Branching Out

The Role of Selection in Bank Branch Entry and Economic Growth

Anthony M. Zdrojewski III

July 2024

Abstract

I study the relationship between bank branch entry and local economic growth by exploiting variation in completion across planned branches. Areas where a branch was planned but withdrawn exhibit higher growth in light emissions, an economic proxy, compared to similar areas (selection effect). However, locations where a branch was opened do not significantly outgrow locations where a branch was planned but withdrawn (treatment effect). Using zip code level business formations or SBA-7a loan amounts as outcome variables yields similar results. Selection effects outweigh treatment effects by a factor of between six to twenty-five. These findings challenge previous studies reporting positive treatment effects of bank branches, and instead emphasize banks' skill in selecting locations poised for growth.

1 Introduction

Does financial development increase economic growth or do financial systems develop to support expected growth? A large literature has argued convincingly that both of these mechanisms contribute to the positive relationship between financial and economic development. Therefore, the resulting debate largely centers on the relative importance of the two effects. Bank branch establishment is a natural context to consider this question: does branch establishment increase local economic growth (treatment effect), or do banks establish branches where they expect greater growth (selection effect)? This fundamental question is largely unaddressed due to two main issues: lack of sufficiently granular economic data and identification problems inherent to bank branching decisions. To "shed light" on this question, I employ a recently proposed granular proxy for economic activity, nighttime light emissions (Henderson et al. (2012)), and introduce a novel identification strategy to disentangle treatment and selection effects of branch entry. I find a positive but insignificant treatment effect of branch entry on economic growth that is overshadowed by the selection effect in both magnitude and statistical strength. I obtain similar results studying other outcome variables that are available at a zip code level, including new businesses registrations and SBA 7A lending.

Conventional measures of economic output are generally not sufficiently granular to detect local (i.e. sub-county) effects. Nighttime light emissions observed by satellite are arguably an appropriate measure to bring to bear in this setting. Many economic activities lead to an increase in light emissions: participation in nightlife venues such as bars or restaurants, plant operation, or residential consumption. In fact, Henderson et al. (2012) show that in developing countries, an optimal estimate of economic growth is a roughly equal-weighted blend of conventional GDP growth and growth predicted by contemporaneous changes in (night) lights. Although less interpretable than more standard economic measures, light emissions are available at a fine geographic granularity; the main data in my analysis are available at a resolution of less than one kilometer. This allows me to zoom-in at a hyper-local resolution to identify economic effects that are concentrated around entering branches. I validate this proxy's usefulness by demonstrating that US county-level changes in light emission positively correlate with changes in GDP per capita, both contemporaneously and in subsequent periods. Based on extant literature and my own supplementary analyses, light emissions seem like a credible measure of local economic activity to ascertain the effect of opening a bank branch.

Even with a suitable proxy of local economic activity, identifying the impact of opening a branch is not straightforward because a bank's decision to open a branch is endogenous. In particular, one might think that banks choose to open branches in locations where they expect the economies to grow, even if these locations would grow without the branch. Disentangling this selection effect and quantifying it relative to the actual treatment effect of the branch opening is challenging; researchers generally only observe units (locations) that are affected by both selection and treatment, or by neither. In an ideal setting, a researcher might experimentally manipulate these dimensions by obtaining a set of locations where a bank hopes to establish a branch, and randomly assigning a subset of the locations to be treated (i.e. open a branch) and others to serve as controls. Although such an experiment is obviously infeasible in practice, it is a useful conceptual benchmark by which to evaluate the merit of a candidate natural experiment.

To recreate the key features of this idealized experiment, I exploit data from the OCC's Corporate Application Search (CAS), which records every application to establish a new branch by the banks the OCC governs. Most of these applications are approved and consummated; but, a small subset of the applications are rejected, withdrawn, or allowed to expire. The corresponding locations are subject to a selection effect on the part of the bank, but not the treatment effect of the branch actually opening. If the decision to consummate a submitted application is exogenous, then differences in economic outcomes between opening and "nearly opening" locations are credibly attributable to the treatment effect of the branch's opening. Similarly, differences in outcomes between nearly opening locations and observably similar locations where branch establishment is not sought are attributable to selection. Although banks are not required to provide a rationale for not completing planned branches, anecdotal evidence from the OCC suggests that withdrawals are "usually...due to delays in construction" (Office of the Comptroller of the Currency Licensing Office (2023)), and therefore arguably unrelated to local economic conditions. The exogeneity of decisions to withdraw is further supported by the stringent application requirements, high rates of conversion into operational banks, and the short timeline over which effective withdrawals must occur. I discuss these and other institutional details at length in subsequent sections.

From CAS, I compile a list of locations where a bank branch establishment is sought and I define withdrawal (opening) events as the first time an application is withdrawn (consummated) at a particular location. I employ a carefully designed difference-in-differences approach to avoid the biases that can occur in a setting with staggered treatment (Goodman-Bacon (2021)). The unbiased approach proposed by Callaway and Sant'Anna (2021) constructs counterfactuals for each treatment cohort using units that are never and/or not yet treated. I adapt this approach to my setting by instead using the contemporaneous cohort for which treatment (branch entry) was withdrawn as counterfactuals for the treated cohort. Compared to other approaches this method is computationally inexpensive and allows the researcher to make informed choices of how to weight estimated effects across cohorts and event-time. It also allows for visual confirmation that there are no pre-treatment divergences in outcome across my treated and control units, lending crucial support to the parallel trends assumption under which my estimates are unbiased.

This procedure reveals a slight, generally statistically insignificant treatment effect of opening a bank branch on a location's economic activity. On average, opening a bank branch causes light emissions to grow by about .77 percentage points after seven years. However, these estimates are statistically insignificant for every event year except for year one. Nonetheless, back-of-the-envelope calculations suggest that this seven-year treatment effect corresponds to an increase in seven-year per-capita GDP growth of roughly .06 percentage points. Repeating the procedure with withdrawing locations as "treated" units and arbitrary economic locations (post offices extant prior to my data's coverage) as comparison units reveals a selection effect; it shows how different economic growth is in places that nearly open a bank (but do not) from otherwise similar locations unconditional on planned bank entry. This effect, perhaps unsurprisingly, is much larger than the estimated treatment effect of the bank; the effect is statistically significant in each of the seven years I study post-withdrawal, and steadily increases to about 3.6 percentage points seven years after opening. This corresponds to a striking increase in seven-year per-capita GDP growth of roughly .29 percentage points. Together, these results indicate that banks open branches where they expect subsequent economic growth but that the branch itself does little, if anything, to cause that growth to increase.

To bolster confidence in my interpretation of increased luminosity as an economic effect, I next consider new business registrants within a zip code as an outcome variable. I find no evidence of a treatment effect of branch entry on new business formation rates. Once again, however, the selection effect is striking; both per-year and average effects are statistically significant across all seven years post-treatment. As a fraction of the number of businesses formed in the year before withdrawal, withdrawing zip codes see an increase in their number of new business registrants of 5% (each year) more than the counterfactual consisting of all other zip codes. In other words, banks select to branch in zip codes whose economies will see an increase new business formation, but don't seem to actually cause that increase by branching there.

Although the lights and new business data reflect a weak treatment effect of branch entry, if any, I next consider a possible mechanism through which entry might affect economic outcomes; small business lending. From the Small Business Administration (SBA), I compile zip code level data on the amount of loans made to borrowers in the entered zip code. Repeating the procedure reveals that there might be a slightly positive treatment effect. The selection effect, however, is about twenty times as large in magnitude and is significant over every event year; as a fraction of the amount of loans issued in the year before treatment, zip codes with withdrawn branches see an increase in loans issued in each of the ensuing seven years of 74% more than their counterfactuals. The selection effect of branch entry is larger than the treatment effect not only for variables measuring local economic output, but also in terms of the *amount of capital* provided.

There is a lengthy theoretical literature on the relationship between financial development and economic growth. Schumpeter (1911) argued that financial development is crucial to promote innovative growth. Robinson (1952) and others, however, believe that financial systems develop in anticipation of future growth, but do not actually cause it. Lucas (1988) abstracts from monetary matters in his models of economic growth, but notes that "insofar as the development of financial institutions is a limiting factor in development more generally conceived I will be falsifying the picture, and I have no clear idea as to how badly." Since then, decades of empirical literature has attempted to quantify just how severe Lucas (1988)'s abstraction was.

A number of papers, starting with Jayaratne and Strahan (1996), have attempted to address the question through a lens of bank branching deregulation. Although the deregulation events provide a pathbreaking natural experiment, this literature [1] may still be plagued by selection biases due to the endogenous choice of a state to deregulate (Freeman (2002), Kroszner and Strahan (1999)). Moreover, it largely fails to address the more fundamental and granular question of what happens to a local economy as a result of branch entry, the topic of this paper.

[1] Recent econometric developments (e.g. Goodman-Bacon (2021), Callaway and Sant'Anna (2021)) have cast doubt on the interpretation of many of the treatment effects estimated in this literature, as emphasized by Zdrojewski and Butler (forthcoming)

This paper also relates to the literature beginning with Petersen and Rajan (2002) which examines the importance of distance between lenders and borrowers in credit markets. Gilje et al. (2016) show that branch networks can transmit economic growth to improve access to finance across county lines. Nguyen (2019) uses overlap in branch networks following bank mergers as a instrument to show that branch closures cause local business lending to decrease. Although these papers make causal claims relevant to my topic, none are able to disentangle selection and treatment effects of bank entry, or access to finance more broadly. To my knowledge, mine is the first paper claiming to make this contribution. Moreover, though my results are not contradictory to theirs, they do seem to paint access to finance as less important for economic growth than suggested by this literature.

This paper makes several key contributions. At the broadest level, it offers a clean answer to the age-old debate of whether finance causes growth (Schumpeter (1911)) or vice versa (Robinson (1952)), albeit in a particular and limited setting. In the context of bank branching in the US, anyway, financial development seems to occur in expectation of economic growth, but does relatively little to cause it. This finding should dissuade any attempt to artificially encourage branch establishment in underdeveloped locations to cause growth. Second, the paper shows that the selection effect is dominant not only in terms of output, but also in terms of the amount of capital provided. This suggests that areas with promising economic futures are likely to receive substantial loans regardless of branch entry, casting doubt on the importance of the geographic location of capital in the modern economy. Third, the paper further confirms the usefulness of night lights as an economic proxy through not only my validation exercises, but also the similarity in observed effects to those using more traditional economic variables.

2 Data and Background

This paper uses data from several standard sources. Data on economic variables such as population or GDP are taken from the Bureau of Economic Analysis (BEA), and branch-level deposit data are taken from the Federal Deposit Insurance Commission (FDIC)'s summary of deposits data, which spans 1994 to 2023. Zip code level data on new business registration comes from the Startup Cartography project, which is introduced in Andrews et al. (2022), and cover years 1988 to 2016. SBA-7a loan data come from the Small Business Administration and cover years 1991 to 2024. From the Harvard Dataverse (Blevins and

Helbock (2021)), I collect a list of post office locations and their establishment dates for use as reference points.

The list of bank branch applications that serve as the primary units of analysis in this paper are scraped from the OCC's Corporate Application Search (CAS), and span years 1990-2023. The following subsections discuss this data in detail, along with my preferred proxy for local economic activity, night-time light emissions observed via satellite, which follows Henderson et al. (2012).

2.1 OCC Branch Establishment Applications

The OCC is the assigned regulator for all nationally chartered banks, federal savings associations, and federal branches and agencies of foreign banks. These institutions represent a sizeable portion of the branching footprint within the United States; as shown in the middle panel of Figure 1, roughly half of US branches belong to OCC-regulated banks. Moreover, the figure also shows that more than half of deposits are with OCC-regulated banks. These institutions are generally required to submit applications to the OCC pertinent to a number of their business operations, including branch establishments, closures, and relocations. This requirement results in a paper trail documenting locations where banks intended to establish a branch, even if this branch is not actually realized. The OCC makes the data on these applications available in its Corporate Applications Search[2], allowing me to observe variation in completion across a sizeable portion of planned branches in the US.

