristics of median household income deciles and share of households with children that have a child aged 6 to 17 in the tract within each of these CEP rollout waves.[9] The full sample for analysis includes 1,527,253 treated and 2,166,202 control individuals, with 88,929,396 person-quarters. Of these, 22,154 individuals are observed in only one quarter, largely because they file for bankruptcy in their first quarter and their later quarters are dropped, and they drop out of estimation because the person fixed effect cannot be estimated for them. That leaves 1,517,103 treated and 2,154,198 control individuals and 88,907,242 person-quarters. In the main sample, 17,234 individuals drop out in the same way, leaving 58,036,372 person-quarters. Controls are drawn 1:1, without replacement, at random within cells defined by pool, baseline tract income decile and tercile of the tract share of households with children that have a child aged 6 to 17, both measured at a person's first observation; the terciles are cut on each pool's treated distribution, and a cell with fewer controls than treated persons keeps every control available. The treated and control counts differ from a 1:1 draw because the draw counts as treated some individuals who are never covered in the analysis, as footnote 9 explains.

The main analytic sample is restricted to the states first eligible for CEP in the final rollout, school year 2014–2015, which excludes individuals who resided in the ten states and the

[9] Controls are drawn 1:1 within cells, with shortfalls of 206,361 controls in the main pool and 83,139 in the pilot pool, where some cells had too few controls. The draw counts as treated anyone ever observed in a tract that adopts CEP, while the analysis counts an individual as ever-CEP only if they lived in a tract while it was covered, so 519,327 individuals drawn as treated are never-CEP in the analysis. The build then removes individuals living only in New York City, which we exclude because citywide school choice means residence does not identify a school, or in tracts without coverage data, 12,569 drawn as treated and 122,774 drawn as controls.

District of Columbia that adopted CEP during the phase-in, even if they later moved to a state outside it. Adoption in those phase-in states is staggered, which makes the comparison group time-varying in a way that complicates estimation. The baseline statistics from both samples, split between those that are ever covered by CEP and those that are never covered by CEP, are shown in Table 1.

3.4 Sample Characteristics

The characteristics of both the full sample, which spans all years, and the main sample, restricted to the states first eligible in 2014, are described in Table 1. There are three methods for capturing CEP coverage, which is defined as the share of school-age children in a tract who likely attend a school covered by CEP. The primary analysis uses coverage at onset, that is, the dose of CEP in a tract when it is first treated, and the gain in CEP in each year after the initial baseline as more schools likely attended by students in the tract adopt CEP. An alternative specification uses contemporaneous coverage, the share covered in the current quarter. Finally, in the models that require a binary variable, there is an indicator that turns on in a person's onset quarter and stays on. In the main sample the average dose of CEP at baseline is 66.83 percent and the average gain is 5.85 percentage points. In the full sample the corresponding figures are 67.26 percent and 7.57 percentage points. Across all quarters observed for ever-CEP individuals, 39.67 percent fall after onset in the main sample and 42.73 percent in the full sample.

Table 1 and the counts here describe the estimation sample, the person-quarters the regressions use. The full sample comprises 3,671,301 persons observed across 88,907,242 person-quarters between 2010 Q1 and 2019 Q2. Ever-CEP individuals account for 1,517,103 persons and 41,615,075 person-quarters, and never-CEP individuals for 2,154,198 persons and 47,292,167 person-quarters. The main sample, restricted to states first eligible in the 2014–15

school year, comprises 2,477,422 persons and 58,036,372 person-quarters, of which 944,602 persons and 25,649,637 person-quarters are ever-CEP. Ever-CEP individuals are observed for longer on average, 27.2 quarters against 21.1 in the main sample, because the never-CEP group is drawn to match on tract characteristics rather than on panel length.

Consistent with the program's targeting, ever-CEP individuals are observably poorer and less financially integrated than never-CEP individuals. The mortgage gap is the widest of the four trade lines and the auto gap the narrowest. Ever-CEP individuals hold fewer accounts on average, and severe delinquency rates are higher across every debt category except mortgage, where the two groups are indistinguishable. The average median household income decile is 4.8 for ever-CEP individuals against decile 5.4 for never-CEP individuals. The average age in the ever-CEP group is over three years lower than in the never-CEP group. The same patterns hold in the full sample.

Level differences of this kind are easier to read once they are scaled by the dispersion of the variables themselves. We therefore compute normalized differences, the difference in group means divided by the root mean squared standard deviation of the two groups, and treat as large any variable exceeding 0.25 in absolute value on either sample, a rule of thumb from Imbens and Rubin (2015) for the point beyond which linear regression estimates tend to be sensitive to the specification (see also Imbens and Wooldridge, 2009). Only two variables cross the threshold, age and the Equifax Risk Score. Birth cohort is absorbed by the person fixed effects and a common age trend by the state-by-quarter effects, so only differences in how credit changes with age remain, which the age-band-by-quarter test in Section 7.1 addresses. The Equifax Risk Score is an outcome rather than a covariate, so the threshold does not bear on it in the same way. A difference-in-differences design requires the two groups to share a trend in the outcome, not a

level, and person fixed effects absorb the gap in levels. A gap this large does raise the question of whether the two groups were on different trajectories before CEP, which the pre-trend tests in the event studies of Section 7.1 examine.

These pre-adoption differences motivate the stratified sampling design and the person fixed effects, which absorb time-invariant individual and tract characteristics. They do not, on their own, threaten identification, since the design asks whether the two groups move in parallel rather than whether they start in the same place. Identification related issues are further discussed in Section 7.

Coverage among treated tracts is concentrated toward full adoption rather than spread evenly across the unit interval. In the main sample the mean of baseline coverage among ever-CEP individuals is 0.6683 with a standard deviation of 0.2911, and contemporaneous coverage averages 0.6562 with a standard deviation of 0.3374. The gain term is small by comparison, averaging 0.0585 with a standard deviation of 0.1554, so most of the coverage a person experiences arrives at onset rather than accumulating afterwards. This is the empirical basis for the two-term specification in equation (4.2), which separates the level of coverage a person first experiences from any subsequent expansion, and it is why the binary indicator used by the staggered-robust estimators captures most of what the continuous dose does.

4. Empirical Framework

The starting point of our analysis is a standard difference-in-differences design. We compare how financial outcomes changed for individuals in census tracts that gained access to CEP relative to households in tracts that did not, over the same time period. Any trends that would have occurred regardless of CEP, such as business cycle fluctuations, national credit market

conditions, or demographic change, affect both groups equally and cancel out in the comparison. In its simplest form:

$$Y_{ict} = \alpha_i + \delta_t + \beta \cdot T_{ct} + \varepsilon_{ict}, \qquad (4.1)$$

where Y_{ict} is the financial outcome for individual i in census tract c at time t, α_i is the individual fixed effect, which absorbs time-invariant differences across individuals, δ_t is the year-quarter fixed effect, which absorbs common shocks affecting all individuals in a given quarter, and β is the causal effect of interest. T_{ct}, the treatment variable, is introduced here as a single time-varying indicator of CEP exposure for exposition.

Equation (4.1) abstracts from two features of the data. First, common shocks that hit a state in a given quarter, such as a state EITC change, a regional labor-market swing, or a state-level credit-reporting practice, could be correlated with CEP adoption timing if adoption itself tracked state-level conditions. Second, a person's exposure to CEP is not a single on/off switch. Coverage at a person's tract can rise after they first become exposed as more schools in the same tract elect the provision in later years. We address both by adding a state×quarter fixed effect and by splitting exposure into two terms, the share of students in CEP schools in the first exposure and the increase in the share over time. The estimating equation is now:

$$Y_{ict} = \alpha_i + \beta_{abs} \cdot A_{it} + \beta_{gain} \cdot G_{it} + \gamma_{s(c)t} + \varepsilon_{ict}, \qquad (4.2)$$

where A_{it} is coverage at individual i's own onset quarter, i.e. the first quarter person i's tract of residence has any CEP coverage. It is zero in every quarter before onset and is held fixed from that quarter forward. G_{it} is any further coverage the tract acquires after person i's onset, net of A_{it} and zero before and at onset.[10] In addition, $\gamma_{s(c)t}$ is a state×quarter fixed effect, where $s(c)$ denotes

[10] Note that coverage does fall in some tract-quarters, but almost never for policy reasons: only 1.9 percent of treated tracts ever lapse, and among person-quarters where coverage falls the median in-tract decline is 0.0047 against 0.4000 for individuals who moved. Since the loss margin is identified almost entirely on individuals who moved,

the state containing tract *c*. It absorbs shocks common to a state in a quarter, including a state's own EITC or TANF generosity in that quarter. The standard errors are clustered at the county level, as school districts are not available in the CCP.

The two CEP variables measure different things. A_{it} asks whether the level of tract coverage a person first experienced matters while G_{it} asks whether tract coverage that keeps expanding, such as when an elementary school's CEP election is followed some years later by the middle and high school in the same feeder pattern, does additional work per unit of coverage. The combined effect of the coverage a person actually receives is the sum of the two coefficients weighted by how much of each a person actually receives on average, $0.6683 \times \beta_{abs} + 0.0585 \times \beta_{gain}$, reflecting a sample mean onset coverage of 0.6683 and a mean subsequent gain of 0.0585. We report this combined marginal effect as the headline number for each outcome, alongside the two component coefficients.

5. Results

5.1 Full Sample and Goodman-Bacon Decomposition

The first step for this analysis is to determine if the staggered implementation of the CEP programs has implications for the analysis. Combining all states and estimating equation (4.2) as a single regression produces a pooled two-way fixed effects estimator that could be problematic. Goodman-Bacon (2021) shows this equals a weighted average of all 2×2 DiD comparisons:

$$\hat{\pi_{TWFE}} = \Sigma_{k \neq l}\, \omega_{kl} \cdot \hat{\pi_{kl}} \qquad (5.1)$$

with k denoting an early treatment group and l denoting a later-treated group. The weights ω_{kl} are nonnegative and sum to one, but some of the comparisons they weight use

which is plausibly endogenous, and policy lapse is too rare to identify on its own, the loss margin is not included in the analysis.

already-treated tracts as controls for later-adopting ones, and when effects change over time those comparisons can carry the wrong sign. Baker, Larcker, and Wang (2022) show this contamination can be quantitatively large. For a continuous treatment like CEP coverage, Callaway, Goodman-Bacon, and Sant'Anna (2024) show that the TWFE coefficient is additionally hard to interpret because the linear specification conflates effects at different dose levels.

The Goodman-Bacon decomposition utilizes the binary CEP variable. Thus, we replace the baseline and gain CEP variables with an indicator of any CEP coverage in the tract and estimate the Goodman-Bacon decomposition on the full sample, 2010 through 2019 Q2. When analyzing the pooled TWFE estimate on the full sample, the decomposition assigns 83.96 percent of the estimate's weight to comparisons between never-treated and eventually-treated tracts, 10.07 percent to valid early-versus-late comparisons, and a relatively small 5.97 percent to the forbidden late-versus-early comparisons. However, the forbidden component is wrong-signed relative to the valid ones on 29 of 31 outcomes, the two exceptions being the share of mortgage debt in SD and the number of auto accounts. Removing the forbidden comparison raises the absolute value of the estimate on every outcome, by 7.97 percent on average (median 7.65 percent), with increases ranging from 5.6 percent to 11.0 percent. Thus, although it is small by weight, the consistent direction of the sign suggests that a change in methodology is in order.[11]

To help avoid the forbidden comparison, the main estimation sample is limited to the states that only became eligible in 2014, thus removing the ten pilot states and the District of Columbia. Limiting the sample reduces the forbidden comparison by 42 percent, from 5.97 to 3.44 percent of the weight, with 87.56 percent on comparisons against never-treated tracts and

[11] The full results of the binary CEP model run on the full sample are shown in Online Appendix Figure OA1, as the full pool row.

9.00 percent on valid early-versus-late comparisons, and the sign is wrong on 17 of the 31 outcomes, rather than 29, but does not fully remove it as there remains the issue of the schools that are eligible for coverage and enroll later, or tracts becoming eligible in the future. The results for the dose estimation, marginal effects that combine the impact of coverage at onset and the gain in CEP over time, are available in Figure 1, with the mortgage outcomes available in Appendix Figure A1. The full dose coefficients are in Appendix Table A1, the binary results on both samples are in Online Appendix Figure OA1 and are discussed in Section 6.2, and the Goodman-Bacon decomposition by outcome is in Appendix Table A2.

5.2 Main Results: Credit Access

CEP coverage is associated with increased participation in credit markets. The probability of having any credit account increased by 1.06 percentage points, on a baseline of just under 80 percent for the treated group, as shown in Figure 1. Comparing across trade lines, revolving account holding increases by 1.18 percentage points and auto by 0.88 percentage points. Mortgage account holding falls, but this result is not robust, reversing sign under the three staggered-robust estimators we report in Section 6.4. However, as a share of baseline, auto has the largest proportional gain, rising 2.5 percent of baseline compared to 1.75 percent in revolving. These effects are on people who already have a credit file, about four in five of whom hold an account before exposure, rather than on people outside the credit system.

The number of accounts follows the same pattern and allows us to separate entry from accumulation of accounts. The number of accounts increases by 0.086 accounts in aggregate, 0.055 in revolving, and 0.007 in auto. Combining the two access results allows us to explore how many people have an account versus how many they hold. Overall and for revolving accounts, the increase in the number of accounts dominates and suggests that households already in the

market are taking on more credit accounts. For auto, it is the opposite, with the entry effect dominating, suggesting that new auto lines are being taken up by households without auto loans initially.

However, the dynamics of credit access differ from the margins. The shape of each effect over time comes from an event study, in which the two dose terms are replaced by indicators for each quarter relative to a person's onset, each multiplied by the person's coverage at onset, held at that value in every quarter. The quarter before onset is omitted so every coefficient is read against it, and quarters nine or more before or after onset are pooled into the end points. Because each indicator is multiplied by coverage at onset, the event-study coefficients are effects of full coverage. At mean coverage at onset, 0.6683, they are about two-thirds as large, so the step at onset for any credit account, 1.01 percentage points, corresponds to about 0.68 at mean exposure. These results are shown in Figure 2. Holding any accounts is a sharp step up that appears at onset and stays flat, i.e. a one-time entry into the credit market with no adjustment period. The effect on the number of accounts follows a similar but slightly delayed path, as it takes about three quarters to adjust, but then remains flat. This result suggests that after households are exposed to CEP they enter the credit market once and accumulate new accounts over a short period of time. Section 7 shows that this step survives every check on the pre-trend, while the level of the effect two years out does not survive all of them.

5.3 Main Results: Severe Delinquency

The marginal effects from the delinquency analyses are presented in Figure 1. Because each measure captures a different aspect of delinquency, and a higher delinquent balance can reflect more debt as well as more fragility, we do not focus on any one measure but on the entire family of delinquency measures by trade line. Within each measure of credit, there are three measures

that are in levels, including whether an individual has an account in severe delinquency (SD), how many accounts are in SD, and the dollar amount of SD, in 2014 dollars. The other two measures are shares, including the share of accounts in SD and the share of total debt in SD.

In the aggregated debt measures, shown in the first panel of Figure 1, the levels all increase while the shares are unchanged on average. The share of households with any SD debt increases by 0.40 percentage points from a baseline of 20.09 percent, which is 1.97 percent of baseline. The number of accounts in SD increases by 0.019 accounts per person from a baseline of 0.40 accounts, which is an increase of 4.63 percent. Delinquent balances increased by 4.6 percent. The shares measure delinquency per account and per dollar. In the dose specification they do not move, but that average combines a dip and a rise. The share of accounts in severe delinquency falls by 0.48 percentage points at onset, as new accounts start out current, and is 0.41 percentage points above its pre-onset level two years later. Once the pre-adoption trend is continued, the share rises by 0.33 to 0.95 percentage points on average over the first nine quarters in the Rambachan and Roth (2023) bounds described in Section 7.2, which allow the pre-trend to continue after onset. Our results suggest a modest worsening in repayment, which does not appear among people observed continuously before and after onset.

The results by type of debt are reported in the remaining panels of Figure 1 for revolving and auto.[12] Revolving has the same shape as the aggregate, although at lower levels. The probability of having any SD revolving debt increased by 1.07 percent of baseline (13.09 percent) compared to 1.97 in aggregate. The number of accounts in SD increased 2.63 percent of baseline and the total debt in SD increased by 1.46 percent. The shares are both flat to slightly

[12] The seven mortgage outcomes are in Appendix Figure A1.

negative, but not significantly so. There is no significant marginal impact on the mortgage delinquencies.

Auto is the only debt type that saw all five measures increase after the CEP treatment began. In addition, as the SD baseline in each delinquency measure for auto is so low, the proportional effect is quite large for similar marginal effects. The incidence of SD increased by 8.44 percent, which is the largest proportional delinquency effect, but the baseline is a very small 4.01 percent, compared to the double-digit bases for revolving and aggregate debt. The number of accounts in SD increases by 0.004 accounts per person, or 9.1 percent of baseline and the amount of debt in SD for autos increased by 3.25 percent. In addition, the share of auto accounts SD rises 0.248 percentage points on a base of 3.55, and the share of auto dollars SD rises by 0.221 percentage points on a base of 3.44 percent, which is 6.99 and 6.41 percent of base, respectively.

Across the family of measures, the three levels of severe delinquency rise in aggregate, in revolving and in auto credit, the two shares rise only in auto, and no mortgage measure changes significantly. The increases are small, 2 to 5 percent of baseline in aggregate and 1 to 3 percent in revolving credit. They are larger in auto, 3 to 9 percent, but from low bases, with any auto severe delinquency rising 0.34 percentage points from 4.0 percent. The shares rise modestly in the other trade lines once the pre-adoption trend is continued, and our results suggest that the rise in severe delinquency does not come from people observed continuously before and after onset, and may come from people new to the credit panel.

5.4 Creditworthiness and Catastrophic Events

The final three outcomes abstract away from credit accounts, per se, and focus on how the change in account behavior affects the person's overall credit status. The first outcome is the

Equifax Risk Score, as shown in the final panel of Figure 1. On average, the Equifax Risk Score falls by 0.94 score points on a base of 652.5, which is 0.14 percent. A decline of this size is not inconsistent with the opening of new accounts, since each brings a new inquiry and lowers the average age of the accounts on file, both of which lower credit scores. However, even as individuals add accounts and severe delinquency rises, the Equifax Risk Score, which measures the probability of severe delinquency over the next two years, falls only slightly. Furthermore, as shown in Figure 2, the impact is somewhat transitory. In event time, the effect is -0.61 at onset, troughs at -1.83 two quarters later and recovers to -0.92 by the eighth, so half of it is gone within two years. In addition, the Equifax Risk Score moves opposite to delinquency in event time. Any severe delinquency rises over exactly the quarters in which the Equifax Risk Score recovers. Because a new severe delinquency lowers a credit score, the pattern is a puzzle, and we cannot explain it with these data.

Bankruptcy and foreclosure both rise, by 0.0184 and 0.0043 percentage points, respectively. Foreclosure is measured among people with no foreclosure on file when they enter the panel, and the 73,160 people in the full sample, 2.0 percent, who entered with one are excluded, since their first foreclosure came before we observe them. Against baselines of 0.2476 and 0.1092 percent these are 7.43 and 3.9 percent of their own bases, but the proportional figures are a poor guide to magnitude when the base rate is this low. In absolute terms, the increases amount to roughly seven additional bankruptcies and two additional foreclosures per ten thousand individuals per year. For comparison, the share holding any credit account rises by about 106 per ten thousand, although that is a stock rather than an annual flow. We are also cautious in reading the increase in filings as deterioration. Filing for bankruptcy is itself costly, and Gross, Notowidigdo, and Wang (2014) show that households constrained out of filing do so

when liquidity arrives, without the additional liquidity causing bankruptcies that would not otherwise have occurred. A small increase in filings among households whose recurring food costs have fallen is consistent with that mechanism rather than with a worsening of the underlying financial position. The bankruptcy event study should not be read as an effect. Because bankruptcy is a first-filing indicator and a person's later quarters are dropped, a filing is always the last quarter in which that person is observed, and person fixed effects can turn that into a pattern across event time. Among ever-CEP individuals, the raw filing rate is lower after onset than before, about 0.10 percent a quarter against 0.14 to 0.22 percent, apart from a rise in the onset quarter itself. We therefore rely on the dose estimate for bankruptcy. The foreclosure effect, in turn, is a small step at onset that fades within the first year, survives neither smoothness bound, which limits how far the pre-trend may bend after onset, and only the smallest relative-magnitudes bound, which limits the violation after onset to a multiple of the largest change before it, so we do not treat it as a finding.

5.5 Heterogeneity

As noted in Section 2.1, where we set out why the cash-flow benefit of CEP should be largest where the fewest students were already eligible for free or reduced-price meals, the effect of CEP is expected to vary with prior eligibility. To test this, we estimate the event study separately for ever-CEP individuals in each quartile of their tract's 2010 free-and-reduced-price share, from the Common Core of Data and measured before onset. This comparison is within ever-CEP individuals, among tracts that adopted at different times, so it is not the same parameter as the dose estimates, but it shows directly how the response varies with prior eligibility. Figure 3 shows the step at onset and the effect two years after onset net of each quartile's pre-trend for four headline outcomes, and Online Appendix Figure OA2 shows all 31. The step at onset is the

change in the outcome between two quarters before onset and the onset quarter, and the detrended effect is the effect two years after onset net of the pre-onset trend carried forward, as Section 7.1 describes.

The step at onset falls steadily with prior eligibility. For any credit account, for example, it is about four times as large in the lowest quartile as in the highest, and the number of accounts and the Equifax Risk Score follow the same ordering. Across outcomes, fewer steps are significant as the share of students already eligible rises, and the step is largest in the lowest quartile for 15 of the 31.[13][14]

Two years after onset the pattern is less clear, because the pre-trend is steepest in the lowest quartile, and carrying it forward ten quarters removes more of the effect there. For any credit account the detrended effect in the lowest quartile is negative, while the number of accounts remains higher in the lowest two quartiles. The rise in severe delinquency net of the pre-trend is concentrated in the lowest quartile, where any severe delinquency is about 1.5 percentage points above its trend. As for the headline results, we therefore rely on the step at onset for how the response varies with prior eligibility.

6. Robustness

6.1 Operationalizing CEP Coverage

The treatment in the main analysis is the share of students in schools associated with each tract who are covered by CEP. It is zero before any school in the tract elects coverage, takes a level at onset, and can rise as further schools elect. The main specification splits this into coverage at

[13] Estimates by tract median household income, in Online Appendix Figure OA3, show the access gains peaking in the sixth to eighth deciles, consistent with gains being largest where more families had incomes just above the cutoffs for free and reduced-price meals.

[14] For any credit account, the step at onset is 2.18 percentage points in the lowest quartile, 1.31 in the second, 0.71 in the third and 0.54 in the highest. The number of outcomes with a significant step is 16, 13, 11 and 8.

onset, held forward, and coverage gained afterwards. This section considers three alternative parameterizations, both to test how far the result depends on that split and to establish the binary specification the robustness apparatus requires. Table 2 reports all four definitions for every outcome.

The first alternative, the absorbing specification in Table 2, drops the gain term and asks only what being exposed to CEP does, using coverage at onset held forward as a single continuous regressor, known as the absorbing version of CEP. Scaled by the mean of its own regressor it returns 0.909 of the main model's estimate at the median, with a range of 0.800 to 0.962 across the 22 outcomes significant under both. The gain term therefore contributes a consistent tenth rather than something that differs by outcome, so the two-term split refines the headline rather than producing a different result.

The second alternative, the contemporaneous specification in Table 2, replaces coverage at onset with coverage in the current quarter, so the treatment follows a person rather than remaining fixed once it arrives. This specification has no parameter for arrival. Where the main model identifies a permanent level shift at onset, current coverage fluctuates and pools that shift with everything that follows it. Scaled by the mean of its own regressor it delivers about a quarter of the main results on credit access, with a median of 0.269, although mortgage access and accounts run three to three and a half times the main estimate, and about three quarters on delinquency, with a median of 0.745. That contrast suggests that the access effect is almost entirely a permanent level shift at onset, while delinquency tracks current coverage much more closely.

The third alternative measure of CEP, the binary specification in Table 2, replaces the dose terms with a single absorbing indicator that turns on when a person is first exposed and

remains on. This is the specification the staggered-robust estimators require, since each organizes the data around the period a unit switches from untreated to treated and none accepts a continuous dose. It captures the onset level shift the first two specifications identify as the main driver and discards the intensity of coverage.

The binary and the dose agree in sign on 29 of the 31 outcomes. The two exceptions are the share of mortgage accounts and the share of mortgage debt in SD, and neither is distinguishable from zero under either specification. The disagreement is between two estimates of zero falling either side of it and not between two findings.

Where both specifications are significant, the binary is, in general, the smaller of the two by a margin that differs across trade lines. Revolving is reproduced almost exactly, at 0.97 to 1.15 of the dose estimate with a median of 1.02. Auto is understated, at 0.69 to 0.89 with a median of 0.73. The aggregate falls between them at 0.87 to 0.93, median 0.90. The binary assigns every covered tract the same treatment, whatever its coverage and however much it gains later, so it falls furthest short where the effect depends most on intensity. That is auto, which of the trade lines follows current coverage most closely, and least where the effect is a step at onset, as in revolving credit. These ranges cover the outcomes significant under both specifications as the ratio is uninformative where the denominator is an estimate of zero. Mortgage access is significant and negative under both specifications, at -0.26 and -0.15 percentage points. The five mortgage delinquency measures are too small under either specification to be meaningful.

6.2 Impact of Sample Restriction

The restriction to states first eligible in 2014 was justified in Section 5.1 by the Goodman-Bacon decomposition. In this section, we explore the implications of that choice. Estimating the binary specification on the restricted sample and on the full pool holds the treatment, the unit, the fixed

effects and the clustering fixed, so any difference is the sample alone. These results are shown in Figure 4 and Online Appendix Figure OA1.

The estimates from these two samples agree in sign on 29 of the 31 outcomes. Twenty are significant in both samples and eight in neither. Among the twenty significant in both, the full-pool estimate is 1.02 times the restricted one at the median, with the middle half of the ratios between 0.94 and 1.08 and the full range from 0.60 to 1.54. The two sign disagreements are mortgage share measures estimated in the fifth decimal place and significant in neither sample.

Three outcomes disagree on significance. Both mortgage access measures are significant and negative on the restricted sample and indistinguishable from zero on the full pool, which is the first of several indications that the negative mortgage result is a property of the restriction rather than of the program. The third is the share of revolving accounts in SD, negative in both samples but significant only on the full pool, where the larger sample estimates the same effect more precisely.

In general, including the pilot states raises the access effects. The delinquency effects are similar, somewhat larger for accounts and debt in severe delinquency across all trade lines and somewhat smaller in revolving and auto credit. Our results suggest the restriction costs very little, and that allows the remaining comparisons to be about estimators rather than about samples.

6.3 Staggered-Adoption Estimators: Methodology

Even after restricting the headline sample to states first eligible in 2014, the staggered timing of tract adoption within that cohort raises concerns about the pooled two-way fixed effects estimator. We therefore implement three staggered-robust estimators, the Callaway and Sant'Anna (2021) group-time average treatment effect (CS), the Sun and Abraham (2021)

interaction-weighted estimator (SA), and the de Chaisemartin-D'Haultfœuille (2020, 2026) heterogeneity-robust event-study estimator (dCDH). Each avoids the forbidden comparison described in Section 5.1, in which already-treated units serve as implicit controls for later-treated units, and each does so by identifying every treated cell against never-treated units rather than against other treated cohorts. The differences between them are what make the comparison in Section 6.4 informative.