For an OCC-regulated bank to establish a new branch, they must undergo an involved multi-step process which is likely to dissuade unserious applications. As part of the submission process banks "must publish notice of the proposed branch establishment or branch/main/home office relocation in a newspaper of general circulation in the community or communities in which the applicant proposes to establish ... the branch" Office of the Comptroller of the Currency (2023). After such a publication, a 30-day public comment period is open, at which point any person can write the Director of Licensing at the nearest OCC office to provide comment on the proposed branch.

In addition to providing proof of public notification of intent to establish the branch, banks also must verify that the branch establishment will not be in violation of either state or federal laws, and in many cases will be required to submit legal opinions to support this

[2] https://apps.occ.gov/CAAS_CATS/

argument. Additionally, the bank has to affirm that the branch will not create a prohibited management interlock, have an effect on the human environment (within the confines of the National Environmental Policy Act), or affect any historical location protected under the National Historic Preservation Act (Office of the Comptroller of the Currency (2023)). If the proposed branch fails any of these criteria, the bank must identify how, and in some cases provide a plan of rectification. In short, applying to open a branch requires banks to step through a series of pain points involving public notification, premises suitability, and legal compliance. These institutional details of the application process strongly suggest that banks are unlikely to apply to establish a new branch in a location without first carefully researching it and establishing its suitability.

Each application to establish a new branch must provide a precise planned location. Figure 2 shows an example record from the CAS database of a completed branch in Newberry, Florida. The record details the precise location of the branch, as well as the the beginning and end dates of the associated comment period and the days on which the application was received, approved, and consummated. Figure 3 is similar, but instead shows an example of a branch application which was planned for Ocoee, Florida, but was eventually withdrawn.

Figure 4 provides a flowchart of the application process to help visualize the common outcomes of applications, their frequencies, and the approximate timing between subsequent steps in the application. As indicated in the leftmost node, the CAS provides data on 42,396 branch establishment applications. Subsequent nodes report the percentage of these applications which reach each step, and the nodes are scaled in proportion to the total number of applications received. For 96.6% of these applications (the next node), the first recorded step in the application process begins with receipt (and the rest seem to be data errors concentrated in the first year of data, 1990). After the application is received, a decision is typically made by the OCC within approximately a month. 92.9% of applications are received and then approved, whereas only .1% (less than 100) are denied, and another 2.7% are withdrawn before a decision from the OCC is even reached. Following approval, banks have 18 months to begin operations at the branch prior to the application expiring. As Figure 4 indicates, 84.4% of the applications are approved and effective. 4.9% of applications are approved but eventually withdrawn, and 2.6% are approved but allowed to expire.

It is worth emphasizing how rare it is for an application to be denied. Out of all applications which have a documented decision from the OCC, 99.8% are approved. This fact suggests that banks diligently ensure compliance with all state and federal branching restrictions prior to application, bolstering the credibility of application as a serious signal of intent to branch. Although rejected applications generate variation that is cleanly exogenous from the perspective of the establishing bank, they are of limited use to researchers because of how few of them there are. In this paper I therefore focus on the more common occurrence of an application which is either rejected, withdrawn, or allowed to expire. For the sake of simplicity I refer to these collectively as unfulfilled or withdrawn applications. In other words, my identifying variation comes from the applications represented in Figure 4 by the "Denied" node, the "Expired" node, and the two "Withdrawn" nodes.

As demonstrated in Figure 4, applications receive their decisions in short-order (typically shortly after the 30-day comment period has ended). Once an application is approved, the bank has 18 months to begin operations at the specified location. If after 18 months the branch is not conducting business operations, then the bank must either apply for an extension, re-apply, or forfeit their plans to establish a branch. This relatively short required turnaround is a useful feature of the setting. It makes it more likely that banks are already working on the arrangements needed to begin operations while the application is being reviewed, which further supports the seriousness of their intent to branch. Moreover the quick timeline means that local economic conditions around the branch are unlikely to have changed from the time the bank signaled its interest in establishing a branch to the time of the withdrawal.

In Figure 5, I present the number of total applications to establish a branch each year (top panel) and the percentage which are ultimately unfulfilled (bottom panel). The number of applications each year peaks in the late 1990's and generally decreases thereafter, with a substantial drop around the Great Financial Crisis in 2008. The percentage of applications which are ultimately unfulfilled varies over time, with an average of roughly 10% across years. Although the unfulfilled percentage increased to 20% in 2009, this does not seem to be straightforwardly attributable to the recession when compared with the similar rate in 2011, or the even higher rate in 2015. Rather, the drop in applications overall following the financial crisis seems to have increased the volatility of the withdrawal rate by way of reducing its denominator. Overall, the figure does not seem to provide compelling evidence that withdrawal rates fluctuate as a function of the macro-economy, which is reassuring for my identification claims.

Figure 6 plots on a US map the location of the withdrawal events that are used in my main analysis, which uses nightlights as an outcome variable. Reassuringly, the map is unsurprising; withdrawal events are clustered around population centers, but there are also a number dispersed through more rural areas. The map mainly shows that the withdrawal events are not condensed in particular odd regions, which might obfuscate their suitability as comparison units for places that fulfill the application.

2.2 Night Lights

2.2.1 Data Description

In the 1970s, the United States Air Force Defense Meteorological Satellite Program (DMSP) began using satellites to detect moonlit clouds via Operational Linescan System (OLS) sensors. Although the program's intent was meteorological, the sensors also recorded light emitted by human activity. Each satellite circles the earth 14 times each day, registering a luminosity observation between 8:30 and 10:00 pm local time every day. The digital records of these observations are available dating back to 1992. Prior to releasing the data to the public, the National Oceanic and Atmospheric Administration's (NOAA) National Geophysical Data Center pre-processes this data to remove daily observations affected by natural variation in light, including moon and sunlight, natural fires, and auroral activity (Henderson et al. (2012)). The data center then averages data from each orbit over time to create satellite-year level datasets, which are publicly available through Google Earth Engine for years 1992 through 2013.

The data for each satellite-year is a grid which records light intensity at a 30 arc-second granularity. Henderson et al. (2012) contextualize this measurement by noting that this grid results in pixels which are approximately 0.86 square kilometers at the equator. Further north or south, the arc-second grid tightens so that these pixels become even smaller. In particular, Henderson et al. (2012) note that pixel size varies in proportion to the cosine of latitude so that a pixel in London (51.5 degrees north) is roughly .53 square kilometers. A similar calculation with a US reference point shows that a pixel in Washington, D.C. (38.9 degrees north) will be roughly .67 square kilometers. This data therefore has a much finer geographic granularity than traditional economic data, allowing the researcher to examine more localized units of observation. However, the technology employed to capture these images are dated, and therefore the data present some unique challenges for an economist.

First, observations taken from the same satellite in different years are not immediately comparable due to the changing position of the satellite in its orbit. Thus, time series analysis requires an intercalibration between separate satellite-years. I calibrate individual pixels following Elvidge et al. (2009), and take the mean luminosity of these pixels from the given annual composite. For years where there are multiple satellites available, I calculate the luminosity in that year as the average value recorded by the two satellites.

Another difficulty with DMSP data is that its sensors can only score each pixel from zero (darkest) to sixty-three (brightest). As a result, variation in extremely bright areas is unobservable, limiting the number of locations which are useful for making economic inferences. I therefore discard extremely bright areas, those with average luminosity greater than fifty-five, from most of the analyses to examine only those where variation in light is truly observable.

In Figure 7, I demonstrate visually how I calculate lights around an individual branch. It depicts and aerial view of the Newberry, Florida branch from Figure 2, indicated by the black dot at the center. In my primary analysis, I calculate the average luminosity of pixels within five kilometers of each branch. In the example figure, this area is shaded as pink. In secondary analyses, I then consider outer rings ("doughnuts") which are farther away from the branch. For example, the area shaded in blue in the figure depicts a ring with an outer radius of ten kilometers, and an inner radius of five kilometers. Using this approach I can explore the degree to which the treatment or selection effects are localized. I do this using rings defined by inner and outer radius pairs of 0 and 5 km, 5 and 10 km, 10 and 20 km, and 20 and 40 km.

2.2.2 Economic Interpretation

The main advantage of the nighttime lights data in my setting is its geographic granularity, as discussed in the previous section. This advantage comes with the obvious cost of not being readily interpretable in terms of standard economic quantities. Henderson et al. (2012) show that in developing countries the information contained in this measure can improve estimates of economic quantities like GDP. In this section, I expand upon these results to argue that nighttime lights are a useful proxy for local economic activity within developed countries.

Because my paper primarily deals with the measurement of economic growth across localities, I begin with cross-sectional (ie Fama-Macbeth) regressions of county-level GDP growth on growth in lights in that county. Holding light growth constant in a given year, I calculate GDP growth contemporaneously (t), over the ensuing year (t+1), and over the ensuing five years (t+5). Table 1 reports the results using GDP growth as the dependent variable and Table 2 reports the same but using GDP on a per capita basis. The first column of Table 1 indicates that a one percentage point growth in lights corresponds to a 3 basis point increase in GDP growth in the same year. This is not simply attributable to increasing population, however, as the first column in Table 2 reports a statistically significant per capita effect of two basis points as well. Moreover, the statistically significant and economically meaningful coefficients in the second and third columns of Table 1 show that increases in lights are also predictive of a county's future GDP growth (although not significant on a per capita basis).

I next consider a lower-frequency analysis to demonstrate the long-run cross-sectional relation between light growth and economic growth. I calculate growth rates in GDP and in lights for each county over the period where the data overlap, 2001-2013. Figure 8 plots each county in blue with its per-capita GDP growth over this time against its growth in lights. There is a clear positive relationship between the two, captured visually by the best-fit line whose slope is a statistically significant .21. This means that on average, a one percentage point increase in luminosity between 2001-2013 corresponds to a 21 basis point increase in GDP. I label several counties which are clear visual outliers, and note that these are all beneficiaries of the US shale boom in this period. I therefore also run a median regression to reduce the influence of these points, and obtain a more modest but still significant estimate coefficient of .08.

While unorthodox, the use of nightlights as a proxy for economic activity has a strong intuitive basis, as nearly all consumption activities at night require light emissions. Henderson et al. (2012) build on this intuition to show that it is a useful economic proxy in developing nations. In Table 1, Table 2, and Figure 8, I validate the usefulness of the proxy at a subnational (county) level. With its validation as a proxy and a rough sketch of its economic interpretation, light emissions seem like a promising outcome variable to study the relationship between bank branch entry and local economic growth.

3 Methods

This paper employees an event-study approach to identify treatment and selection effects of branch entry. I focus on the first time an application to establish a branch is withdrawn/completed at a particular location. For analyses using lights as a variable, I define a location as a latitude/longitude coordinate pair, and for each of the other variables I use a zip code level definition.

Most approaches consider treated units and estimate a counterfactual for them consisting of untreated units. In my case, however, the number of treated units (locations where a branch enters) far exceeds the number of untreated units (locations where a branch application is withdrawn) each year. Table 3 shows the number withdrawal events that are used in my primary analysis, using lights as an outcome variable, broken out by year. There are ~ 40 withdrawals and ~ 625 opening locations in the average year. As a result, matching each opening location to one or more withdrawal locations is likely to produce unreliable counterfactuals. I instead consider each withdrawal event and match it to a counterfactual consisting of similar opening events. The difference between this counterfactual and its corresponding withdrawal event is thus my estimate of the treatment effect of branch entry at that location.

Similarly, I identify the selection effect at each location by constructing a counterfactual for each withdrawal that instead uses a pool of comparison units which are not conditioned on branch entry plans. The difference in outcome between the withdrawing location and this counterfactual is my estimate of the selection effect: the difference in outcome at that location that is associated with selection by banks for branch entry. For zip code level variables like new business registrants or SBA-7a loan amounts, I simply use the set of other zip codes as comparison units to find the selection effect. Using lights as an outcome variable is less straightforward, however, because randomly selected geographic coordinates are likely to be very different from those in the choice set of locations for entering banks. For instance, a point in a remote and uninhabited area is probably not useful for comparison. So, a bit of careful consideration is merited in selecting comparison locations.

Intuitively, good comparison points will be accessible to a nearby population maintaining a reasonable level of economic activity. To identify such locations, I compile a list of post offices that were extant in 1991, prior to the start of my lights data. With their exact locations I can calculate the light emissions around post offices in exactly the same manner as I do for bank branches (see Figure 7). The first branch entries I consider using the lights data are in 1996, so by using post offices extant in 1991 I am assuaging concerns that these locations may be subject to entry timing by the USPS. The presence of a post office ensures that there is a baseline level of nearby economic activity, making these locations more suitable for comparison than randomly selected geographic coordinates. In fact, in 2021 there was a postal banking pilot program launched that allowed customers to effectively cash business and payroll checks at certain post office locations (Anthony (2024)) [3]. Although customers seem to have had little interest in banking with the USPS, the plausibility of the program speaks to the suitability of using post offices as comparison locations for bank branches.

I employ a staggered difference-in-difference approach to estimate both the treatment and selection effects of branch entry. A key advantage of such a setting is that the staggered nature of interventions reduces concerns that there might be a single event that confounds treatment by simultaneously affecting all of the observations. However, a burgeoning literature has emerged pointing out that some commonly used empirical approaches in a staggered difference-in-difference setting may generate biased estimates of treatment effects (Goodman-Bacon (2021), Sun and Abraham (2021), de Chaisemartin and D'Haultfœuille (2020)). A major reason for concern is that some approaches, like two-way fixed-effect regressions, implicitly compare treated units to a control group which includes previously treated units. With dynamic treatment effects, this can result in estimates may not only be biased, but even obtain the wrong sign (Goodman-Bacon (2021)).

To avoid making these inappropriate comparisons to previously treated units, Callaway and Sant'Anna (2021) propose a methodology that constructs counterfactuals for each treatment cohort using units that are never and/or not yet treated. I adapt this approach to my setting by making my comparisons within treatment/withdrawal cohorts. That is, I compare the outcomes of would-be branches which are withdrawn in a given year with those of branches that are completed in the same year. This allows me to estimate what the treatment effect of branch entry would have been for those locations which are "nearly" entered. Compared to other approaches this method is computationally inexpensive and allows the researcher to make informed choices of how to weight estimated effects across cohorts and event-time. It allows for an intuitive approach to adding covariates and ensures that these covariates do not affect estimates by themselves changing as a result of the treatment. The approach also allows for visual confirmation that there are no pre-treatment divergences in outcome across treated and control units, lending crucial support to the parallel trends assumption under which the estimates are unbiased.

The following discussion formalizes my adaptation of the approach employed by Callaway and Sant'Anna (2021). To make the intuition of the approach clear, I discuss what the

[3] https://www.cato.org/blog/no-customers-no-success-postal-banking-failure-exposed

estimates and assumptions look like in a case without covariates. Then I describe intuitively how inclusion of covariates affects the estimation.

3.1 Adaptation of Callaway and Sant'Anna (2021)

Let G index the year that a particular location fulfills or withdraws a branch application. W is an indicator variable equal to one if the location withdraws its application, and F is an indicator variable equal to one if the location fulfills its application. I adopt a potential outcomes framework; $Y_t(g,w)$ is what the outcome would be at time t if the unit belongs to cohort g and has $W = w$. I define the Average Treatment effect on Withdrawn units (ATW - analagous to ATT of Callaway and Sant'Anna (2021)):

$$ATW(g,t) = \mathbb{E}[Y_t(g,0) - Y_t(g,1)|W + F = 1, G = g, W = 1] \tag{1}$$

This is the expected difference in outcome at time t between a unit that completes versus withdraws its branch establishment in cohort g. The conditional expression in the above simply requires that all applications are either completed or withdrawn, and that this application is withdrawn in occur in year g.

The most important assumption on which my approach relies is analogous to the parallel trend condition required by Callaway and Sant'Anna (2021). I assume that the following holds for each treatment group g and each time $t \geq g$ and $s \geq t$:

$$\begin{aligned} &\mathbb{E}[Y_t(g,1) - Y_{t-1}(g,1)|W + F = 1, G = g, W = 1] \\ &\quad = \mathbb{E}[Y_t(g,1) - Y_{t-1}(g,1)|W + F = 1, G = g, W = 0] \end{aligned} \tag{2}$$

The above states that, in expectation, if the locations that withdraw at time g had instead completed application at that time, they would have changed the same amount between $t-1$ and t as did the locations that actually completed. Under this assumption (and more mundane assumptions analogous to those of Callaway and Sant'Anna (2021) which I omit), $ATW(g,t)$ can be identified by:

$$\begin{aligned} ATW(g,t) = &\mathbb{E}[Y_t(g,0) - Y_{g-1}(g,0)|W + F = 1, G = g, W = 0] \\ &- \mathbb{E}[Y_t(g,1) - Y_{g-1}(g,1)|W + F = 1, G = g, W = 1] \end{aligned} \tag{3}$$

Abusing notation slightly, the above can be simply estimated by the appropriate difference in differences of sample means:

$$\widehat{ATW(g},t) = \overline{Y}_t(g, W=0) - \overline{Y}_{g-1}(g, W=0) - (\overline{Y}_t(g, W=1) - \overline{Y}_{g-1}(g, W=1)) \tag{4}$$

That is, a simple estimate of the cohort-time treatment effect is the difference in changes in average outcome for fulfilling versus withdrawing units. Estimation of selection effects is similar, but uses withdrawing locations in the first difference and in the second uses other suitable locations regardless of whether a branch application existed there.

3.2 Inclusion of Covariates

In this section, I build on the framework of the previous section to provide an intuitive explanation of how I modify the approach to account for covariates. Callaway and Sant'Anna (2021) propose an inverse probability weighting (*ipw*) approach which essentially re-weights control units so that they are more similar to the treated units in terms of the covariates. Let W_g be an indicator variable equal to one when a unit withdraws in cohort g, and F_g be an indicator variable equal to one when a unit is fulfilled/effective in cohort g. $p_g(X)$ is the probability that a unit from cohort g is withdrawn, conditional on the covariates, X. Then, the $ATW(g,t)$'s can be estimated by:

$$ATW_{ipw}(g,t) = \mathbb{E}\left[\left(\frac{\frac{p_g(X)F_g}{1-p_g(X)}}{\mathbb{E}\left[\frac{p_g(X)F_g}{1-p_g(X)}\right]} - \frac{W_g}{\mathbb{E}[W_g]}\right)(Y_t - Y_{g-1})\middle| W + F = 1, G = g\right] \tag{5}$$

The above is similar to Equation 3 in that the product of differences it contains can be re-written as a difference-in-differences. However, the first difference is re-weighted to emphasize fulfilled units in each cohort which are most like the withdrawing units of that cohort. To see that, consider that as $p_g(X)$ increases for a given unit, $\frac{p_g(X)F_g}{1-p_g(X)}$ also increases, assigning more weight to that unit. The division by the expected value of this term, $\mathbb{E}\left[\frac{p_g(X)F_g}{1-p_g(X)}\right]$, ensures that across fulfilled units these weights sum to one.

To estimate Equation 5, one simply replaces the given expectations with their sample means, and replaces $p_g(X)$ with $\widehat{p_g(X)}$, which I estimate using a logistic regression approach following Callaway and Sant'Anna (2021).

4 Results

4.1 Main Results

In this section, I use variation in the completion of planned bank branches to disentangle the selection and treatment effects of branch entry with respect to three main economic variables: night lights observed around the branch, new business registrants within the branch's zip code, and SBA-7a loan amounts to borrowers within the branch's zip code. To avoid overlapping events, I include only first withdrawals and completions at a particular location. For each variable, I require four years of pre-treatment data and eight years of post-treatment data. I employ my adaptation (described in the previous section) of the approach in Callaway and Sant'Anna (2021), controlling for the outcome variable in each of the four years pre-treatment.

For consistency and computational simplicity, I restrict the pool of comparison units used in estimating the selection effect to have the same size as that used in estimating the treatment effect. To decide which units to keep in the selection pool, I rank them on their similarity to the withdrawing units' level of the outcome variable pre-treatment. Then I keep the closest n observations, where n is the number of units with fulfilled applications in that year.

In each of the plots, I calculate event-time effects by weighting the estimated cohort-time treatment effects ($ATW(g,t)$) by the number of withdrawals. Thus, my scheme weights my estimates in proportion to how much of my experimental variation occurs in each cohort. A simple equal-weighting scheme across cohorts gives similar results.

4.1.1 Night Lights

To explore the relationship between branch entry and economic growth, my first analysis uses nighttime light emissions within five kilometers of the proposed location. Table 3 lists the number of withdrawn branches used in this analysis from each year, along with the number of fulfilled branches in the donor pool. With 432 withdrawal events occurring over eleven years, I have substantial variation to test for an effect.

In Figure 9 I show the estimated per year selection effects (blue) and treatment (green) of branch entry on the log-level of lights. Neither effect shows evidence of a trend pre-treatment, partly as a result of controlling for the pre-treatment outcomes. Post-treatment, a strong selection effect immediately manifests, and increases over time. By year seven, the selection effect is .036, meaning that withdrawing locations grow their lights by 3.6 percentage points more after seven years than post offices that saw similar growth in lights pre-treatment. The estimated selection effect in each year is positive and statistically significant at all levels with an average per-year effect of .028.

The treatment effect, by comparison, is much smaller. Only year one shows a statistically significant treatment effect, and the average effect across years is only .004, or .4 percentage points. In year seven, the effect is .0077 or .77 percentage points, but statistically insignificant. Across years, the average selection effect is roughly six times that of the treatment effect. Though crude, the coefficients from Figure 8 provide one way to contextualize these effects in terms of more standard economic figures. Using the median-regression estimates, the year-seven selection effect of .036 corresponds to a 3.6*.08 = .29 percentage point increase in the seven-year growth rate in per capita GDP. Similarly, the year-seven treatment effect of .0077 corresponds to a .77*.08 = .06 percentage point increase in the seven-year growth rate in per capita GDP. In terms of economic magnitudes, the selection effect is striking whereas the treatment effect is negligible.

The estimates presented in Figure 9 are valid under the assumption that withdrawn branch applications do not occur as a function of expected local economic growth. Section 2.1 makes the argument that this is a credible assumption, but it is worth considering what it means if the assumption does not hold. Intuitively, it seems unlikely that withdrawals would be positively correlated with expectations of economic growth; stronger local growth will probably lead to greater demand for capital to finance new ventures, increasing bank profitability and therefore incentivizing the bank to complete its branch. On the other hand, one might imagine that withdrawals are triggered by a bank choosing to not complete entry in locations for which they have lower expectations of growth, possibly due to an unforeseen change affecting the economy after the submission of the application. In this sense, if the assumption of application withdrawn being unrelated to expectations of economic growth is incorrect, it is likely in the direction of mechanically over-estimating the treatment effect, and under-estimating the selection effect. Thus, if my estimates are invalid, the bias is likely to be working *against* my findings, as I document an economically trivial treatment effect and a large selection effect. Moreover, the positive and strong selection effect makes this bias unlikely; if withdraws occur due to an economically substantial downgrading of the local economy, we would expect to see a negative selection effect.

The finding that branch entry has an economically small and statistically weak effect on light growth seems to run counter to prior literature. It has emerged as a stylized fact of the finance-growth nexus that access to finance improves economic outcomes (Jayaratne and Strahan (1996), Chava et al. (2013), Butler and Cornaggia (2011), many others), so why do I fail to find an economically compelling effect in this setting? One possibility is that the effect of opening a branch only manifests in areas with limited access to capital ex ante. To test this, I calculate the dollar value of deposits within X km of each proposed branch, and repeat the procedure limiting the sample to observations in the bottom quartile of that value for each year. The results (ADD) are little changed. Another plausible explanation is that I am simply using a poor proxy for economic growth which makes detecting the effect difficult. Given how thoroughly established nightlights is as a proxy in the development literature, this seems unlikely. Nevertheless, in the next subsection I consider a more orthodox economic variable, zip-code level new business registrants, and find similar results.