SA augments the TWFE regression with cohort-by-event-time interactions and is estimated at the person level, with person and state-by-quarter fixed effects; saturating cohort by event time is not feasible on the full sample, so it is fitted on a cell-stratified draw of 22,475,942 person-quarters, 38.7 percent of the estimation sample. CS compares each cohort's change from its own base period with the same change among never-treated units. dCDH compares groups whose treatment changes with groups whose treatment does not, chains the comparisons into dynamic effects, and builds placebos the same way before the change; it is the only one of the three that accepts a dose, so we fit it on coverage at onset. CS and dCDH are estimated on tract-quarter means weighted by person count, because CEP varies by schools and coverage is assigned to tracts, and defining cohorts by a person's own first exposure would date a mover from the quarter they arrive. All three estimators use the main 2014 sample. SA shares the unit, treatment and fixed effects of the TWFE binary specification, so a disagreement between those two is the estimator alone; CS and dCDH also differ in unit, and cannot absorb a shock specific to one state in one quarter. Because the estimators' default summaries average different things, we report each as an average of its effects from onset, over the nine quarters from k = 0 to 8 for SA and CS and over the first eight horizons for dCDH, which begin at onset, with dCDH's standard error computed by the delta method from its own covariance matrix. The bounds in

Section 7.2 instead use the Callaway-Sant'Anna average over the twelve quarters from onset, with twelve quarters before it.

6.4 Staggered-Adoption Estimators: Results

Figure 4 reports all six specifications for selected outcomes, with the remainder in Online Appendix Figure OA1. Eighteen outcomes are significant with the same sign under both the TWFE binary specification, whose event studies are in Online Appendix Figure OA4, and SA. The magnitudes are of the same order, at 0.95 against 0.69 percentage points for any credit account, 0.080 against 0.087 accounts, 0.35 against 0.48 percentage points for any severe delinquency, 0.017 against 0.021 accounts in severe delinquency, and 3.9 against 4.5 percent for delinquent balances. Of the 22 outcomes significant under the TWFE binary specification, 18 are also significant with the same sign under SA, 19 under CS and 17 under dCDH. Every result on access and on the levels of delinquency that we report as a finding falls in this group. The person-level estimators find the effects arriving as a step at onset, and the tract-level ones as a steadier ramp, but their averages from onset agree (Online Appendix Figures OA5 to OA7).

Four outcomes are significant under Sun-Abraham and dCDH but not under TWFE, the shares of accounts and of debt in severe delinquency, across all trade lines and in revolving credit. Callaway-Sant'Anna finds the same four and five mortgage outcomes, the shares of mortgage accounts and of mortgage debt in severe delinquency, any mortgage severe delinquency, mortgage accounts in severe delinquency and mortgage debt in severe delinquency. For the three mortgage delinquency measures the signs agree with the TWFE binary estimates, but for the mortgage shares they do not. The staggered-robust estimators find the aggregate and revolving shares rising, by 0.13 to 0.23 and 0.09 to 0.34 percentage points, where the TWFE binary estimate is small and not significant. Because the shares dip at onset and rise afterwards,

forbidden comparisons pull the pooled estimate toward zero. Both the staggered estimators and the detrended effects therefore point to a modest rise in the shares in severe delinquency after onset, which does not appear among people with an unbroken credit file in the balanced-window comparison in Section 7.1. The staggered-robust estimator resolves an effect the pooled specification cannot distinguish from zero. Four outcomes significant under TWFE are not significant under dCDH, which is estimated on a switcher set less than half the size of the full treated group. They are bankruptcy, foreclosure, the number of mortgage accounts and the Equifax Risk Score, and mortgage access is significant under both with opposite signs. We read these as differences in precision rather than in conclusion. Only two sets of results, mortgage access and the Equifax Risk Score, are ever estimated with confidence in opposite directions, and both are taken up below.

Two sets of outcomes divide the estimators. The first is mortgage access. The TWFE binary specification finds small, significant falls in mortgage access and in the number of mortgage accounts; all three staggered-robust estimators reverse them, the fall does not survive the full pool, and mortgage access has, with auto access, the lowest relative-magnitudes breakdown value of the access outcomes. We do not treat it as a finding.

The second is the Equifax Risk Score. The person-level estimators find it falling, by 0.94 score points under the dose specification, 0.81 under the TWFE binary specification and 1.23 under Sun-Abraham, while Callaway-Sant'Anna, at the tract level, finds it rising by 0.59 score points. We examined two explanations for the disagreement. The first is that the tract averages describe a different set of scored people. Re-weighting the tract collapse by the number of people with a score leaves the tract-level estimate essentially unchanged, and restricting to people scored in every quarter they appear moves it further from the person-level ones, while the

person-level estimate on that restricted sample barely changes, so this does not appear to be the explanation. The second is movement in the Equifax Risk Score that preceded CEP. The tract-level estimators compare tracts across states without state-by-quarter effects, and their placebo tests reject for many outcomes, the dCDH placebos for 20 of the 31. For most outcomes the placebos are small, 0.2 to 2.3 percent of baseline, and smaller than the effect, at 14 to 72 percent of it. For the Equifax Risk Score alone they exceed the effect, so the tract-level estimate cannot separate CEP from movement that came before it. dCDH has the Equifax Risk Score falling at the first two horizons, as the person-level estimates do, and rising after the first year, so that its mean over the eight quarters is not distinguishable from zero. We therefore rely on the person-level estimate, whose pre-trend slope is flat in the person-level event-study tests of the parallel-trends section, although the joint test that all its pre-adoption coefficients are zero rejects, as it does for every outcome but foreclosure.

7. Parallel Trends and Sensitivity

7.1 Parallel Trends

The identifying assumption is that covered and uncovered tracts would have followed parallel paths without CEP. The event studies in Figure 2 and Online Appendix Figure OA8 provide the conventional test, which rejects that assumption. The joint null that all pre-adoption coefficients are zero is rejected for 30 of the 31 outcomes, all but foreclosure, and the dCDH placebos reject at the five percent level for 20. At 58 million person-quarters these tests detect departures far smaller than any effect we estimate, so they cannot tell us which results are affected. We therefore ask of each result whether it survives the pre-trend carrying on after adoption.

We summarize each event study with three contrasts, each answering a different question. The first is the pre-trend slope, fitted over the quarters from eight to two before onset. It asks

whether outcomes were already diverging before CEP arrived. It is significant for 13 of the 31 outcomes. That is fewer than the 30 for which the joint test rejects, because the joint test rejects for any pre-onset movement, including the annual school-year cycle, while the slope test asks only whether there is a steady trend. The second contrast is the effect two years after onset, net of that slope carried forward. It asks whether the effect exceeds what the pre-trend alone would have produced. The third is the change in the outcome between two quarters before onset and the quarter CEP begins, the step at onset. It asks whether the outcome jumps when CEP arrives. Figure 5 shows the second and third for four headline outcomes, and Online Appendix Figure OA9 shows all three, with the seasonal comparisons described below, for every outcome.

To take any credit account as an example, the pre-trend is real, at 0.056 percentage points a quarter ($t = 4.4$). The effect two years after onset survives carrying that trend forward. It is 0.44 percentage points ($t = 3.3$), about half the raw effect. The step at onset is 1.01 percentage points ($t = 17.7$). Across outcomes, all 25 that are significant two years after onset before the trend is removed remain significant after. The step is the least exposed of the three event-study contrasts to the pre-trend, because it spans only two quarters. In the same example, a trend of 0.056 percentage points a quarter would account for only about 0.11 of the 1.01 percentage-point step. The step is significant for 18 of the 31 outcomes. For those, the pre-trend would account for about a tenth of the step at the median and at most about a third, for total debt in severe delinquency, so every step remains in the same direction. Measured from the quarter before onset instead, the coefficient at onset for any credit account is 0.84 percentage points ($t = 21.3$), and it is significant for 21 of the 31 outcomes, so the step does not depend on the change into that quarter.

Credit outcomes also follow an annual cycle tied to the school year, so we check the step at onset against it. Any credit account, for example, rises by about 0.2 percentage points each time a school year starts, both in the school year before onset and in the one after. Jumps of this kind are significant for 10 of the 31 outcomes before onset and 13 after. Because onset is dated to the start of a school year, the step at onset contains one of these jumps. To remove it, we compare only quarters at the same point in the school year. The seasonal step is the change from four quarters before onset to the onset quarter, net of the same change a year earlier, and it removes both the annual cycle and a steady trend. It asks whether the outcome jumps at onset by more than it does at the start of any school year. In the same example, the seasonal step for any credit account is 0.77 percentage points ($t = 10.0$), about three quarters of the unadjusted step. Across outcomes, the seasonal step is significant for 15 of the 31, against 18 for the unadjusted step.

We apply the same comparison two years after onset, setting that quarter against the same quarter before onset and removing the pre-onset trend over the three years between them. On that basis the effect on any credit account is no longer significant. Revolving access, at 0.94 percentage points ($t = 4.5$), and the number of accounts, at 0.062 accounts ($t = 3.8$), remain higher. Across outcomes, 22 of the 31 remain significant two years out. We find that whether the effect on any credit account persists depends on how both the trend and the season are handled, while the step at onset does not.

Access steps up at onset and then stays at the new level. Two years later the coefficient on any credit account is where it was at onset, and the number of accounts builds over the first year and then levels off. Because the value of the program does not change once a tract adopts, a one-time step of this kind is the response we would expect, and the step is the estimate we rely

on for access. Whether households would eventually have reached the new level without CEP depends on whether the pre-onset trend would have continued, and our results suggest that, either way, CEP brought them there sooner. The pre-trend is unlikely to produce the step itself. Main-sample onsets begin in the third quarter of 2014, after the sharp rebound in credit from 2009 to 2012, and the slow drift that remains cannot produce a step of this size at each person's own onset.

For revolving credit, whose pre-trend is nearly flat, the new level is clearly above the trend. Two years out, after the seasonal adjustment above, revolving access is 0.94 percentage points higher and the number of accounts 0.062 accounts higher. The lasting gain in the number of accounts may come from people with thin credit files or who enter the panel later. Among people with an unbroken credit file, most of whom already held accounts, the number of accounts falls below its trend while revolving access rises above it, as the balanced-window comparison later in this section shows. Our results suggest this is the pattern a loosened budget constraint would produce.

The pre-trend has a natural explanation in the uneven credit shock of 2006–09 and the recovery from it (Mian, Rao, and Sufi, 2013). If that recovery continued after onset, it would be mistaken for an effect of CEP, so we test what drives it. Three tests narrow it down, and Online Appendix Figures OA10 and OA11 show the step at onset and the detrended effect two years out for every outcome under each of them and under the samples that follow. Comparing CEP tracts only with each other, early adopters against later ones, removes the gap between CEP and non-CEP tracts. It cuts the number of outcomes with a significant pre-trend from 13 to 7, while the detrended effects are significant for 25 outcomes either way. CEP tracts are also poorer on average than other tracts, and poorer tracts may have recovered from the credit shock on a

different path. If so, the pre-trend would reflect income rather than CEP. We test this by adding effects for each baseline tract-income decile in each quarter, over and above the state-by-quarter effects, which let tracts at each income level follow their own path. If the pre-trend came from an income-related recovery, those effects would absorb it. Our results suggest they do not. The pre-trends are significant for 12 of the 31 outcomes, against 13 in the main specification, and the slope for any credit account is unchanged ($t = 4.6$ against 4.4). With those effects the detrended effects shrink, to a median of three quarters of their published size, and three delinquency measures across all trade lines lose significance, while the step at onset barely moves, at 0.96 percentage points for any credit account against 1.01. We find, therefore, that the pre-trend does not appear to be an income-related recovery.

Ever-CEP individuals are also younger than never-CEP ones, and credit outcomes rise fastest early in the life cycle. The pre-trend could therefore reflect younger people moving along a steeper age profile. We test this by adding effects for each five-year age band in each quarter, which let every age group follow its own path. If the pre-trend came from age, those effects would absorb it. Our results suggest they do not. The pre-trends are significant for 12 of the 31 outcomes, against 13, and the slope for any credit account is unchanged. The detrended effects keep a median of nine-tenths of their published size, and the pooled dose effects hold as well, also at about nine-tenths. We find that the pre-trend does not appear to reflect differences in age.

The sample requires anyone in the panel before 2014 Q2 to have a credit file in that quarter, so the pre-2014 quarters include only people known to have a file later, which could by itself produce a rising pre-trend in access. Our results suggest that it does not. Among people observed in every quarter from $k = -8$ to $k = 8$, the pre-trend in any credit account is unchanged, 0.053 percentage points a quarter against 0.056, and across all 31 outcomes the balanced window

has as many significant pre-trends as the full sample, 14 against 13. The step at onset survives. It is significant with the same sign in both samples for 13 outcomes, at about half its full-sample size for access, as expected among people with an unbroken file, most of whom already hold an account. Two years out, any credit account is not distinguishable from its trend, and the number of accounts is significantly below it, by 0.047 accounts, while revolving access remains above it (0.87 percentage points, t = 3.9). And among people with an unbroken file, delinquency rises steadily before onset, any severe delinquency by 0.11 percentage points a quarter, so continuing that trend puts the detrended delinquency effects below zero, and any severe delinquency is 1.5 percentage points lower two years out (t = -6.5). Our results suggest that the modest rise in the delinquency shares after detrending in the full sample does not come from people with an unbroken credit file, and it may come from people new to the credit panel.