4.1.2 New Business Registrants

The interpretation of the results from Section 4.1.1 are only as good as my proxy for economic activity, namely nighttime light emissions. Although like prior literature (e.g. Henderson et al. (2012)), I document the usefulness of the proxy, there might persist concerns that the economically null treatment effect I observe is due to my choice of proxy, rather than a lack of a true economic effect. Therefore, I extend the results by considering a more conventional economic proxy, number of new business registrants, which is available at a zip code level courtesy of the Startup Cartography Project (Andrews et al. (2022)). In Table 4 I document the number of withdrawals that occur each year, along with the number of zip codes that are available to use as comparison units. Once again, I have substantial variation coming from 513 withdrawal events, dispersed across 18 years.

Figure 10 presents the per-year selection and treatment effects associated with a branch entering a zip code on the log of its new business registrations, and largely tells the same story as Figure 9. I find strong evidence of a selection effect, with all but one of the individual yearly effects being statistically positive. The effects are large; the average effect across years is .05, meaning that in the seven years following withdrawal, zip codes with a withdrawn branch see an average of 5 percent more new business registrants than other similar zip codes, as a portion of the number in the year before withdrawal. As was the case in the nightlights analysis, however, there is essentially no evidence of a treatment effect. None of the individual yearly effects are statistically significant, and two of the eight are in fact negative. The average treatment effect across years is a statistically insignificant .002, or .2 percent. Across years, the average selection effect is 25 times that of the average treatment effect. In short, using zip code level data on new business registrants, I again find no evidence of a treatment effect of branch entry on economic growth. Taken together, the results of Section Figure 10 assuage concerns that effects reported in 4.1.1 are simply attributable to a poor choice in outcome variable.

Sections 4.1.1 and 4.1.2 show that banks selectively enter into high-growth areas but are not themselves a significant cause of local growth. If bank branches are a reasonable metric of financial development, then these results seem to conflict with the battery of empirical papers that document positive effects of financial development on growth. This apparent disparity motivates my next analysis, which tests whether bank branch entry increases access to capital (SBA-7a loan amounts), and whether areas selected by banks receive greater capital even when the branch is not realized.

4.1.3 SBA-7a Loan Amounts

In this section, I explore whether branch entry impacts the amount of capital available for small businesses in the entered area by using zip code level data on total SBA-7a loan issuance amounts. Table 5 reports the number of withdrawing zip codes and available comparison zip codes within each calendar year cohort. Using this dataset, I have about half of the number of withdrawals as were available in the prior analyses. This is because there were many zip codes that had values of zero in this data, making them unsuitable for use as logged values. Nevertheless, I have substantial variation stemming from 271 withdrawal events with about ten times as many comparison zip codes, staggered over 22 years.

Figure 11 shows the event-year selection and treatment effects of branch entry on the log of the total loan amounts issued within a zip code. The results are similar to those that rely on nightlights or new business registration data; selection effects are strikingly large and significant, with much weaker evidence of a treatment effect. The average selection effect across years is estimated at .739, which is roughly 13 times the size of the average treatment effect of .056. In the average post-treatment year, branch entry caused the dollar value of SBA-7a loans to borrowers in the entered zip code to increase by 5.6% of the pre-treatment value. On the other hand, withdrawing zip codes see 74% greater loan amounts, as a fraction of the pre-treatment value, in the year following withdrawal than other zip codes with similarly tending loan amounts (unconditional on their plans for branch entry). Although the overall treatment effect is significant, only one of the per-year treatment effects is significant and the overall effect is smaller and less statistically reliable than the selection effect.

The results of Figure 11 help to resolve the seeming contradiction between the previously reported results and the extant literature, and so deserve special discussion. Not only is the selection effect dominant in terms of economic growth, but it even dominates the main channel through which one might expect an effect, increased access to capital. This suggests that areas with promising economic futures are likely to receive substantial loans *regardless of branch entry*. Therefore, my previous results do not declare a null effect of access to finance on economic growth, but rather suggest that bank branches are not actually a good measure for local access to capital. As pertains to the effect of access to finance on economic growth, the focus this literature (Jayaratne and Strahan (1996), Chava et al. (2013), etc.) has placed on bank branches seems to have been misplaced.

4.2 Effect Localization

My paper is most closely related to the seminal work of Jayaratne and Strahan (1996) and the family of papers it has spawned (Jayaratne and Strahan (1998), Huang (2008), Chava et al. (2013), Berger et al. (2021), etc.) which use state-level regulation changes as plausibly exogenous shocks to branching networks. These papers all report positive effects of loosening branching restrictions, whereas my paper finds minimal effect of branch entry. The disparity between these results highlights an important value-add of both my identification strategy and my data: the ability to study *local* effects of branch entry, rather than aggregate/state level changes in response to deregulation. The ability to "zoom-in" at a local level uncovers two novel results: that banks select to branch in locations which are primed for growth, and that the branches themselves contribute little to this growth. In this section, I further exploit the geographic granularity of the nightlights data to explore how localized these results are.

Around each branch, I calculate the mean luminosity observed within a radius of 5 km, and then over rings/doughnuts with inner and outer diameter pairs of 5 and 10 km, 10 and 20 km, and 20 and 40 km. Figure 7 shows an aerial view of an example branch, illustrating how theses areas are drawn. The black dot at the center of the figure represents the branch, whereas the innermost circle of 5 km is shaded in pink and the ring with an inner radius of 5 km and an outer radius of 10 km is shaded in blue. For each of these areas, I calculate the mean luminosity observed across pictures, repeating this for both completed and withdrawn branches as well as post office locations (for comparison, as in Section 4.1.1). Using this approach I can explore the degree to which the treatment or selection effects are localized.

Figure 12 shows the results. The top left panel recreates Figure 9, using a simple circle of radius 5 km, and then from left-to-right the remaining panels repeat the exercise using increasingly expansive rings, as described previously. Several facts emerge from this exercise. First, the selection effect maintains roughly the same magnitude across all four panels, perhaps slightly increasing at the longer range distances in the bottom panels. This indicates that the selection effect is not hyper-localized, but that banks in fact pick general areas which are primed for growth, rather than hyper-specifically. The treatment effect, on the other hand, decreases in the longer distance specifications, even becoming generally negative at the farthest specification. In addition to being economically weak, the effect of branch entry is hyper-localized, attenuating quickly outside of 5 km.

Overall, the results of Figure 12 provide a novel insight to the literature on the bank branching and growth nexus: studies using aggregate (i.e. state-level) shocks may suffer from built-in biases because they are affected by selection issues. By focusing on localized effects, my findings reveal that banks tend to establish branches in areas already primed for growth, and that the branches themselves do little to spur further development. This suggests that aggregate-level analyses could be misleading, as they may conflate the inherent growth potential of certain areas with the actual impact of policy changes. For instance, when state-level regulation changes result in branching deregulation, as seen in studies by Jayaratne and Strahan (1996) and Chava et al. (2013), positive economic outcomes might be incorrectly attributed to the policy change itself. Instead, banks may be strategically entering regions already on a growth trajectory, thus conflating the effects. My localized approach highlights the importance of considering these selection biases, providing a clearer picture of the true effects of banking deregulation on local economic growth.

4.3 Selection Effects: Bank Branches vs. Walmarts

Previous sections have established that banks select high growth areas for entry, but do not contribute much to growth themselves. This raises the question of just how good banks are at location selection, compared to other types of businesses. In this section, I explore this question by examining the total (selection and treatment) effects of businesses which are, arguably, unlikely to have a substantial treatment effect: Walmarts. I calculate the luminosity around newly entering Walmart locations, and compare the subsequent changes with those at extant post offices, as in previous sections. Because I do not observe planned Walmart locations but only those that are completed, the differences I calculate will combine both selection and treatment effects associated with new Walmarts.

Figure 13 shows the estimated total effects of Walmart (green) along with the selection effects associated with branch entry (blue) at various distances. As in Figure 12, the top left panel uses a simple circle of radius 5 km, and then from left-to-right the remaining panels extend the exercise using increasingly expansive rings, as described previously. The top panel shows that both effects are positive, and of a similar magnitude. As the rings become more expansive in the following panels, the total effect associated with Walmart entry then decreases whereas the selection effect of branch entry remains relatively stable. The bottom right panel reports the effects observed at the farthest distances, between 20 and 40 km, and in most years the bank selection effect is about twice the total effect of Walmart entry.

This exercise yields several insights. First, the comparable magnitudes of the bank branches' selection effect and Walmart's total effect, especially at the shortest distance of 5 km, serves to bolster the plausibility of the estimates in previous sections. Second, the fact that the Walmart effect dissipates at longer distances and banks' do not highlights their contrasting strategic focuses. For Walmart, the substantial local effect at short distances suggests a business model focused on capturing retail demand directly around each store. The diminishing effect at longer distances reflects their reliance on proximity to consumer traffic to drive profitability. Conversely, the sustained bank selection effect at greater distances indicates that banks are not only choosing high-growth areas but are also selecting on broader regional trends. This suggests that banks, with their focus on financial services and network-building, anticipate economic conditions over a larger area than Walmarts. Finally, the results of Figure 13 highlight the particular importance of location selection for banks at a regional level in comparison to other businesses. This emphasis further suggests, as in

Section 4.2, that studies using broad geographic shocks may be influenced more by selection effects in promising areas than by true causal effects of branch entry.

5 Conclusions

This paper takes a new angle in addressing a classic finance-growth topic: the impact of bank branching on economic growth. Using variation in completion across planned branches and granular proxies for economic growth, I disentangle selection and treatment effects at a local level. Using a series of outcome variables (nighttime light emissions, new business registrations, and SBA-7a loan amounts) I reveal a striking pattern: banks are highly selective in their branch placements, favoring areas with pre-existing growth potential, but their actual contribution to economic development is minimal.

The divergence between selection and treatment effects sheds light on economic aspects of branching that aggregate-level analyses. The extensive literature studying state-level shocks to branching regulations (Jayaratne and Strahan (1996), Jayaratne and Strahan (1998), Chava et al. (2013), Huang (2008), Berger et al. (2021)) may inadvertently conflate selection effects with the causal impacts of branch entry. By focusing on state-level shocks, these studies might overestimate the true economic benefits of deregulation, as the observed positive effects could be more reflective of the inherent growth potential of the selected areas rather than the impact of the bank branches themselves.

Ultimately, this research calls for a reassessment of how we measure and interpret the economic impacts of financial institutions. The dominant selection effect observed suggests that the placement of bank branches may not be as influential on local economic growth as previously thought. This emphasizes the need to critically evaluate the assumptions underlying aggregate-level analyses and consider the nuanced dynamics of financial development at a more granular level.

From a policy perspective, these findings imply that efforts to stimulate economic growth through increased bank branching may be less effective than anticipated. Policymakers should consider complementary strategies that address underlying economic conditions and support broader access to capital, particularly in areas not currently targeted by banks. Furthermore, future research should explore other dimensions of financial development and their localized impacts, potentially examining different types of financial institutions or services. By refining our understanding of these dynamics, policymakers and researchers can better discern the true effects of financial deregulation and craft more effective strategies to foster sustainable economic growth.

References

Andrews, RJ, Catherine Fazio, Jorge Guzman, Yupeng Liu, and Scott Stern, 2022, The startup cartography project: Measuring and mapping entrepreneurial ecosystems, *Research Policy* 51, 104437.

Anthony, Nicholas, 2024, No customers, no success: The postal banking failure exposed, *CATO Institute* .

Berger, Elizabeth A., Alexander W. Butler, Edwin Hu, and Morad Zekhnini, 2021, Financial integration and credit democratization: Linking banking deregulation to economic growth, *Journal of Financial Intermediation* 45, 100857.

Blevins, Cameron, and Richard W. Helbock, 2021, US Post Offices.

Butler, Alexander W., and Jess Cornaggia, 2011, Does access to external finance improve productivity? evidence from a natural experiment, *Journal of Financial Economics* 99, 184–203.

Callaway, Brantly, and Pedro H.C. Sant'Anna, 2021, Difference-in-differences with multiple time periods, *Journal of Econometrics* 225, 200–230, Themed Issue: Treatment Effect 1.

Chava, Sudheer, Alexander Oettl, Ajay Subramanian, and Krishnamurthy Subramanian, 2013, Banking deregulation and innovation, *Journal of Financial Economics* 109, 759–774.

de Chaisemartin, Clément, and Xavier D'Haultfœuille, 2020, Two-way fixed effects estimators with heterogeneous treatment effects, *American Economic Review* 110, 2964–96.

Elvidge, Christopher D., Daniel Ziskin, Kimberly E. Baugh, Benjamin T. Tuttle, Tilottama Ghosh, Dee W. Pack, Edward H. Erwin, and Mikhail Zhizhin, 2009, A fifteen year record of global natural gas flaring derived from satellite data, *Energies* 2, 595–622.

Freeman, Donald G, 2002, Did state bank branching deregulation produce large growth effects?, *Economics Letters* 75, 383–389.

Gilje, Erik P., Elena Loutskina, and Philip E. Strahan, 2016, Exporting liquidity: Branch banking and financial integration, *The Journal of Finance* 71, 1159–1184.

Goodman-Bacon, Andrew, 2021, Difference-in-differences with variation in treatment timing, *Journal of Econometrics* 225, 254–277, Themed Issue: Treatment Effect 1.

Henderson, J. Vernon, Adam Storeygard, and David N. Weil, 2012, Measuring economic growth from outer space, *American Economic Review* 102, 994–1028.

Huang, Rocco R., 2008, Evaluating the real effect of bank branching deregulation: Comparing contiguous counties across us state borders, *Journal of Financial Economics* 87, 678–705.

Jayaratne, Jith, and Philip E. Strahan, 1996, The finance-growth nexus: Evidence from bank branch deregulation, *The Quarterly Journal of Economics* 111, 639–670.

Jayaratne, Jith, and Philip E. Strahan, 1998, Entry restrictions, industry evolution, and dynamic efficiency: Evidence from commercial banking, *The Journal of Law Economics* 41, 239–274.

Kroszner, Randall S., and Philip E. Strahan, 1999, What Drives Deregulation? Economics and Politics of the Relaxation of Bank Branching Restrictions, *The Quarterly Journal of Economics* 114, 1437–1467.

Lucas, Robert E., 1988, On the mechanics of economic development, *Journal of Monetary Economics* 22, 3–42.

Nguyen, Hoai-Luu Q., 2019, Are credit markets still local? evidence from bank branch closings, *American Economic Journal: Applied Economics* 11, 1–32.

Office of the Comptroller of the Currency, 2023, Branch and Relocation Application.

Office of the Comptroller of the Currency Licensing Office, 2023, Private Email Exchange.

Petersen, Mitchell A., and Raghuram G. Rajan, 2002, Does distance still matter? the information revolution in small business lending, *The Journal of Finance* 57, 2533–2570.

Robinson, Joan, 1952, The rate of interest and other essays .

Schumpeter, Joseph A., 1911, The theory of economic development .

Sun, Liyang, and Sarah Abraham, 2021, Estimating dynamic treatment effects in event studies with heterogeneous treatment effects, *Journal of Econometrics* 225, 175–199, Themed Issue: Treatment Effect 1.

Zdrojewski, Anthony, and Alexander W. Butler, forthcoming, Are Two-Way Fixed-Effect Difference-in-Differences Estimates Blowing Smoke? A Cautionary Tale from State-Level Bank Branching Deregulation, *Critical Finance Review* 12.

Table 1 : County-Level GDP Growth on Lights: 2001-2013

Table 1 reports Fama-Macbeth style regressions of county-level GDP growth on growth in county lights. GDP growth is calculated contemporaneously (t), over the next year (t+1), and over the next five years (t+5). Table 2 reports the same, but GDP growth is now reported on a per capita basis. Data are 2001-2013, and the 50 brightest counties are excluded.

Table 1: County GDP Growth on Lights Growth

	GDP growth (t)	GDP growth (t+1)	GDP growth (t+5)
Lights Growth	0.03***	0.02***	0.06***
	(6.52)	(2.94)	(5.2)
Average R-squared	0.01	0.01	0.01

Table 2: County GDP Per Capita Growth on Lights Growth

	GDP growth (t)	GDP growth (t+1)	GDP growth (t+5)
Lights Growth	0.02***	0.01	0.02
	(4.02)	(0.74)	(1.09)
Average R-squared	0.01	0.01	0.01

Table 3 : Number of Withdrawn Branches and Donor Pool Observations with Lights Data: 1996-2006

The table reports the number of units from each year that are included in the difference-in-differences analyses (see Figure 9) using log-lights as an outcome variable. For each of the each of the 432 withdrawal events covered by the data, I estimate what the average treatment effect across these locations would have been had the branch been completed. The size of the donor pool for this counterfactual is given in the column "Comparison Locations". For consistency, I restrict the donor pool used to estimate the selection effect of each cohort to be the same size as that used to estimate the treatment effect.

Table 3: Number of Withdrawing Locations and Comparison Locations with Lights Data: 1996-2006

Year	Withdrawing Locations	Comparison Locations
1996	38	750
1997	32	864
1998	57	667
1999	41	714
2000	44	478
2001	33	531
2002	24	501
2003	36	532
2004	37	682
2005	39	604
2006	51	559
Total	432	6882

Table 4 : Number of Withdrawn Branches and Donor Pool Observations with New Business Registrant Data: 1992-2009

The table reports the number of units from each year that are included in the difference-in-differences analyses using the log of new business registrants as an outcome variable, which are reported in Figure 10. For each of the each of the 513 withdrawal events covered by the data, I estimate what the average treatment effect across these locations would have been had the branch been completed. The size of the donor pool for this counterfactual is given in the column "Comparison Zip Codes". For consistency, I restrict the donor pool used to estimate the selection effect of each cohort to be the same size as that used to estimate the treatment effect.

Table 4: Number of Withdrawing Zip Codes and Comparison Zip Codes with New Business Registrant Data: 1992-2009

Year	Withdrawing Zip Codes	Comparison Zip Codes
1992	11	157
1993	4	197
1994	14	290
1995	15	424
1996	42	494
1997	22	598
1998	61	450
1999	37	382
2000	38	318
2001	27	317
2002	20	312
2003	26	350
2004	28	433
2005	31	363
2006	38	350
2007	32	363
2008	36	299
2009	31	200
Total	513	6297

Table 5 : Number of Withdrawn Branches and Donor Pool Observations with SBA-7A Data: 1995-2017

The table reports the number of units from each year that are included in the difference-in-differences analyses (see Figure 11) using SBA-7a lending to borrowers in the entered zip code as an outcome variable. For each of the each of the 271 withdrawal events covered by the data, I estimate what the average treatment effect across these locations would have been had the branch been completed. The size of the donor pool for this counterfactual is given in the column "Comparison Zip Codes". For consistency, I restrict the donor pool used to estimate the selection effect of each cohort to be the same size as that used to estimate the treatment effect.

Table 5: Number of Withdrawing Zip Codes and Comparison Zip Codes with SBA-7a Data: 1995-2017

Year	Withdrawing Zip Codes	Comparison Zip Codes
1995	7	135
1996	21	172
1997	11	239
1998	34	176
1999	15	148
2000	18	150
2001	7	149
2002	9	121
2003	13	175
2004	14	225
2005	14	171
2006	19	152
2007	11	197
2008	22	157
2009	12	115
2010	5	89
2011	4	108
2012	6	69
2013	7	61
2014	10	52
2015	9	30
2016	3	34
Total	271	2925

Figure 1 : Number of Branches and Banks by Regulator

The figure displays the number of banks (top), branches (middle), and deposits (bottom) that are under regulation by each of the major regulators each year from 1994 to 2023. In each plot, the bottom bar (blue) represents the OCC.

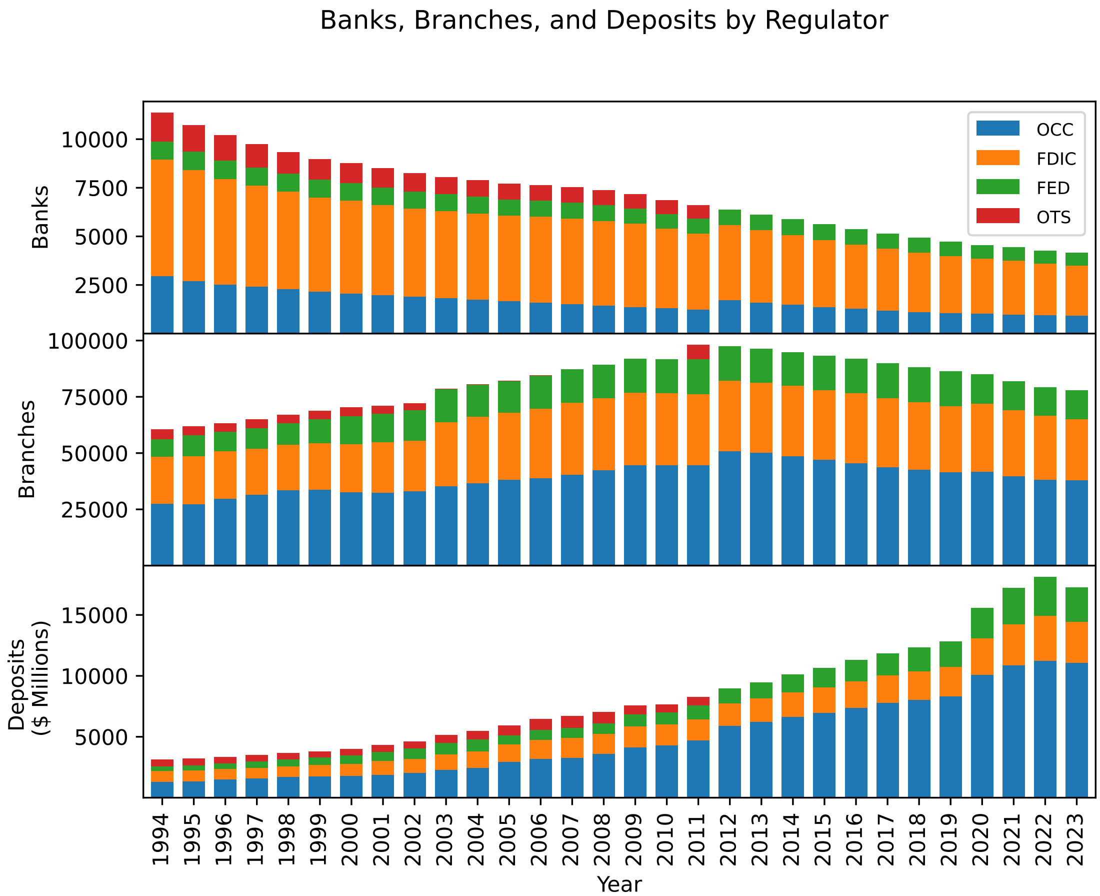

Figure 2 : Example Fulfilled Application - Newberry, Florida

The figure displays an example entry from the CAS detailing information on a completed branch in Newberry, Florida. The top section, "Details For OCC Control Number: 2000-SE-05-0027", contains basic information about the applying bank, including its name, charter number, and headquarters location. Under "Proposed Branches:" the system provides the proposed location of the new branch. "Public Comment Information:" details the start and end dates between which the public could comment on the proposed branch. "Filing status:" documents the receipt, approval, and effective dates of the branch's application. The aerial view of the branch, and a demonstration of the area over which light emissions are calculated in Figure 7.

Corporate Applications Search Result Details

Details For OCC Control Number: **2000-SE-05-0027** Return to List

Application Type:	Branch Establishment
Transaction Form:	Branch Establishment - Staffed Branch
Bank:	**The First National Bank of Alachua**
Charter/License #:	**8980**
Bank Headquarters Location:	15000 NW 140TH STREET ALACHUA , FL 32616 County: Alachua

Proposed Branches:

Branch Name	Street Address	Suite	City	State	Zip	County	Certification #
NEWBERRY BRANCH	24202 WEST NEWBERRY ROAD, SUITE F		NEWBERRY	FL	32669	Alachua	117659A

Public Comment Information:

Comment Period Start Date	Comment Period End Date	Adjusted Period Start Date	Adjusted Period End Date	OCC Contact
2000-02-17	2000-03-17			Southeast District Office Contact Info

Filing Status:

Action	Date
Receipt	2000-02-15
Approved	2000-03-23
Consummated/Effective	2000-05-08

Figure 3 : Example Withdrawn Application - Ocoee, Florida

The figure displays an example entry from the CAS detailing information on a withdrawn branch in Ocoee, Florida. The top section, "Details For OCC Control Number: 1999-ML-05-0220", contains basic information about the applying bank, including its name, charter number, and headquarters location. Under "Proposed Branches:" the system provides the proposed location of the new branch. "Public Comment Information:" details the start and end dates between which the public could comment on the proposed branch. "Filing status:" documents the receipt, approval, and withdrawal dates of the branch's application.