A stricter test keeps the same people but drops their quarters before 2014 Q2, so that no observation depends on a credit file observed later. Dropping those quarters leaves the tracts that adopted in 2014 with almost no quarters before onset, so the pre-trend in this sample comes from tracts that adopted in later years, whose quarters before onset fall after 2014 Q2 and are free of the selection. A pre-trend remains, and is steeper. In any credit account, for example, it is 0.13 percentage points a quarter against 0.056 (t = 5.7). Across outcomes, the pre-trend is significant for 14 of the 31, against 13 in the main specification. The step at onset for any credit account is about the size it is in the full sample, 1.09 against 1.01 percentage points, and for the number of accounts about 30 percent larger, 0.063 against 0.048 accounts. It is significant for 15 outcomes, against 18. Two years out, the number of accounts is 0.081 accounts above its trend, the same as in the full sample (t = 5.7). As in the balanced window, any credit account is not distinguishable from its trend. Revolving access is 0.75 percentage points above its trend, just short of

significance (t = 1.94). The detrended effects remain significant for 16 outcomes, against 25, which may partly reflect the smaller sample. Our results suggest that the selection is not needed to produce a pre-trend, although this test cannot rule out that it contributes to the pre-trend among the tracts that adopted in 2014.

A final possible explanation for the pre-trend is anticipation, with households taking on credit before CEP arrived because they expected it. Anticipation depends both on how likely CEP was and on how many families stood to gain from it. In tracts whose schools had the highest shares of students eligible for free and reduced-price meals, CEP was most expected, but few families stood to gain, since most already received free or reduced-price meals. Where the share was lowest, more families stood to gain, since most were paying for school meals, but CEP was least expected. Anticipation should therefore be strongest between the two, in the middle quartiles, where CEP was likely and many families still stood to gain. We split ever-CEP tracts by quartile of their 2010 free-and-reduced-price share. The pre-trend falls steadily across the quartiles. In any credit account, for example, it is 0.18 percentage points a quarter in the lowest quartile and 0.05 in the second, is absent in the third, and reverses to a decline of 0.12 in the highest. Across outcomes, significant pre-trends number 14 in the lowest quartile, 10 in the second, 3 in the third and 9 in the highest, where all nine are declines. The pre-trend is therefore largest where CEP was least expected, rather than in the middle quartiles, where anticipation should be strongest. The step at onset falls in the same order, as the heterogeneity by prior eligibility shows. The pre-trend in any severe delinquency is flat in all four quartiles, although in the second quartile several auto and revolving delinquency measures have significant pre-trends. Our results suggest that the pre-trend does not reflect anticipation based on how likely CEP was.

The test cannot rule out a response once a district had elected CEP, which we take up in Section 7.2.

We find that the jump in credit access at onset is robust to every check, while whether the effect stays at the same level relative to trend two years out depends on how the pre-trend is extrapolated and on the season. That is consistent with the dynamics in the main event studies, in which households enter credit markets when CEP arrives, add accounts over the following year and then level off.

7.2 Sensitivity

Rambachan and Roth (2023) bound the effect under restrictions on how far the post-adoption violation of parallel trends may depart from the pre-adoption one. We report two restrictions, relative magnitudes and smoothness, and then examine what the smoothness bound rests on, by comparing it with a straight line through the quarter before onset and with a line fitted through all the pre-onset quarters. Relative magnitudes, in Table A3, bounds the post-adoption violation by a multiple of the largest quarter-to-quarter change before adoption. It suits pre-trends that fluctuate, but it rewards smooth ones. When the pre-trend drifts steadily, its quarterly changes are small, the yardstick is small, and a given violation counts as a large multiple. Any credit account, the smoothest pre-trend, has the highest breakdown value under it, 1.25, shared with the number of accounts. Smoothness instead bounds how much the slope of the trend may change each quarter; at zero it is exact linear extrapolation of the pre-trend. It is the restriction designed for a steady drift, which is what the pre-periods show.

The Callaway-Sant'Anna bounds rest on its coefficients before onset, which show the same school-year cycle as the person-level event studies, in which credit outcomes rise each time a school year starts. Around a straight-line trend, about half of their movement before onset

follows the school year for the median outcome. The cycle does not bias the average effect after onset, because that average spans three full school years, so each school year's seasonal rise and fall offset each other. It does enlarge the movement before onset from which the relative-magnitudes bounds are scaled, so our results suggest those bounds are, if anything, conservative.

Under smoothness, 17 of the 31 outcomes survive a straight-line continuation of the pre-trend, shown in Online Appendix Figure OA12. They are all 14 aggregate and revolving outcomes, both mortgage access measures and the tract-level Equifax Risk Score, although the tract-level estimate for the Equifax Risk Score is a rise, which we set aside because its placebos exceed its effect. None of the auto outcomes, mortgage delinquency measures, bankruptcy or foreclosure survive. Among the survivors, the bend each tolerates is largest for any credit account and for revolving accounts in severe delinquency, at about a sixth of the largest quarter-to-quarter bend before onset, and smallest for delinquent balances and the tract-level Equifax Risk Score, at a fortieth.

These bounds use the Callaway-Sant'Anna event study on tract-quarter means. The same bounds on the person-level TWFE event study in Figure 2, for the average over the first nine quarters, give a broadly similar picture. Of the 31 outcomes, 16 survive a straight-line continuation of the pre-trend, including every revolving outcome, six of the seven aggregate outcomes, auto delinquent balances, the Equifax Risk Score and bankruptcy, although the bankruptcy event study reflects how that outcome is built, as the discussion of bankruptcy in Section 5.4 explains, so we do not read its survival as evidence. Five outcomes differ from the Callaway-Sant'Anna bounds. Any credit account and both mortgage access measures survive under Callaway-Sant'Anna but not under TWFE, and auto delinquent balances and bankruptcy survive under TWFE but not under Callaway-Sant'Anna. The other auto outcomes, the mortgage

delinquency measures and foreclosure fail under both. For any credit account, the average effect does not survive a straight line under the TWFE event study, although the step at onset and the detrended effect two years out do, as the contrasts in the event studies show. The share measures, flat in the dose specification, rise once the pre-trend is continued, the share of accounts in severe delinquency by 0.33 to 0.95 percentage points and the revolving shares by more. The person-level decline in the Equifax Risk Score survives, at 0.88 to 1.50 score points. The straight-line bound turns on a single quarter. Because the continued trend must pass through zero in the reference quarter, its slope is set by the change from two quarters before onset to the quarter just before it, and for 23 of the 31 outcomes under TWFE and 26 of the 31 under Callaway-Sant'Anna the reported interval matches carrying that one change forward. For credit access that change is a rise several times the typical quarterly change earlier in the pre-period, concentrated in the tracts with the lowest free-lunch shares, where the pre-trend is concentrated. That quarter, April to June, is when districts must elect CEP for the coming school year, so the rise may reflect a pre-trend that sped up or people responding once their district elected CEP, and the data cannot distinguish the two. In 2014–15 and 2015–16, however, the deadline was August 31, so many elections fell in the onset quarter itself. Either way, the step does not depend on that quarter, since the coefficient at onset, measured from the quarter before onset, is 0.84 percentage points for any credit account. A line fitted through all the pre-onset quarters, the assumption behind the detrended effects two years out reported earlier in this section, leaves 25 of the 31 outcomes significant under TWFE and 22 of the 31 under Callaway-Sant'Anna. Under TWFE the difference is any credit account, both mortgage access measures and six auto outcomes, and under Callaway-Sant'Anna it is the seven auto outcomes, while the two shares across all trade lines survive the straight-line bound but not the fitted line. The mortgage delinquency measures

and foreclosure fail under both lines. Which reading is right depends on whether the pre-trend sped up in the quarter before onset, which the data cannot settle. For auto credit, whose long pre-trend in access is flat, the fitted line leaves every access and delinquency measure significant under both estimators.

We do not claim the pre-trends are satisfactory. What we claim is narrower. The step in credit access at onset survives every check we apply. The step at onset in the Equifax Risk Score is a fall of less than a fifth of one percent of its mean. That decline holds in every person-level sample and specification but not under the tract-level estimators, where Callaway-Sant'Anna finds a rise of similar size and dCDH finds no significant change. We rely on the person-level estimates because the tract-level placebos exceed the effect for this outcome. Whether credit access stays above its trend two years after onset does not survive every check, nor do the auto and mortgage results, and we place less weight on them. The share measures are flat in the dose specification but rise once the pre-trend is taken into account, and our results suggest a modest worsening in repayment that does not appear among people with an unbroken credit file.

8. Conclusion

The Community Eligibility Provision removes a recurring food cost from households, without an application and without a means test at the household level, and it reaches many families who may receive nothing else from the safety net. The effects of schoolwide free meals on children have been studied extensively (Schwartz and Rothbart, 2020; Ruffini, 2022), but what the household does with money it no longer has to spend on food has not.

Our results suggest that exposure to CEP draws households further into credit markets, and it does so at onset. Access steps up when CEP arrives, the number of accounts builds over the following year, and the gains are most durable in revolving credit, the kind that lets

households smooth consumption. The composition of what households owe changes little. The share of accounts and the share of dollars in severe delinquency do not move in the dose specification, and rise modestly once the pre-adoption trend is taken into account or the staggered-robust estimators are used, a rise that our results suggest does not come from people observed throughout. Our results suggest that the rise in severe delinquency is modest, and that among people with an unbroken credit file it does not rise. The gains are largest where the fewest students were already eligible for free or reduced-price meals. The Equifax Risk Score, the standard measure of creditworthiness, barely moves. The dose estimate is a fall of about a seventh of one percent of its base, and in the event study it recovers about half of its largest decline within two years. In auto, the shares of accounts and of debt in severe delinquency rise in the main dose estimates, before any adjustment for the pre-adoption trend, where in the other trade lines they rise only once that trend is taken into account. The auto rise holds when the pre-trend is carried forward but not under the straight-line bound through the quarter before onset, and it does not appear among people with an unbroken credit file. That is not unexpected. Our results suggest that auto, unlike the other trade lines, is where the expansion comes mainly from people who did not hold an auto loan taking one out, rather than from existing borrowers adding loans. Our results suggest that CEP expands households' use of credit, especially revolving credit, while the Equifax Risk Score barely moves and repayment worsens modestly, though not among people with an unbroken credit file.

We read the expansion in credit holding as coming from the demand side, although we offer this as a conjecture. Lenders extending more credit because households had become better risks would show up as an improving Equifax Risk Score, and the score shows no improvement at any point. It falls at onset, reaches its lowest point two quarters later and recovers about half of

that decline within two years. A loosening of lending standards would not have to move the Equifax Risk Score, and ruling it out would require inquiries, originations by score band and limit changes on existing cards. We put weight on the absence of improvement rather than on the decline, because credit scores in general are influenced by the age of accounts and by utilization, and can therefore fall for reasons unrelated to creditworthiness. This suggests that households are taking up credit once the budget constraint loosens rather than being offered it.

The existing evidence on nutrition assistance and credit runs in one direction, with increased debt, delinquencies and lower credit scores associated with loss of benefits. In these cases, borrowing rises because resources were withdrawn and credit substitutes for the lost transfers. CEP runs the other way yet produces the same increase in borrowing. That symmetry suggests that substitution is not a general account of the relationship. If credit rose only when benefits were withdrawn, borrowing would be a distress signal; rising on both sides suggests instead that the transfer changes what a household is able to carry.

The health insurance literature is a further comparison. Studies of Medicaid and the Affordable Care Act expansions find reductions in medical debt and financial strain (Finkelstein et al., 2012; Hu et al., 2018; Brevoort et al., 2020; Miller et al., 2021). Insurance works by removing exposure to bills a household could not have anticipated or budgeted for, so much of the measured reduction likely follows from the liability never arising rather than from any change in household behavior. CEP removes a recurring cost that is known and already in the household budget. Nothing about uncertainty changes, and any effect on credit records must therefore run through what the household chooses to do with the freed resources. We therefore see the nutrition assistance results as the closer comparison. We also do not read our results as turning on the in-kind form of the transfer, which is inframarginal for households that were

already buying lunches (Hoynes and Schanzenbach, 2009). What distinguishes CEP is that it arrives daily throughout the school year rather than as a lump sum or a monthly payment, and that it reaches households without a means test.

The conventional pre-trend tests reject for nearly every outcome. The step in credit access at onset survives detrending, the comparison of adopting tracts only, a balanced sample and a sample beginning in 2014 Q2; the persistence of the level effect two years out does not survive all of them. Exposure is assigned at the tract level from the schools a tract's residents are most likely to attend, so we read the magnitudes as lower bounds on the effect for genuinely exposed households. We place less weight on three results for reasons other than the sensitivity analysis. They are mortgage access, because it does not hold when the ten pilot states and the District of Columbia are returned to the sample and the estimators reverse its sign, the tract-level Equifax Risk Score because its placebos exceed its effect, and bankruptcy because the absolute magnitude is small and the cost of filing gives a benign reading. Foreclosure survives neither smoothness bound and only the smallest relative-magnitudes bound.

The auto result is where interpretation is the most difficult. The credit panel observes trade lines rather than vehicles, so a new auto loan tells us that financing was acquired, not that transport was, and we cannot separate a first vehicle from a financed replacement of one owned outright. Either reading is consistent with the transportation channel that prior work identifies as consequential for employment[15], and both create a monthly obligation where there was none, which is a direct reason auto delinquency rises as it does. Because employment is not observed, we cannot test the employment effect and thus the welfare reading is genuinely ambiguous. A household that finances a vehicle and falls behind on the loan may be worse off or may have

[15] See Raphael and Rice, 2002; Ong, 2002; and Gurley and Bruce, 2005.

exchanged access to work and future earnings for current debt. They may have also changed their exposure to risk. An older vehicle that they owned outright carries no monthly payment but leaves the household exposed to unexpected repair costs.

Credit is not just a means of financing new purchases. It also lets a household choose when to spend, and the timing of food spending among constrained households is among the better documented facts in this literature. Consumption falls over the course of a benefit cycle, food insecurity rises toward its end, and households turn to charitable sources later in the month.[16] The same cycle governs prices paid, as Zaki and Todd (2023) find that SNAP recipients are most price-conscious and most successful at price-saving shortly after benefits arrive and note that some of those strategies require access to a lump sum. Buying in bulk, or buying on sale, means paying more today for a lower unit price, and that is precisely what a household short of liquidity cannot do (Orhun and Palazzolo, 2019). Zaki (2016) finds that access to short-term credit smooths consumption within the pay cycle directly. Our results suggest that this may be the most useful benefit of the program for household finances. The durable gain is in revolving credit, which lets a household choose when to spend, and it comes with little change in the Equifax Risk Score and no rise in delinquency among households already established in credit markets. Credit on those terms would let a household buy in bulk or on sale and spread the cost, freeing resources beyond the value of the meals themselves, a channel the credit records cannot observe directly.

School meals arrive on a cycle of their own. The benefit is available through the school year and absent for the summer, and the resources it frees are therefore not a smooth annual flow. A household drawn into revolving credit gains a way to carry costs across that gap, and

[16] See Shapiro, 2005; Todd, 2015; Gregory and Smith, 2019; and Byrne and Just, 2021.

more generally to buy when it is cheapest to buy rather than when the money happens to arrive. On this reading the rise in access and the rise in delinquency are not competing findings but two aspects of the same change, with people making more use of credit once resources free up.

References

Argys, Laura M., Andrew I. Friedson, M. Melinda Pitts, and D. Sebastian Tello-Trillo. 2020. "Losing Public Health Insurance: TennCare Reform and Personal Financial Distress." Journal of Public Economics 187: 104202.

Baker, Andrew C., David F. Larcker, and Charles C. Y. Wang. 2022. "How Much Should We Trust Staggered Difference-in-Differences Estimates?" Journal of Financial Economics 144 (2): 370-395.

Bhattacharya, Jayanta, Janet Currie, and Steven Haider. 2006. "Breakfast of Champions? The School Breakfast Program and the Nutrition of Children and Families." Journal of Human Resources 41 (3): 445-466.

Bitler, Marianne, and Hilary Hoynes. 2016. "The More Things Change, the More They Stay the Same? The Safety Net and Poverty in the Great Recession." Journal of Labor Economics 34 (S1): S403-S444.

Board of Governors of the Federal Reserve System. 2026. Economic Well-Being of U.S. Households in 2025. Washington, DC: Board of Governors of the Federal Reserve System.

Brevoort, Kenneth, Daniel Grodzicki, and Martin B. Hackmann. 2020. "The Credit Consequences of Unpaid Medical Bills." Journal of Public Economics 187: 104203.

Byrne, Anne T., and David R. Just. 2021. "The Other Half: An Examination of Monthly Food Pantry Cycles in the Context of SNAP Benefits." Applied Economic Perspectives and Policy 43 (2): 716–731.

Byrne, Anne T., David R. Just, and Christopher B. Barrett. 2023. "But It Came from a Food Pantry: Product Stigma and Quality Perceptions of Food Pantry Offerings." Agricultural Economics 54 (2): 327–344.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. 2024. "Difference-in-Differences with a Continuous Treatment." NBER Working Paper 32117.

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." Journal of Econometrics 225 (2): 200-230.

Center on Budget and Policy Priorities. 2015. "Take Up of Community Eligibility This School Year." Washington, DC: Center on Budget and Policy Priorities, February.

Cotti, Chad, John Gordanier, and Orgul Ozturk. 2021. "Does Distributing SNAP Benefits Later in the Month Smooth Expenditures?" Food Policy 104: 102123.

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

de Chaisemartin, Clément, and Xavier D'Haultfœuille. 2026. "Difference-in-Differences Estimators of Intertemporal Treatment Effects." Review of Economics and Statistics 108 (4): 863-880.

Dodini, Samuel, Jeff Larrimore, and Anna Tranfaglia. 2024. "Financial Repercussions of SNAP Work Requirements." Journal of Public Economics 229: 105034.

Equifax. How Are Credit Scores Calculated? https://www.equifax.com/personal/education/credit/score/articles/-/learn/how-is-credit-score-calculated/

Federal Reserve Bank of New York. Consumer Credit Panel/Equifax. Center for Microeconomic Data, Federal Reserve Bank of New York.

Federal Reserve Bank of New York. 2020. Quarterly Report on Household Debt and Credit, 2019:Q4. Center for Microeconomic Data.

Federal Reserve Bank of New York. 2026. Quarterly Report on Household Debt and Credit, 2026:Q2. Center for Microeconomic Data.

Finkelstein, Amy, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, Katherine Baicker, and the Oregon Health Study Group. 2012. "The Oregon Health Insurance Experiment: Evidence from the First Year." Quarterly Journal of Economics 127 (3): 1057-1106.

Fletcher, Jason M., and David E. Frisvold. 2017. "The Relationship Between the School Breakfast Program and Food Insecurity." Journal of Consumer Affairs 51 (3): 481-500.

Food Research & Action Center. 2020. Community Eligibility: The Key to Hunger-Free Schools, School Year 2019–2020. Washington, DC: Food Research & Action Center.

Fox, Mary Kay, and Elizabeth Gearan. 2019. School Nutrition and Meal Cost Study: Summary of Findings. Alexandria, VA: U.S. Department of Agriculture, Food and Nutrition Service.

Frisvold, David E. 2015. "Nutrition and Cognitive Achievement: An Evaluation of the School Breakfast Program." Journal of Public Economics 124: 91-104.

Fuller, Sarah Crittenden, and Aubrey N. Comperatore. 2020. The Effects of the Community Eligibility Provision on Educational Outcomes for Students in North Carolina. Policy brief, December. Chapel Hill, NC: Education Policy Initiative at Carolina.

Gleason, Philip M., and Carol W. Suitor. 2001. "Children's Diets in the Mid-1990s: Dietary Intake and Its Relationship with School Meal Participation." Special Nutrition Programs Report CN-01-CD1. Alexandria, VA: U.S. Department of Agriculture, Food and Nutrition Service.

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." Journal of Econometrics 225 (2): 254-277.

Goodman-Bacon, Andrew, and Leslie McGranahan. 2008. “How Do EITC Recipients Spend Their Refunds?” Economic Perspectives, Federal Reserve Bank of Chicago, 32 (2): 17–32.

Gordanier, John, Orgul Ozturk, Breyon Williams, and Crystal Zhan. 2020. "Free Lunch for All! The Effect of the Community Eligibility Provision on Academic Outcomes." Economics of Education Review 77: 101999.

Gordon, Nora, and Krista Ruffini. 2021. "Schoolwide Free Meals and Student Discipline: Effects of the Community Eligibility Provision." Education Finance and Policy 16 (3): 418-442.

Gregory, Christian A., and Travis A. Smith. 2019. “Salience, Food Security, and SNAP Receipt.” Journal of Policy Analysis and Management 38 (1): 124–154.

Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang. 2014. "Liquidity Constraints and Consumer Bankruptcy: Evidence from Tax Rebates." Review of Economics and Statistics 96 (3): 431-443.

Gundersen, Craig, Brent Kreider, and John Pepper. 2012. "The Impact of the National School Lunch Program on Child Health: A Nonparametric Bounds Analysis." Journal of Econometrics 166 (1): 79-91.

Gundersen, Craig, and James P. Ziliak. 2003. "The Role of Food Stamps in Consumption Stabilization." Journal of Human Resources 38 (Supplement): 1051-1079.

Gurley, Tami, and Donald Bruce. 2005. "The Effects of Car Access on Employment Outcomes for Welfare Recipients." Journal of Urban Economics 58 (2): 250-272.

Handbury, Jessie, and Sarah Moshary. 2021. "School Food Policy Affects Everyone: Retail Responses to the National School Lunch Program." NBER Working Paper 29384.

Homonoff, Tatiana, Min S. Lee, and Katherine Meckel. 2025. "The Financial Consequences of Being Denied Benefit Access." NBER Working Paper 34434.

Hoynes, Hilary W., and Diane Whitmore Schanzenbach. 2009. "Consumption Responses to In-Kind Transfers: Evidence from the Introduction of the Food Stamp Program." American Economic Journal: Applied Economics 1 (4): 109-139.

Hu, Luojia, Robert Kaestner, Bhashkar Mazumder, Sarah Miller, and Ashley Wong. 2018. "The Effect of the Affordable Care Act Medicaid Expansions on Financial Wellbeing." Journal of Public Economics 163: 99-112.

Imbens, Guido W., and Donald B. Rubin. 2015. Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction. Cambridge: Cambridge University Press.

Imbens, Guido W., and Jeffrey M. Wooldridge. 2009. “Recent Developments in the Econometrics of Program Evaluation.” Journal of Economic Literature 47 (1): 5–86.

Kho, Adam, and Seth B. Hunter. 2026. “Free Meals for All: The Effects of Tennessee’s Community Eligibility Provision on Student Non-Academic Outcomes.” Educational Studies 52 (2): 129-147.

Manson, Steven, Jonathan Schroeder, David Van Riper, and Steven Ruggles. 2019. IPUMS National Historical Geographic Information System: Version 14.0 [dataset]. Minneapolis, MN: IPUMS. http://doi.org/10.18128/D050.V14.0

Marcus, Michelle, and Katherine G. Yewell. 2022. "The Effect of Free School Meals on Household Food Purchases: Evidence from the Community Eligibility Provision." Journal of Health Economics 84: 102646.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." Quarterly Journal of Economics 128 (4): 1687-1726.

Miller, Sarah, Luojia Hu, Robert Kaestner, Bhashkar Mazumder, and Ashley Wong. 2021. "The ACA Medicaid Expansion in Michigan and Financial Health." Journal of Policy Analysis and Management 40 (2): 348–375.

Miller, Sarah, and Cindy K. Soo. 2021. “Do Neighborhoods Affect the Credit Market Decisions of Low-Income Borrowers? Evidence from the Moving to Opportunity Experiment.” Review of Financial Studies 34 (2): 827–863.

National Cancer Institute, Surveillance, Epidemiology, and End Results Program. U.S. Population Data.

National Center for Education Statistics. 2023. Digest of Education Statistics 2023, Table 203.10. Washington, DC: U.S. Department of Education, Institute of Education Sciences.

Ong, Paul M. 2002. "Car Ownership and Welfare-to-Work." Journal of Policy Analysis and Management 21 (2): 239-252.

Orhun, A. Yeşim, and Mike Palazzolo. 2019. "Frugality Is Hard to Afford." Journal of Marketing Research 56 (1): 1-17.

Rambachan, Ashesh, and Jonathan Roth. 2023. "A More Credible Approach to Parallel Trends." Review of Economic Studies 90 (5): 2555-2591.

Raphael, Steven, and Lorien Rice. 2002. "Car Ownership, Employment, and Earnings." Journal of Urban Economics 52 (1): 109-130.

Ruffini, Krista. 2022. "Universal Access to Free School Meals and Student Achievement: Evidence from the Community Eligibility Provision." Journal of Human Resources 57 (3): 776-820.

Ruffini, Krista, Orgül Öztürk, and Pelin Pekgün. 2025. "In-Kind Government Assistance and Crowd-Out of Charitable Services: Evidence from Free School Meals." Journal of Public Economics 248: 105391.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod. 2012. “Check in the Mail or More in the Paycheck: Does the Effectiveness of Fiscal Stimulus Depend on How It Is Delivered?” American Economic Journal: Economic Policy 4 (3): 216–50.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod. 2015. “Balance-Sheet Households and Fiscal Stimulus: Lessons from the Payroll Tax Cut and Its Expiration.” NBER Working Paper 21220.

Schwartz, Amy Ellen, and Michah W. Rothbart. 2020. “Let Them Eat Lunch: The Impact of Universal Free Meals on Student Performance.” Journal of Policy Analysis and Management 39 (2): 376–410.

Shapiro, Jesse M. 2005. "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle." Journal of Public Economics 89 (2–3): 303–325.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." Journal of Econometrics 225 (2): 175-199.

Todd, Jessica E. 2015. "Revisiting the Supplemental Nutrition Assistance Program Cycle of Food Intake: Investigating Heterogeneity, Diet Quality, and a Large Boost in Benefit Amounts." Applied Economic Perspectives and Policy 37 (3): 437-458.

U.S. Code of Federal Regulations. Title 7, Section 245.9: Special assistance certification and reimbursement alternatives.

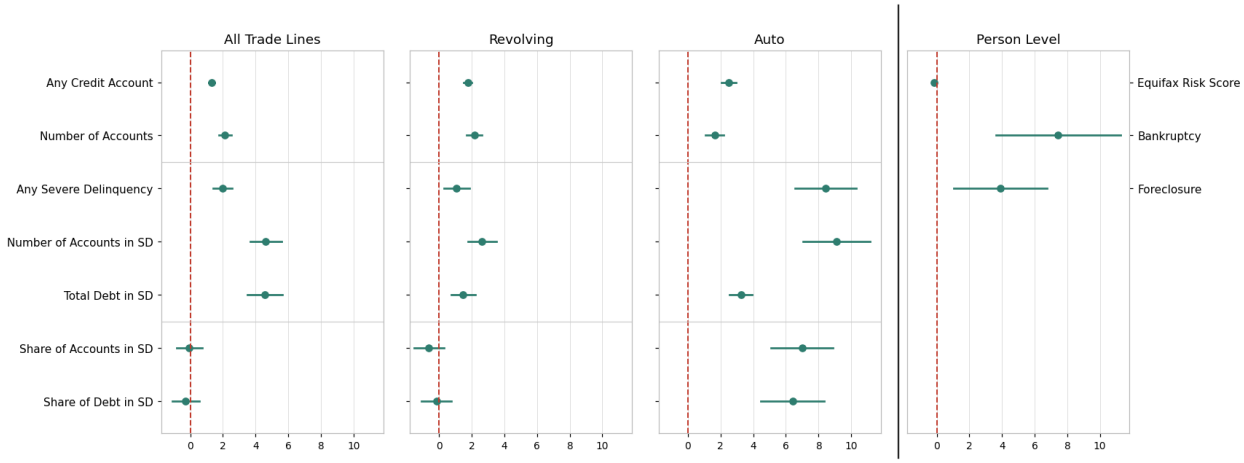
U.S. Department of Agriculture, Food and Nutrition Service. 2015. “Extension of the Deadline for Local Educational Agencies to Elect the Community Eligibility Provision for School Year 2015–16.” Memorandum SP 25-2015, March 4.