Corporate Applications Search Result Details

Details For OCC Control Number: **1999-ML-05-0220** Return to List

Application Type: Branch Establishment
Transaction Form: Branch Establishment - Staffed Branch
Bank: **FIRST UNION NATIONAL BANK**
Charter/License #: **1**
Bank Headquarters Location: FIRST UNION PLAZA
CHARLOTTE , NC 28288
County: Mecklenburg

Proposed Branches:

Branch Name	Street Address	Suite	City	State	Zip	County	Certification #
CLARKE ROAD BRANCH	VICINITY OF SE SILVER STAR AND CLARKE ROAD		OCOEE	FL	32701	Orange	116857A

Public Comment Information:

Comment Period Start Date	Comment Period End Date	Adjusted Period Start Date	Adjusted Period End Date	OCC Contact
1999-09-28	1999-10-27			ML Contact Info

Filing Status:

Action	Date
Receipt	1999-09-28
Approved	1999-10-29
Withdrawn	2000-10-27

Figure 4 : Branch Establishment Application Process

The figure visualizes the portion of branch establishment applications that reach each stage, as well as the amount of time that typically passes between each stage. The process flows from left to right starting with the "Applications" node, which indicates that there are 42,396 applications in the OCC's CAS data. Subsequent nodes report the percentage of these applications that reach the corresponding step (e.g. 96.6% of applications' documentation begins with the OCC receiving the application, whereas the remaining 3.4% of applications are data errors). Each node is scaled in proportion to the number of applications that it represents, and where appropriate the passage of time between steps is indicated at the bottom of the figure.

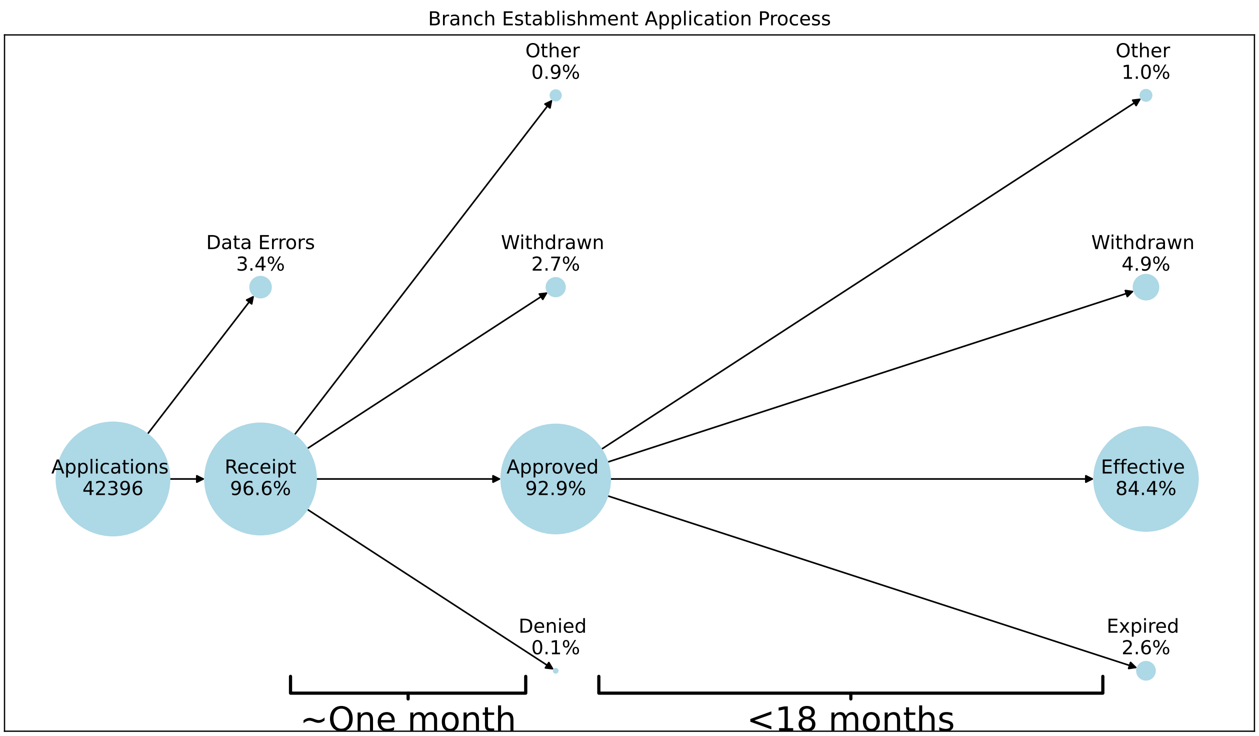

Figure 5 : Number of Bank Branch Applications and Percentage Unfulfilled Over Time

The figure displays the number of branch applications each year (green, solid line) as well as the percentage which are not fulfilled (blue, dashed line).

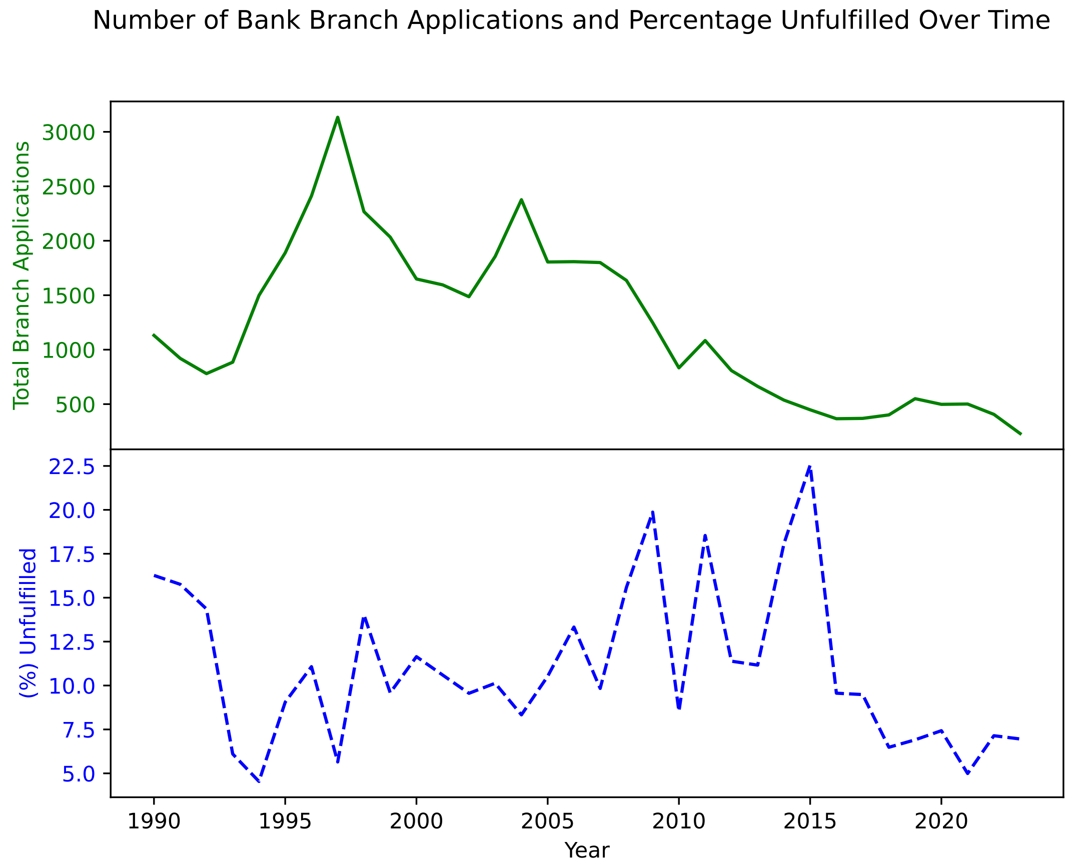

Figure 6 : Locations of First Withdraws

The figure displays the location of each withdrawal that is used in the main analysis (which uses luminosity as an outcome variable).

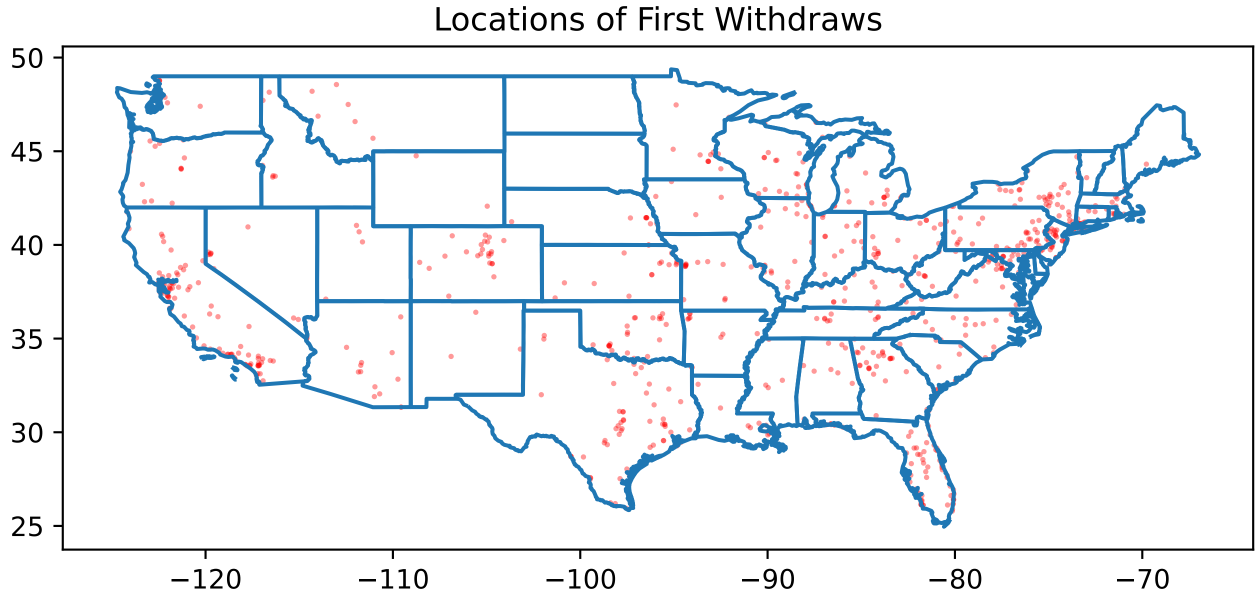

Figure 7 : Newberry, Florida, Aerial View

The figure uses an aerial view of the Newberry, Florida branch from Figure 2 to portray luminosity calculation. The branch is indicated in the center of the map by a black dot. The two concentric circles have radii of 5 and 10 kilometers. I calculate the average light value for both the pink-shaded inner circle, as well as the blue-shaded doughnut/ring shape. Though additional rings are not shown in the picture, I do the same for rings with inner/outer radii pairs of 10/20 and 20/40 km. I employ the same approach to calculate lights around branches that are completed or withdrawn, and around post offices as a reference point.