U.S. Department of Agriculture, Food and Nutrition Service. 2016. “Extension of the Deadline for Local Educational Agencies to Elect the Community Eligibility Provision for School Year 2016–17.” Memorandum SP 30-2016, March 10.

Zaki, Mary. 2016. "Access to Short-Term Credit and Consumption Smoothing Within the Paycycle." Fondazione Eni Enrico Mattei Working Paper 2016.07.

Zaki, Mary, and Jessica E. Todd. 2023. "Price Consciousness at the Peak of 'Impatience'." Journal of Human Resources 58 (3): 1003-1027.

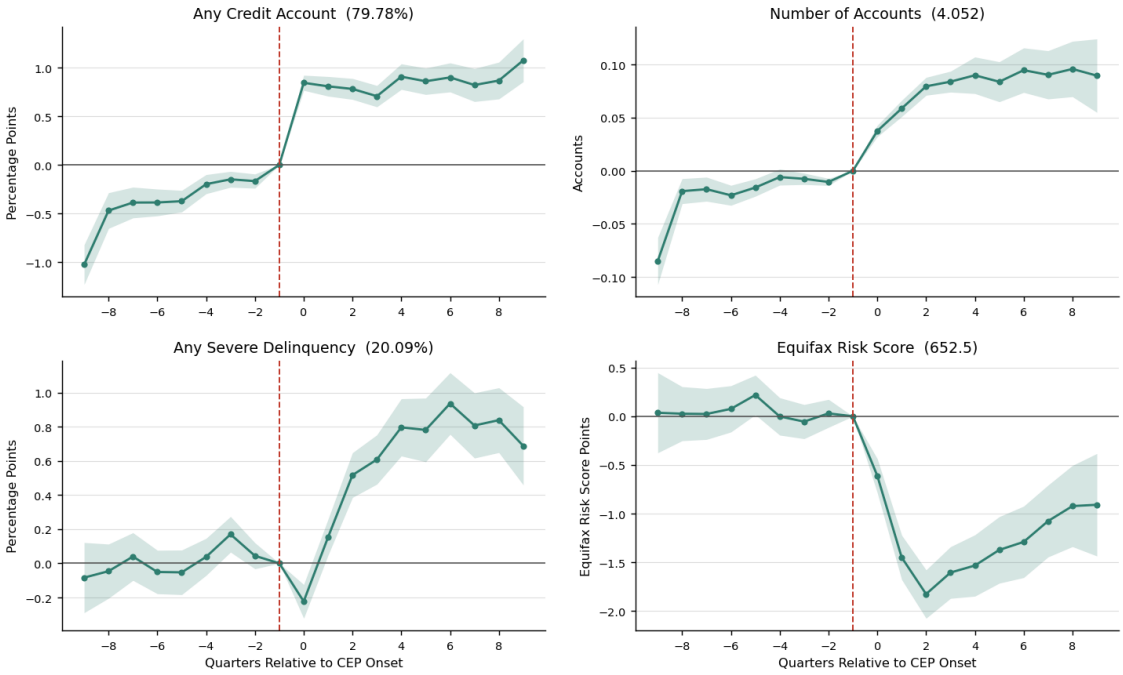
Figure 1. Marginal Effects as a Percentage of Pre-Treatment Baseline



Notes. Combined marginal effect at realized exposure, with 95 percent confidence intervals, expressed as a percentage of each outcome's own pre-treatment baseline. The baseline is the mean among ever-CEP persons in quarters before their own onset. Log-dollar outcomes are converted as exp(b) − 1. Mortgage outcomes are in Appendix Figure A1. Coefficients and significance for all 31 outcomes are in Appendix Table A1.

Source. Federal Reserve Bank of New York Consumer Credit Panel/Equifax; school-level CEP participation data from Ruffini (2022).

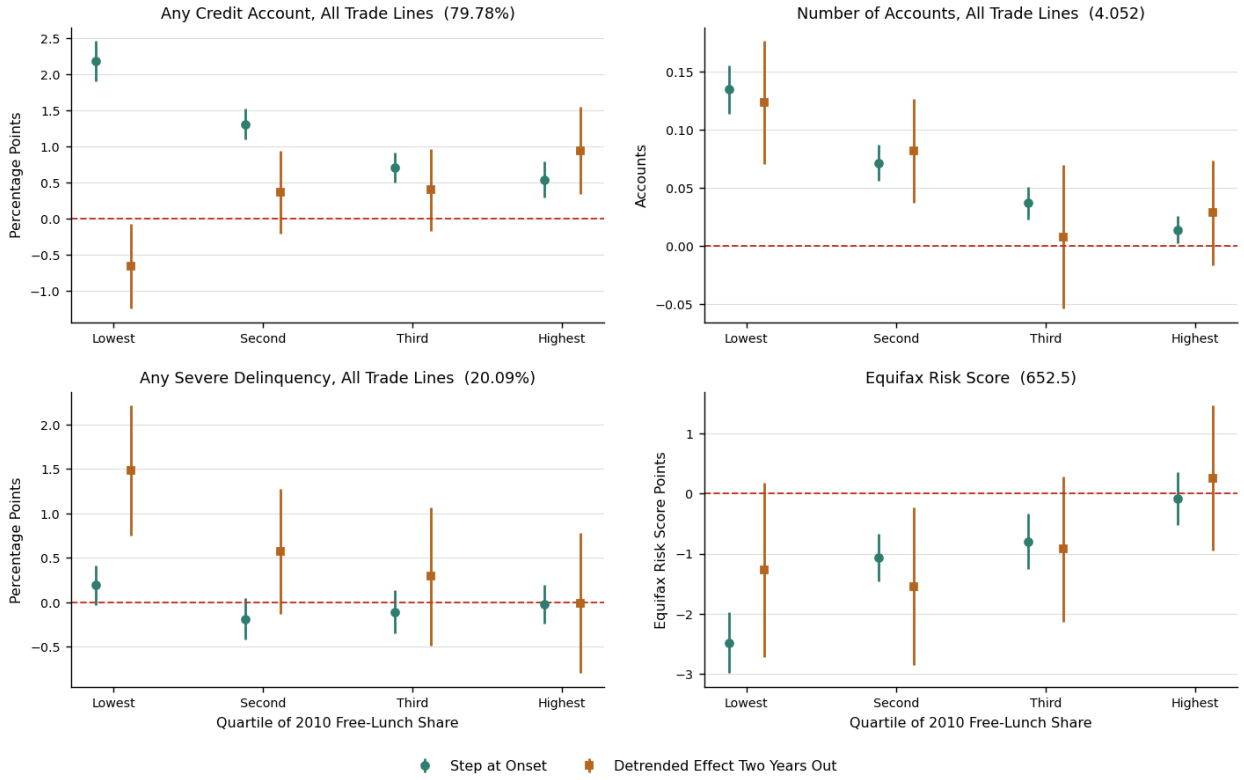
Figure 2. Event Study, Headline Outcomes



Notes. Event-study coefficients with 95 percent confidence intervals shaded. The dashed line marks the omitted reference quarter. Event time is binned at the end points, nine quarters before and after onset, so the first and last points pool every earlier and later quarter; coverage is held at its onset value, and each outcome's pre-treatment baseline is in parentheses. Event studies for the remaining outcomes are in Online Appendix Figure OA8.

Source. Federal Reserve Bank of New York Consumer Credit Panel/Equifax; school-level CEP participation data from Ruffini (2022).

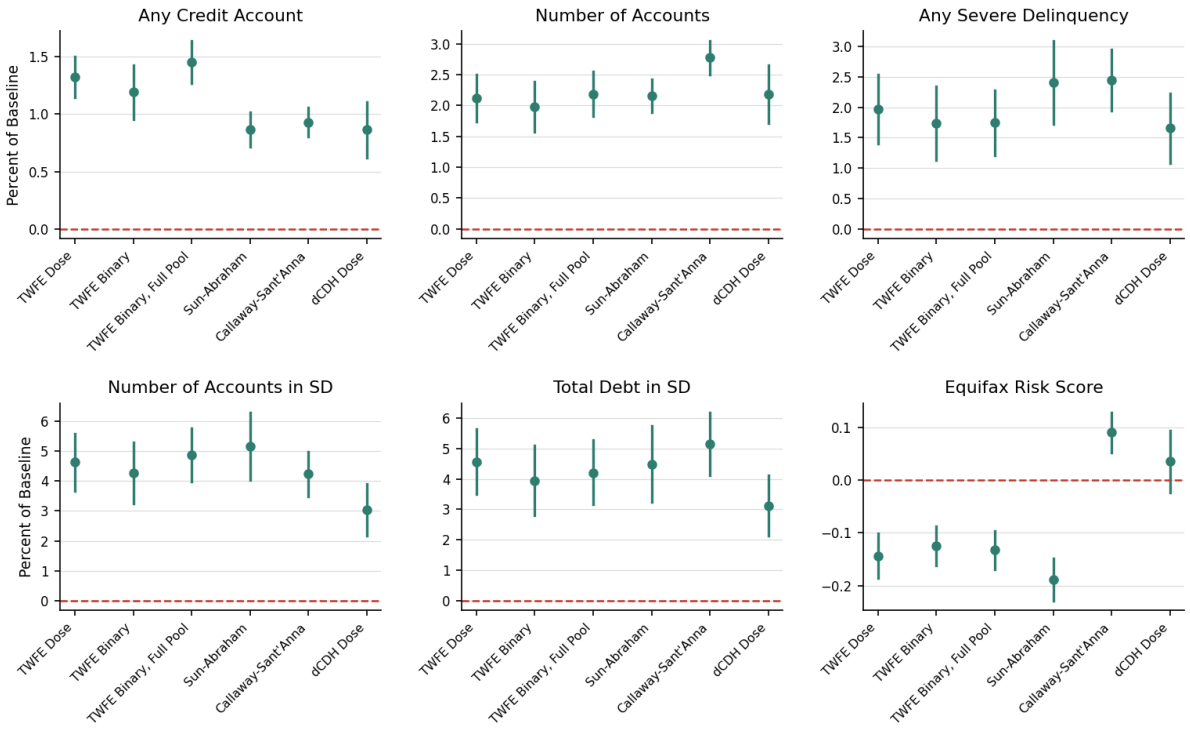
Figure 3. Step at Onset and Detrended Effect by Prior Eligibility, Headline Outcomes



Notes. Event studies estimated separately for ever-CEP individuals in each quartile of their tract's 2010 free-and-reduced-price share, from the lowest to the highest. The step at onset is the change from two quarters before onset to the onset quarter. The detrended effect is the effect two years after onset net of each quartile's pre-onset slope carried forward. Bars are 95 percent confidence intervals. Each vertical axis carries the outcome's own units and the figure in parentheses after each panel title is its pre-treatment baseline.

Source. Federal Reserve Bank of New York Consumer Credit Panel/Equifax; school-level CEP participation data from Ruffini (2022); 2010 free-and-reduced-price shares from the National Center for Education Statistics Common Core of Data, via Ruffini (2022).

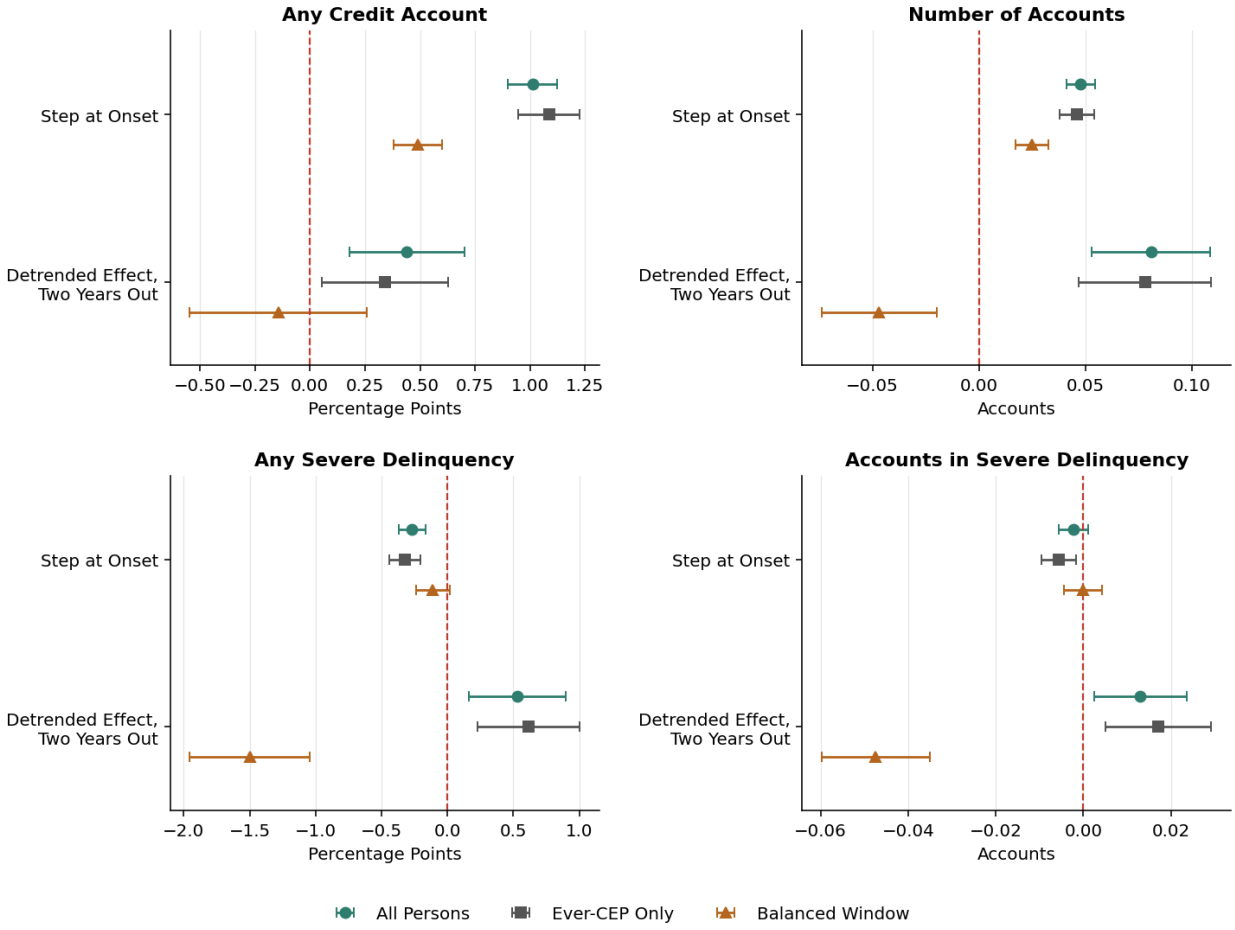
Figure 4. Estimator Comparison, Selected Outcomes



Notes. Each panel reports the same outcome estimated six ways, expressed as a percentage of that outcome's own pre-treatment baseline with 95 percent confidence intervals. TWFE dose is the combined marginal effect at realized exposure; TWFE binary replaces the dose with an absorbing indicator; the full pool row returns the ten pilot states and the District of Columbia to the sample. Sun-Abraham is estimated at the person level, Callaway-Sant'Anna and de Chaisemartin-D'Haultfœuille on tract-quarter means weighted by person count. Sun-Abraham and Callaway-Sant'Anna are averages over the nine quarters from onset; the dCDH figure is the mean of its horizon effects over horizons one through eight, with a delta-method standard error from its own covariance matrix. Log-dollar outcomes are converted as exp(b) − 1. Each panel carries its own vertical scale; the Equifax Risk Score effect is a fraction of one percent of a base near 652 and would be invisible on a shared scale.

Source. Federal Reserve Bank of New York Consumer Credit Panel/Equifax; school-level CEP participation data from Ruffini (2022).

Figure 5. Step at Onset and Detrended Effect, Selected Outcomes



Notes. Contrasts from the dose event study, with 95 percent confidence intervals. The step at onset is the coefficient at k = 0 less the coefficient at k = -2. The detrended effect is the coefficient at k = 8 less the pre-trend slope, fitted over k = -8 to -2, extrapolated forward to k = 8. Both are linear combinations of the event-study coefficients, with standard errors from their joint covariance. All persons is the main specification; ever-CEP only compares CEP tracts with each other; the balanced window keeps people observed in every quarter from k = -8 to 8. Any credit account and any severe delinquency are in percentage points; the counts are in accounts.

Source. Federal Reserve Bank of New York Consumer Credit Panel/Equifax; school-level CEP participation data from Ruffini (2022).

Table 1. Baseline Characteristics

Variables	Full: Ever CEP	Full: Never CEP	Main 2014: Ever CEP	Main 2014: Never CEP
Treatment				
CEP Coverage at Onset	0.6726	–	0.6683	–
	(0.2912)		(0.2911)	
CEP Gain	0.0757	–	0.0585	–
	(0.1778)		(0.1554)	
CEP Contemporaneous	0.6702	–	0.6562	–
	(0.3434)		(0.3374)	
CEP Binary	0.4273	–	0.3967	–
	(0.4947)		(0.4892)	
Demographics				
Age	35.2	38.3	35.0	38.7
	(8.1)	(10.1)	(8.0)	(10.3)
Tract Median Income Decile	4.9	5.5	4.8	5.4
	(2.6)	(2.5)	(2.6)	(2.5)
Share of Households with Children	62.7	62.5	62.4	62.2
with a Child Aged 6 to 17 (%)	(9.0)	(8.5)	(8.9)	(8.4)
Access				
Has Any Credit Account	0.8070	0.8383	0.8107	0.8369
	(0.3815)	(0.3569)	(0.3784)	(0.3580)
Revolving	0.6787	0.7344	0.6823	0.7325
	(0.4522)	(0.4289)	(0.4505)	(0.4298)
Auto	0.3387	0.3515	0.3422	0.3504
	(0.4567)	(0.4597)	(0.4578)	(0.4593)
Mortgage	0.2552	0.3290	0.2506	0.3312
	(0.4265)	(0.4591)	(0.4234)	(0.4596)
Number of Accounts Held	3.8756	4.2835	3.8199	4.2426
	(3.9676)	(4.0663)	(3.8972)	(4.0228)
Revolving	2.3597	2.7645	2.3409	2.7436
	(2.8756)	(3.0879)	(2.8406)	(3.0696)
Auto	0.4070	0.4297	0.4142	0.4312
	(0.6110)	(0.6321)	(0.6190)	(0.6377)
Mortgage	0.3384	0.4544	0.3295	0.4562
	(0.6656)	(0.7673)	(0.6554)	(0.7680)
Delinquency				
Any Severe Delinquency	0.2074	0.1600	0.2129	0.1619
	(0.3788)	(0.3427)	(0.3827)	(0.3443)
Revolving	0.1374	0.1087	0.1433	0.1111
	(0.3200)	(0.2895)	(0.3259)	(0.2923)
Auto	0.0394	0.0285	0.0412	0.0290
	(0.1841)	(0.1572)	(0.1881)	(0.1586)
Mortgage	0.0187	0.0190	0.0182	0.0185
	(0.1210)	(0.1216)	(0.1188)	(0.1196)
Number of Accounts in SD	0.4202	0.3252	0.4319	0.3282
	(1.0862)	(0.9647)	(1.0911)	(0.9624)

Variables	Full: Ever CEP	Full: Never CEP	Main 2014: Ever CEP	Main 2014: Never CEP
Revolving	0.2254	0.1882	0.2393	0.1947
	(0.6874)	(0.6598)	(0.7162)	(0.6758)
Auto	0.0426	0.0307	0.0447	0.0313
	(0.2074)	(0.1763)	(0.2128)	(0.1778)
Mortgage	0.0218	0.0226	0.0213	0.0220
	(0.1551)	(0.1618)	(0.1523)	(0.1576)
Share of Accounts in SD	0.1376	0.1024	0.1406	0.1034
	(0.2888)	(0.2532)	(0.2905)	(0.2540)
Revolving	0.1076	0.0824	0.1115	0.0840
	(0.2710)	(0.2392)	(0.2747)	(0.2411)
Auto	0.0352	0.0253	0.0366	0.0257
	(0.1688)	(0.1436)	(0.1717)	(0.1447)
Mortgage	0.0168	0.0167	0.0163	0.0163
	(0.1117)	(0.1108)	(0.1093)	(0.1086)
Total Debt in SD, Log	1.6983	1.3406	1.7376	1.3521
	(3.2229)	(2.9794)	(3.2425)	(2.9821)
Revolving	0.9833	0.8031	1.0267	0.8221
	(2.3706)	(2.2144)	(2.4146)	(2.2376)
Auto	0.3403	0.2451	0.3563	0.2496
	(1.6030)	(1.3644)	(1.6394)	(1.3765)
Mortgage	0.2160	0.2220	0.2106	0.2166
	(1.4072)	(1.4338)	(1.3827)	(1.4095)
Share of Debt in SD	0.1343	0.1015	0.1366	0.1021
	(0.2981)	(0.2640)	(0.2998)	(0.2644)
Revolving	0.1145	0.0895	0.1189	0.0914
	(0.2858)	(0.2560)	(0.2901)	(0.2582)
Auto	0.0342	0.0246	0.0356	0.0250
	(0.1676)	(0.1426)	(0.1705)	(0.1437)
Mortgage	0.0170	0.0170	0.0165	0.0165
	(0.1129)	(0.1125)	(0.1103)	(0.1102)
Creditworthiness and Catastrophic Events				
Equifax Risk Score	638.5104	667.7865	636.6304	668.1209
	(102.7272)	(103.2532)	(102.2215)	(103.4752)
Bankruptcy (First Filing)	0.2174	0.2278	0.3344	0.3130
(Percentage Points)	(2.9695)	(3.0415)	(3.6757)	(3.5610)
Foreclosure (First Filing)	0.1365	0.1451	0.1362	0.1469
(Percentage Points)	(1.9352)	(1.9847)	(1.9448)	(2.0003)
Sample Size				
Persons	1,517,103	2,154,198	944,602	1,532,820
Person-quarters	41,615,075	47,292,167	25,649,637	32,386,735
Quarters per Person, Mean	27.4	22.0	27.2	21.1

Notes. Standard deviations beneath means. Outcome and covariate rows are each person's mean over the quarters of their first year in the sample, averaged across persons. Treatment rows are computed over post-onset quarters for ever-CEP persons only; CEP binary is the share of all ever-CEP person-quarters that fall after onset. Bankruptcy and foreclosure are first-filing indicators, reported in percentage points. The sample is the estimation sample.
Source. Federal Reserve Bank of New York Consumer Credit Panel/Equifax; school-level CEP participation data from Ruffini (2022).

Table 2. Alternative Treatment Parameterizations

Panel A: Access and Delinquency, by Trade Line

	Any Credit Account	Number of Accounts	Any SD	Number of Accounts in SD	Total Debt in SD	Share of Accounts in SD	Share of Debt in SD
			Aggregate				
Dose (Main)	1.056	0.0860	0.396	0.0187	0.0446	-0.007	-0.036
	(0.071)	(0.0079)	(0.058)	(0.0020)	(0.0052)	(0.052)	(0.055)
Absorbing	0.968	0.0799	0.367	0.0174	0.0410	0.002	-0.023
	(0.069)	(0.0075)	(0.057)	(0.0019)	(0.0050)	(0.050)	(0.054)
Contemporaneous	0.204	0.0209	0.259	0.0137	0.0328	0.043	0.030
	(0.067)	(0.0079)	(0.048)	(0.0018)	(0.0044)	(0.042)	(0.045)
Binary	0.950	0.0800	0.349	0.0172	0.0387	-0.036	-0.060
	(0.095)	(0.0083)	(0.062)	(0.0021)	(0.0056)	(0.053)	(0.056)
			Revolving				
Dose (Main)	1.183	0.0550	0.140	0.0056	0.0145	-0.063	-0.016
	(0.086)	(0.0060)	(0.052)	(0.0010)	(0.0037)	(0.048)	(0.050)
Absorbing	1.089	0.0507	0.120	0.0051	0.0127	-0.058	-0.017
	(0.081)	(0.0058)	(0.050)	(0.0009)	(0.0035)	(0.046)	(0.049)
Contemporaneous	0.055	0.0039	-0.044	0.0008	0.0003	-0.123	-0.104
	(0.079)	(0.0058)	(0.042)	(0.0008)	(0.0030)	(0.039)	(0.040)
Binary	1.142	0.0540	0.161	0.0057	0.0154	-0.048	-0.004
	(0.104)	(0.0061)	(0.053)	(0.0010)	(0.0038)	(0.047)	(0.050)
			Auto				
Dose (Main)	0.882	0.0072	0.338	0.0040	0.0320	0.248	0.221
	(0.079)	(0.0013)	(0.038)	(0.0005)	(0.0034)	(0.035)	(0.034)
Absorbing	0.794	0.0063	0.306	0.0036	0.0290	0.225	0.200
	(0.076)	(0.0012)	(0.037)	(0.0004)	(0.0033)	(0.033)	(0.033)
Contemporaneous	0.237	-0.0003	0.254	0.0030	0.0238	0.195	0.178
	(0.067)	(0.0011)	(0.036)	(0.0004)	(0.0032)	(0.032)	(0.032)
Binary	0.781	0.0052	0.249	0.0029	0.0238	0.176	0.153
	(0.081)	(0.0012)	(0.037)	(0.0004)	(0.0033)	(0.034)	(0.033)

Panel A: Access and Delinquency, by Trade Line

	Any Credit Account	Number of Accounts	Any SD	Number of Accounts in SD	Total Debt in SD	Share of Accounts in SD	Share of Debt in SD
				Mortgage			
Dose (Main)	-0.264	-0.0034	0.004	0.0002	0.0011	0.000	0.001
	(0.057)	(0.0008)	(0.017)	(0.0002)	(0.0018)	(0.015)	(0.015)
Absorbing	-0.211	-0.0027	0.002	0.0001	0.0009	-0.000	0.001
	(0.054)	(0.0008)	(0.016)	(0.0002)	(0.0017)	(0.014)	(0.014)
Contemporaneous	-0.906	-0.0105	0.017	0.0003	0.0024	0.014	0.015
	(0.055)	(0.0007)	(0.014)	(0.0002)	(0.0015)	(0.013)	(0.013)
Binary	-0.147	-0.0022	0.000	0.0001	0.0007	-0.002	-0.001
	(0.059)	(0.0008)	(0.017)	(0.0002)	(0.0020)	(0.015)	(0.016)