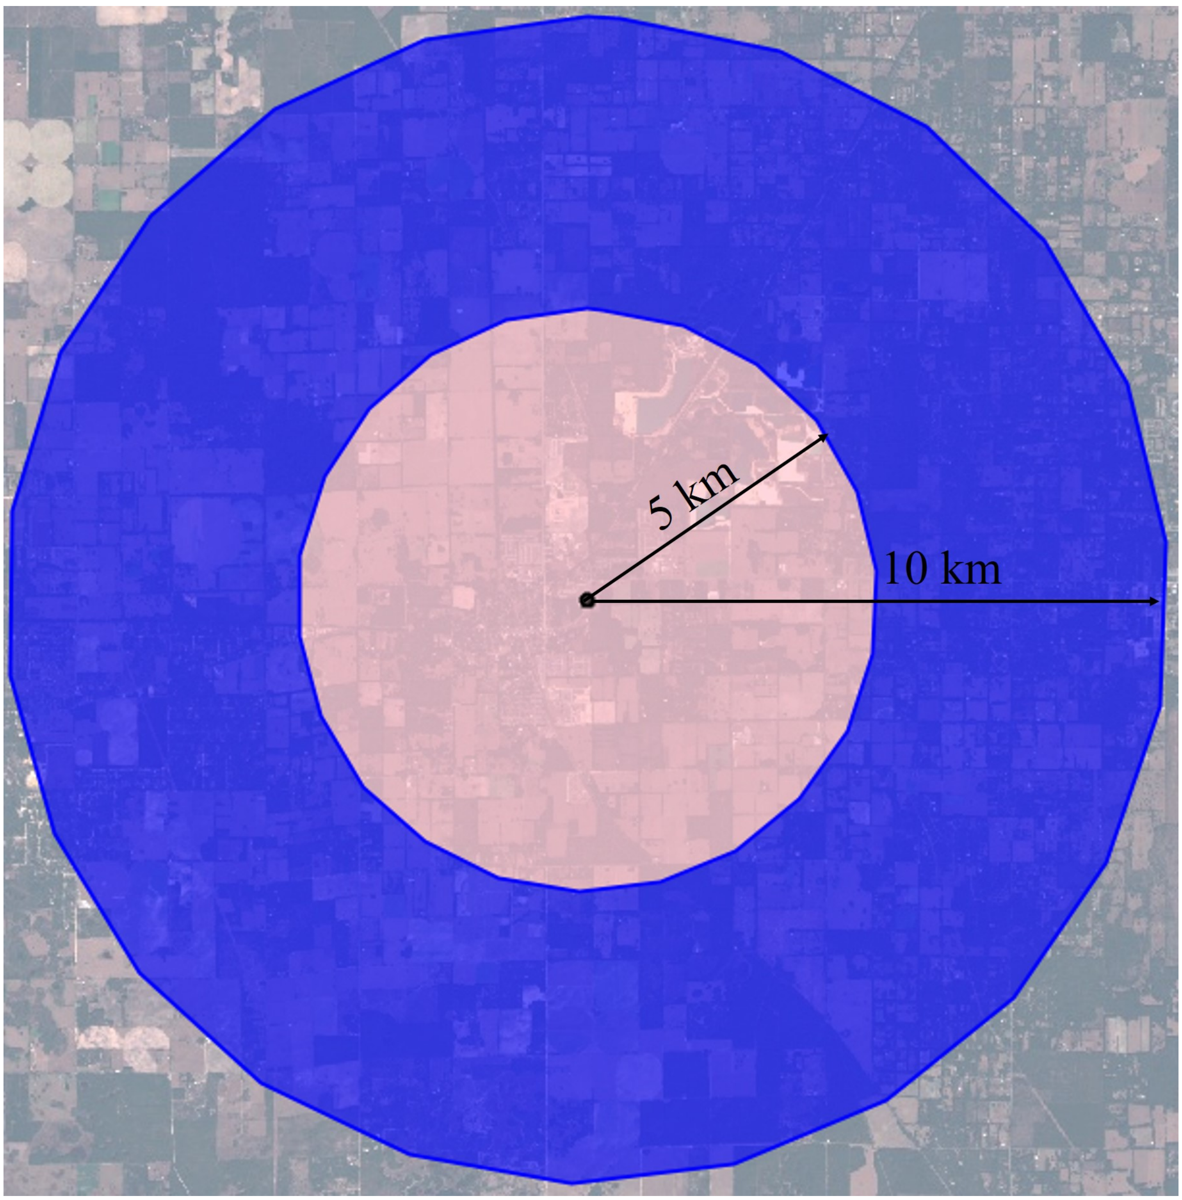

Figure 8 : County-Level Growth in GDP Per Capita vs Growth in Lights (2001-2013)

The figure displays the cross-sectional relationship between county growth in per capita GDP over 2001-2013 and county growth in lights over the same period. Real GDP growth per capita is reported on the vertical axis, with growth in lights reported on the horizontal axis. Each blue dot represents an individual county. Best fit lines are estimated using both OLS (blue line) and median-regression (orange-line) approaches, with these regressions reported below the plot. Extreme outlier counties are labeled for context.

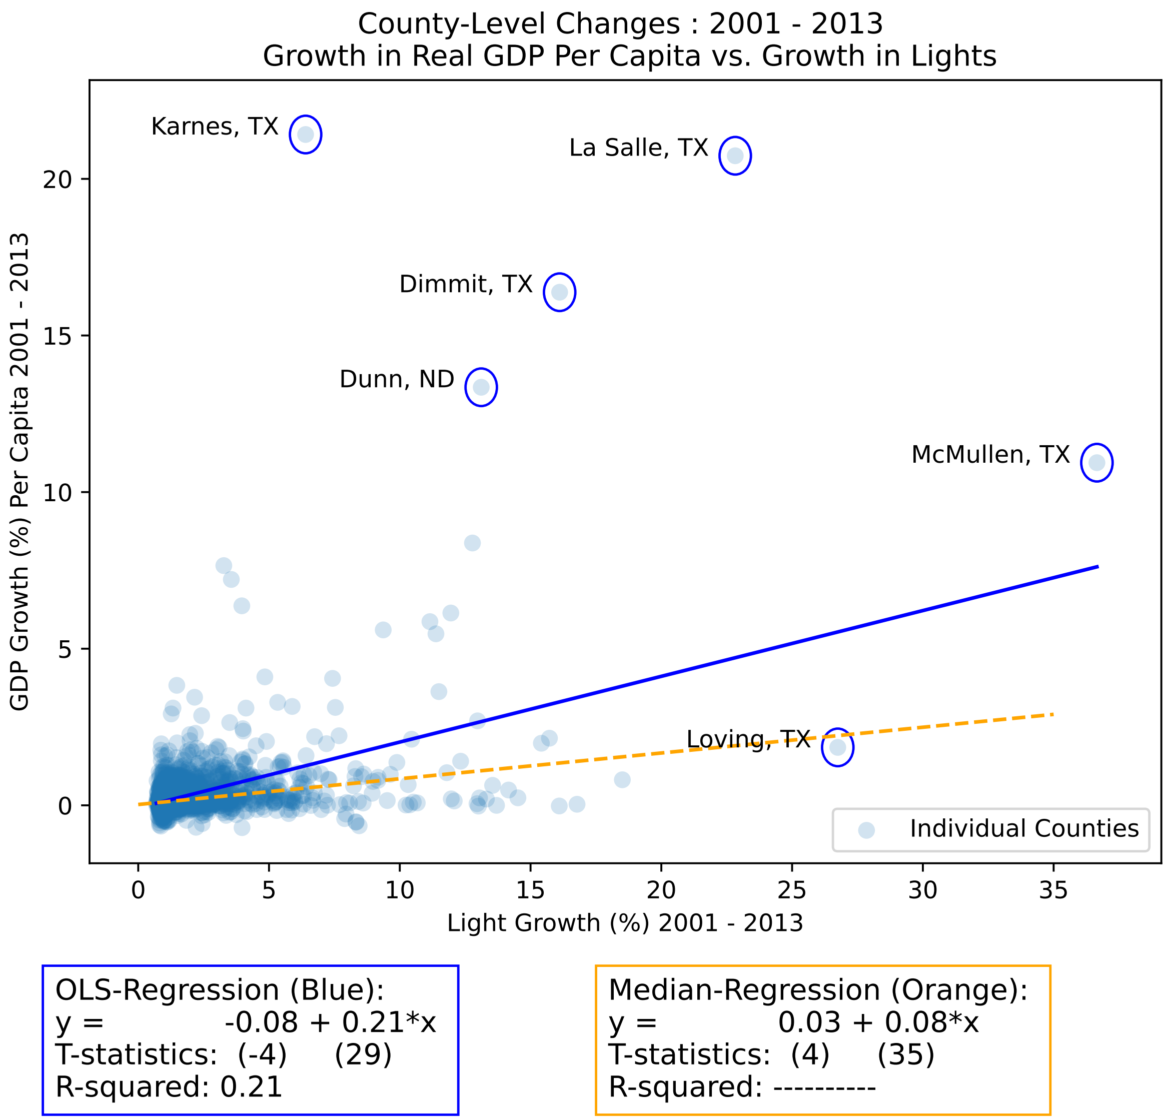

Figure 9 : Selection and Treatment Effects of Opening a Branch on Log-Lights Within 5km

The figure displays estimated selection and treatment effects of branch entry on the log-lights around the branch's location. Effects are estimated using a modification of the Callaway and Sant'Anna (2021) approach. Selection effects compare changes around withdrawing locations to those around post offices, whereas treatment effects compare opening locations to withdrawing locations of the same cohort. Lights are calculated as the mean across pixels within five kilometers of the withdrawn branch. The values of yearly effects are reported on the left-hand axis. 95 and 90% confidence intervals estimated via bootstrapped standard errors.

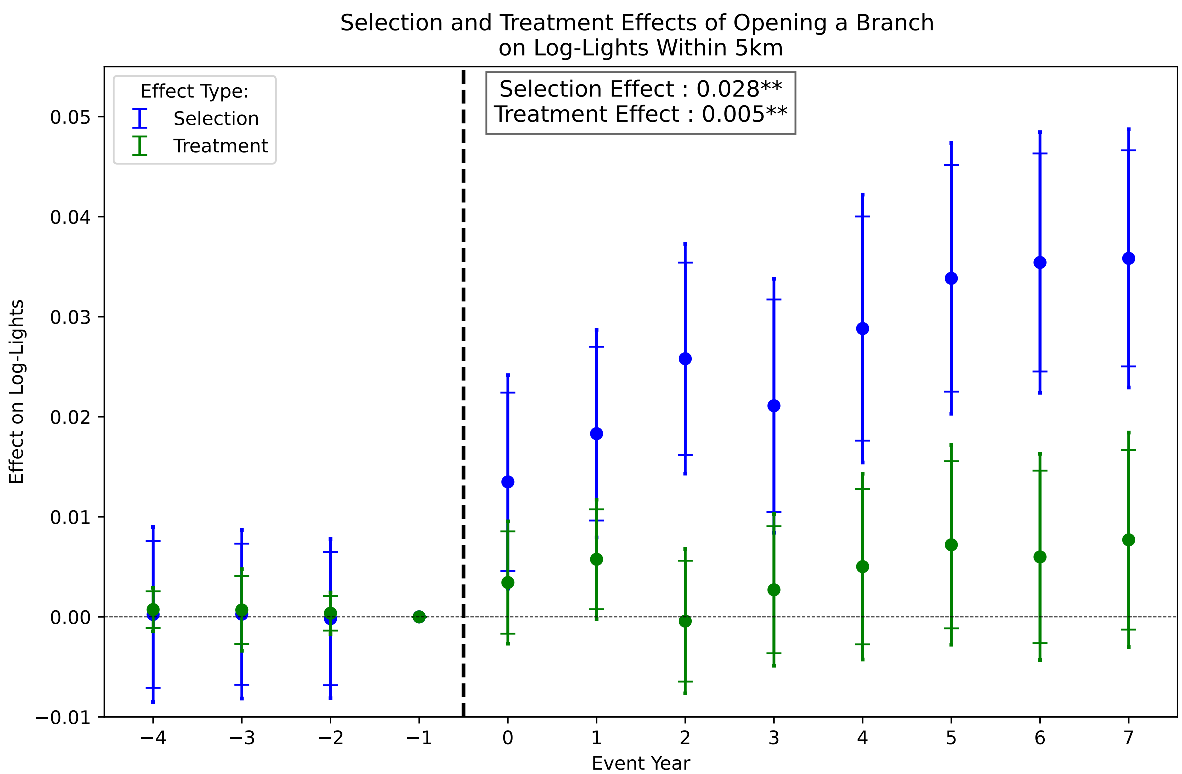

Figure 10 : Selection and Treatment Effects of Opening a Branch on Log of New Business Registrants in that Zip code

The figure displays estimated selection and treatment effects of branch entry on the log-number of new business registrants in the entered zip code. Effects are estimated using a modification of the Callaway and Sant'Anna (2021) approach. Selection effects compare changes in withdrawing zip codes to all other zip codes, whereas treatment effects compare opening zip codes to withdrawing zip codes of the same cohort. The values of yearly effects are reported on the left-hand axis. 95 and 90% confidence intervals estimated via bootstrapped standard errors wherein zip codes are randomly assigned as opening/withdrawing and the procedure repeated. Cumulative effects since period 0 are indicated by the red line, with values reported along the right-hand axis and individual markers indicating statistical significance, where applicable.