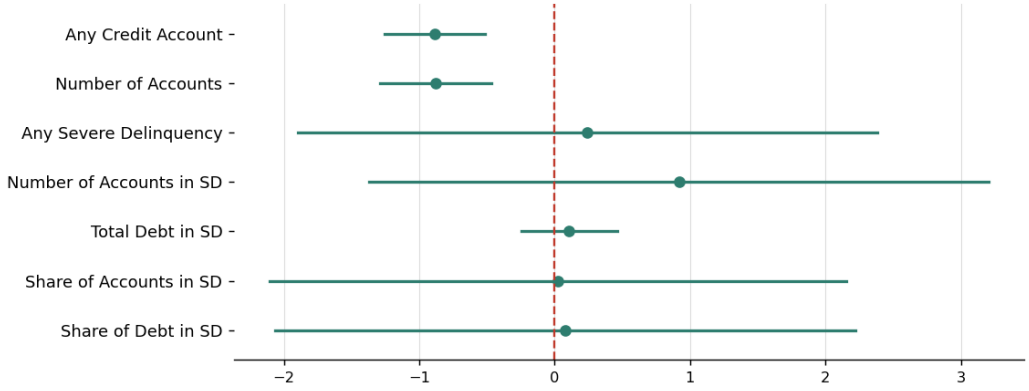
Panel B: Creditworthiness and Catastrophic Events

	Equifax Risk Score	Bankruptcy, Percentage Points	Foreclosure, Percentage Points
Dose (Main)	-0.9376	0.0184	0.0043
	(0.1407)	(0.0048)	(0.0016)
Absorbing	-0.8328	0.0177	0.0040
	(0.1358)	(0.0046)	(0.0016)
Contemporaneous	-0.5866	0.0103	0.0036
	(0.1165)	(0.0042)	(0.0014)
Binary	-0.8141	0.0178	0.0052
	(0.1248)	(0.0044)	(0.0017)

Notes. Each cell is the effect of the treatment definition named in the row on the outcome named in the column, with the standard error beneath. Dose (main) is the combined marginal effect at realized exposure, 0.6683 times the onset coefficient plus 0.0585 times the gain coefficient. Absorbing and Contemporaneous are rescaled to the same realized exposure, by 0.6683 and 0.6562 respectively, so all four rows are comparable. Binary is the effect of a tract adopting at all and needs no rescaling. Total Debt in SD is the log of one plus dollars, so the entry is a change in logs. The Equifax Risk Score, bankruptcy and foreclosure are properties of the person rather than of an account and have no trade-line split. All specifications carry person and state-by-quarter fixed effects on the main 2014 sample of 58,036,372 person-quarters, with standard errors clustered on county. The Equifax Risk Score is estimated on the 52,302,739 person-quarters with a score. Foreclosure is estimated on 56,958,561 person-quarters, because people with a foreclosure on record when they enter are excluded. Probability and share outcomes are in percentage points.
Source. Federal Reserve Bank of New York Consumer Credit Panel/Equifax; school-level CEP participation data from Ruffini (2022).

Appendix

Figure A1. Marginal Effects, Mortgage



Notes. Combined marginal effect at realized exposure, with 95 percent confidence intervals, expressed as a percentage of each outcome's own pre-treatment baseline. Mortgage outcomes only; the other 24 appear in Figure 1. Standard errors are exact, computed from the covariance between the two dose terms.

Source. Federal Reserve Bank of New York Consumer Credit Panel/Equifax; school-level CEP participation data from Ruffini (2022).

Table A1. Headline Dose Coefficients

	Onset Coverage	Coverage Gained	Baseline Mean
Access			
Aggregate			
Any Credit Account	1.427	1.750	79.78
	(0.102)	(0.277)	
Number of Accounts	0.1180	0.1221	4.0517
	(0.0111)	(0.0236)	
Revolving			
Any Credit Account	1.605	1.894	67.51
	(0.119)	(0.341)	
Number of Accounts	0.0747	0.0866	2.5338
	(0.0086)	(0.0166)	
Auto			
Any Credit Account	1.165	1.769	35.03
	(0.113)	(0.278)	
Number of Accounts	0.0092	0.0171	0.4358
	(0.0018)	(0.0039)	
Mortgage			
Any Credit Account	-0.302	-1.060	29.89
	(0.082)	(0.216)	
Number of Accounts	-0.0039	-0.0132	0.3851
	(0.0011)	(0.0032)	
Delinquencies			
Aggregate			
Any SD	0.542	0.588	20.09
	(0.085)	(0.233)	
Number of Accounts in SD	0.0256	0.0268	0.4042
	(0.0028)	(0.0079)	
Total Debt in SD, Log	0.0605	0.0711	3,933
	(0.0075)	(0.0201)	
Share of Accounts in SD	0.006	-0.197	13.42
	(0.075)	(0.177)	
Share of Debt in SD	-0.031	-0.262	12.84
	(0.080)	(0.181)	
Revolving			
Any SD	0.175	0.386	13.09
	(0.075)	(0.195)	
Number of Accounts in SD	0.0075	0.0100	0.2123
	(0.0013)	(0.0042)	
Total Debt in SD, Log	0.0185	0.0372	500
	(0.0053)	(0.0143)	
Share of Accounts in SD	-0.086	-0.094	10.17
	(0.069)	(0.155)	
Share of Debt in SD	-0.026	0.016	10.80
	(0.073)	(0.162)	
Auto			

	Onset Coverage	Coverage Gained	Baseline Mean
Any SD	0.450	0.645	4.01
	(0.055)	(0.123)	
Number of Accounts in SD	0.0053	0.0070	0.0436
	(0.0007)	(0.0015)	
Total Debt in SD, Log	0.0425	0.0613	328
	(0.0049)	(0.0111)	
Share of Accounts in SD	0.331	0.462	3.55
	(0.050)	(0.113)	
Share of Debt in SD	0.294	0.411	3.44
	(0.049)	(0.114)	
Mortgage			
Any SD	0.003	0.026	1.52
	(0.024)	(0.046)	
Number of Accounts in SD	0.0002	0.0005	0.0169
	(0.0003)	(0.0005)	
Total Debt in SD, Log	0.0013	0.0035	2,357
	(0.0026)	(0.0051)	
Share of Accounts in SD	-0.000	0.008	1.37
	(0.021)	(0.042)	
Share of Debt in SD	0.001	0.010	1.36
	(0.021)	(0.042)	
Creditworthiness and Catastrophic Events			
Equifax Risk Score	-1.2179	-2.1146	652.4703
	(0.2005)	(0.4580)	
Bankruptcy, Percentage Points	0.0263	0.0142	0.2476
	(0.0068)	(0.0107)	
Foreclosure, Percentage Points	0.0059	0.0051	0.1092
	(0.0023)	(0.0061)	
Person-quarters	58,036,372	58,036,372	15,481,242

Notes. Restricted sample. Probability and share outcomes are in percentage points, with baselines in percent. Standard errors clustered on county in parentheses. Onset coverage is the share covered in the quarter a tract is first covered, held forward; Coverage gained is any increase above that level. The marginal effect reported in the text is 0.6683 times the onset coefficient plus 0.0585 times the gain coefficient. Total Debt in SD is the log of one plus delinquent dollars and its baseline is the mean delinquent balance in 2014 dollars. Bankruptcy and foreclosure are first-filing indicators with baselines in percent. The Equifax Risk Score is estimated on the 52,302,739 person-quarters with a score. Foreclosure is estimated on 56,958,561 person-quarters, because people with a foreclosure on record when they enter are excluded.
Source. Federal Reserve Bank of New York Consumer Credit Panel/Equifax; school-level CEP participation data from Ruffini (2022).

Table A2. Goodman-Bacon Decomposition, by Outcome

	Full: Pooled	Full: Change (%)	Full: Forbidden	Restricted: Pooled	Restricted: Change (%)	Restricted: Forbidden
Access						
Aggregate						
Any Credit Account	1.506	9.93	-0.848	1.564	4.05	-0.215
Number of Accounts	0.2308	7.19	-0.0305	0.2377	3.64	-0.0052
Revolving						
Any Credit Account	1.196	11.05	-0.885	1.366	4.55	-0.379
Number of Accounts	0.1553	7.30	-0.0232	0.1635	3.79	-0.0106
Auto						
Any Credit Account	0.213	9.46	-0.104	0.420	2.30	0.149
Number of Accounts	0.0038	5.64	0.0004	0.0065	2.36	0.0022
Mortgage						
Any Credit Account	1.948	7.01	-0.205	1.883	3.80	-0.125
Number of Accounts	0.0420	7.04	-0.0046	0.0407	3.86	-0.0035
Delinquencies						
Aggregate						
Any SD	1.065	7.37	-0.172	1.053	3.29	0.079
Number of Accounts in SD	0.0480	7.88	-0.0116	0.0442	4.04	-0.0059
Total Debt in SD, Log	0.1263	7.50	-0.0229	0.1255	3.59	-0.0010
Share of Accounts in SD	0.514	7.49	-0.093	0.491	3.02	0.074
Share of Debt in SD	0.578	7.70	-0.123	0.561	3.18	0.060
Revolving						
Any SD	0.099	8.26	-0.030	0.152	3.16	0.017
Number of Accounts in SD	0.0065	7.69	-0.0014	0.0067	3.64	-0.0002
Total Debt in SD, Log	0.0197	7.66	-0.0041	0.0230	3.70	-0.0009
Share of Accounts in SD	-0.091	7.71	0.020	-0.038	7.71	0.044
Share of Debt in SD	-0.007	7.72	0.001	0.044	2.05	0.019
Auto						
Any SD	0.496	10.20	-0.302	0.547	4.07	-0.079
Number of Accounts in SD	0.0056	10.43	-0.0036	0.0063	4.05	-0.0009
Total Debt in SD, Log	0.0464	10.06	-0.0271	0.0521	4.06	-0.0073
Share of Accounts in SD	0.395	10.69	-0.270	0.438	4.04	-0.060
Share of Debt in SD	0.364	11.01	-0.267	0.405	4.01	-0.051
Mortgage						
Any SD	0.212	6.49	-0.005	0.218	3.51	0.003
Number of Accounts in SD	0.0030	6.62	-0.0001	0.0030	3.59	-0.00002

	Full: Pooled	Full: Change (%)	Full: Forbidden	Restricted: Pooled	Restricted: Change (%)	Restricted: Forbidden
Total Debt in SD, Log	0.0286	6.57	-0.0010	0.0289	3.55	0.0001
Share of Accounts in SD	0.160	6.36	-0.000	0.169	3.36	0.009
Share of Debt in SD	0.166	6.28	0.002	0.176	3.28	0.014
Creditworthiness and Catastrophic Events						
Equifax Risk Score	1.1561	7.65	-0.2362	1.2355	3.19	0.1275
Bankruptcy (Percentage Points)	0.0322	6.54	-0.0010	0.0295	3.46	0.0009
Foreclosure (Percentage Points)	0.0170	6.64	-0.0008	0.0203	3.22	0.0020
Decomposition Weights and Summary						
Weight: Never vs Ever	83.96%			87.56%		
Weight: Early vs Late	10.07%			9.00%		
Weight: Forbidden	5.97%			3.44%		
2x2 Components	57			21		
Opposite-signed, of 31	29			17		
Mean Change (%)	7.97			3.65		

Notes. Goodman-Bacon (2021) decomposition of the pooled two-way estimator, on the binary CEP indicator. The forbidden component covers comparisons in which already-treated tracts serve as controls for later-treated ones. Change is the percentage by which the estimate rises in absolute value when those comparisons are removed. Weights and component counts are identical across outcomes within a sample and are reported once at the foot. The decomposition is estimated on tract-quarter means over a balanced panel of tracts, with tract and quarter fixed effects only, as the method requires, so the pooled estimates differ from the person-level binary estimates in Table 2, which carry person and state-by-quarter fixed effects. Bankruptcy and foreclosure are in percentage points and the other outcomes in their own units, as in Table 2.
Source. Federal Reserve Bank of New York Consumer Credit Panel/Equifax; school-level CEP participation data from Ruffini (2022).

Table A3. Rambachan-Roth Breakdown Values

	Any Credit Account	**Number of Accounts**	**Any SD**	**Number of Accounts in SD**	**Total Debt in SD**	**Share of Accounts in SD**	**Share of Debt in SD**
Aggregate	1.25	1.25	0.50	0.50	0.50	0.25	0.25
Revolving	1.00	1.00	0.25	0.50	0.50	0.25	0.25
Auto	0.75	0.75	0.50	0.50	0.75	0.50	0.50
Mortgage	0.75	0.75	0.25	0.25	0.25	0.25	0.25
	Equifax Risk Score	Bankruptcy	Foreclosure				
Creditworthiness and Catastrophic Events (Tract Level)	0.25	0.25	0.25				

Notes. Rambachan and Roth (2023) relative-magnitudes restriction, alpha of 0.05. The breakdown value is the multiple of the largest pre-adoption violation of parallel trends at which the effect first ceases to be distinguishable from zero, so the last value at which a result survives is one grid step lower, and a value above one means the post-adoption violation would have to exceed anything observed before adoption to overturn the result. The grid runs from 0 to 2 in steps of 0.25, so every value is a multiple of 0.25 and no finer distinction is available. Coefficients are placed on the universal base period, which normalizes the quarter before onset to zero as the method requires. Bounds are computed on the Callaway-Sant'Anna event study, for the average over the twelve quarters from onset, with twelve quarters before it. Bankruptcy and foreclosure are estimated in percentage points, which leaves breakdown values unchanged, and the quarter before onset, the reference quarter, is excluded from the pre-onset window. The row for creditworthiness and catastrophic events is estimated at the tract level; its Equifax Risk Score entry is the tract-level estimate, a rise, which the text sets aside because its placebos exceed its effect.

Source. Federal Reserve Bank of New York Consumer Credit Panel/Equifax; school-level CEP participation data from Ruffini (2022).