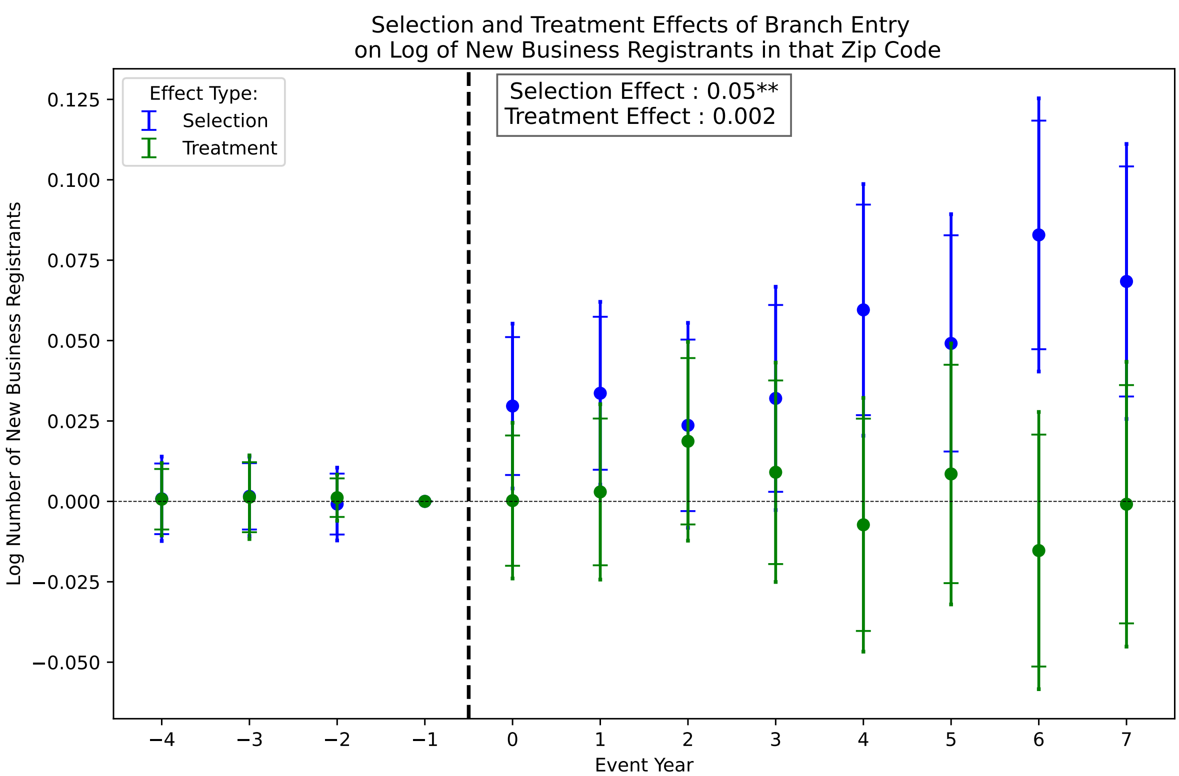

Figure 11 : Selection and Treatment Effects of Opening a Branch on Log of SBA-7a Loan Amounts Issued to Borrowers in that Zip code

The figure displays estimated selection and treatment effects of branch entry on the log of SBA-7a loan amounts issued to borrowers in the entered zip code. Effects are estimated using a modification of the Callaway and Sant'Anna (2021) approach. Selection effects compare changes in withdrawing zip codes to all other zip codes, whereas treatment effects compare opening zip codes to withdrawing zip codes of the same cohort. The values of yearly effects are reported on the left-hand axis. 95 and 90% confidence intervals estimated via bootstrapped standard errors wherein zip codes are randomly assigned as opening/withdrawing and the procedure repeated. Cumulative effects since period 0 are indicated by the red line, with values reported along the right-hand axis and individual markers indicating statistical significance, where applicable.

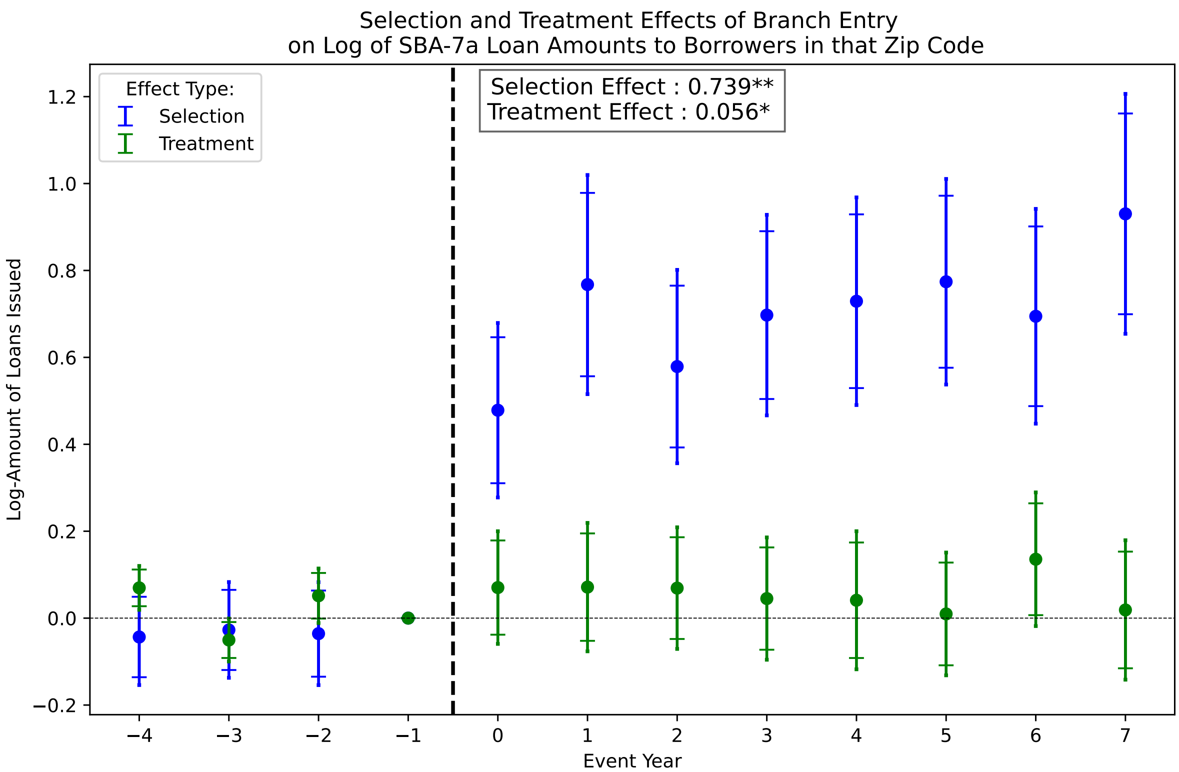

Figure 12 : Selection and Treatment Effects of Opening a Branch on Log of New Business Registrants in that Zip code

The figure displays estimated selection and treatment effects of branch entry on the log-number of new business registrants in the entered zip code. Effects are estimated using a modification of the Callaway and Sant'Anna (2021) approach. Selection effects compare changes in withdrawing zip codes to all other zip codes, whereas treatment effects compare opening zip codes to withdrawing zip codes of the same cohort. The values of yearly effects are reported on the left-hand axis. 95 and 90% confidence intervals estimated via bootstrapped standard errors wherein zip codes are randomly assigned as opening/withdrawing and the procedure repeated. Cumulative effects since period 0 are indicated by the red line, with values reported along the right-hand axis and individual markers indicating statistical significance, where applicable.

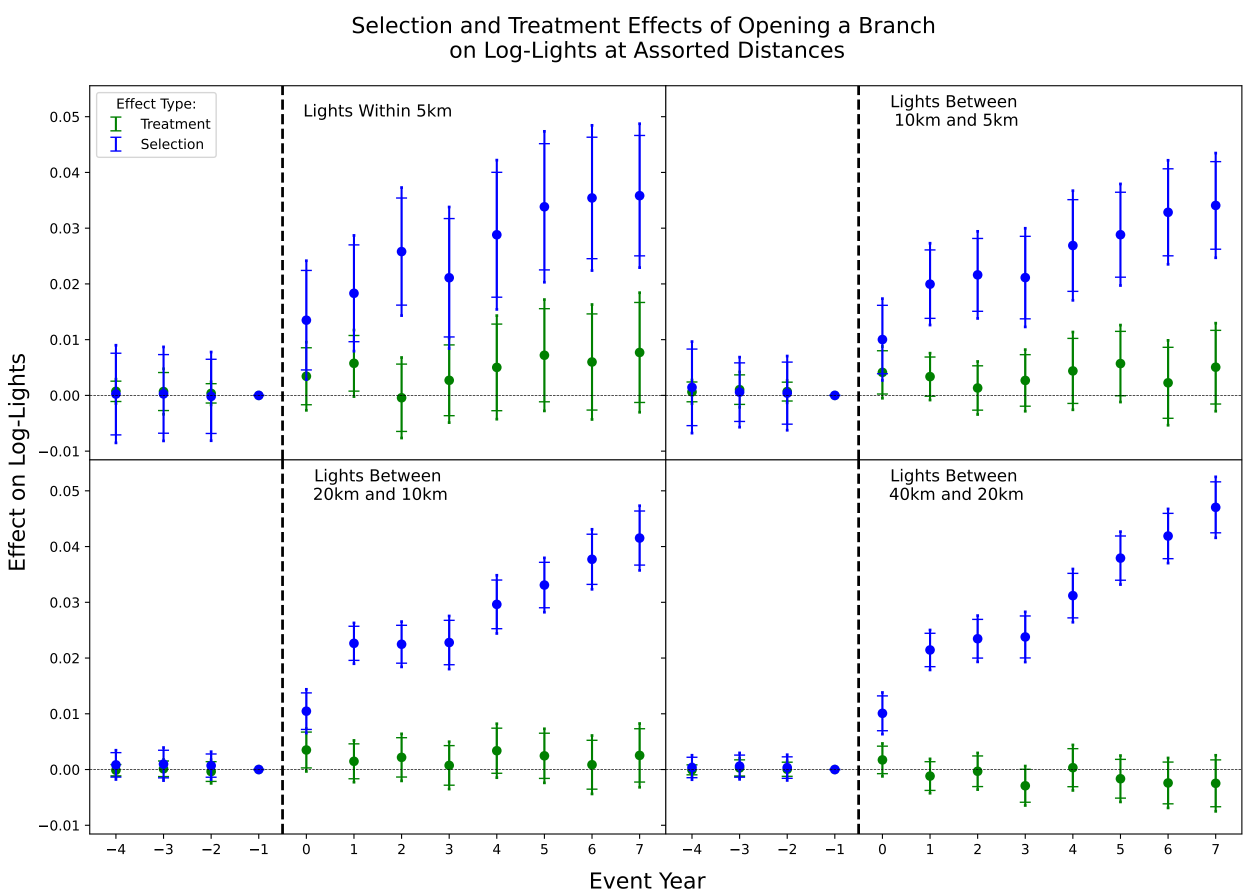

Figure 13 : Selection and Treatment Effects of Opening a Branch on Log of New Business Registrants in that Zip code

The figure displays estimated selection and treatment effects of branch entry on the log-number of new business registrants in the entered zip code. Effects are estimated using a modification of the Callaway and Sant'Anna (2021) approach. Selection effects compare changes in withdrawing zip codes to all other zip codes, whereas treatment effects compare opening zip codes to withdrawing zip codes of the same cohort. The values of yearly effects are reported on the left-hand axis. 95 and 90% confidence intervals estimated via boot-strapped standard errors wherein zip codes are randomly assigned as opening/withdrawing and the procedure repeated. Cumulative effects since period 0 are indicated by the red line, with values reported along the right-hand axis and individual markers indicating statistical significance, where applicable.