The Local Multiplier Effect of Pensions: Evidence from Brazil*

Rodrigo Toneto†

July 2024

PRELIMINARY DRAFT

Abstract

Pensions are the most prevalent form of transfer in most nations. While the impact on recipients' households has been extensively studied, little is known about their aggregate effects on local economies. This question is particularly relevant in contexts of high informality, where generous non-contributory systems can reduce the premium of formal sector jobs and shift labor supply to the informal sector. At the same time, by providing liquidity-constrained households with a higher and more stable source of income, an increase in pension coverage boosts local demand. Leveraging administrative data on the universe of pensioners and formal employment in Brazil, this paper exploits the quasi-random availability of pension agency offices in Brazilian municipalities to identify how increasing pension transfers affects local labor markets' dynamics. Establishing a new office boosts pension take-up by 4.3%, with the increase concentrated in non-contributory benefits. Consistent with local multiplier effects, expanding total pension payments results in 4.2% more jobs in the private formal sector. To further understand the demand component of the employment effects, I document the causal shift in consumption triggered by retirement using household budget surveys. After receiving a pension transfer, households increase consumption mainly in formal stores with an implied marginal propensity to consume of 0.9. Coupled with a null effect on wages at the aggregate level, this suggests that lack of labor demand was the primary constraint to employment growth in the formal sector.

Keywords: Pensions system, informality, local multiplier effects, Brazil

*I would like to thank Lucie Gadenne, François Gerard, Felipe Gonzales, Marco Manacorda, Pedro Souza, Andrea Tesei, and participants of the Applied and Development seminars at Queen Mary University for their comments and suggestions. I also thank the INSS staff, particularly Marcia de Souza, for ensuring data access and for carefully discussing the context of the policy studied in this paper.

†Queen Mary University of London. Email: r.toneto@qmul.ac.uk

1 Introduction

The United Nations projects that in most low- and middle-income countries, the population aged 65 and older will at least double by 2050 (United Nations, 2019). Consequently, public pension expenditures, that represents the most substantial income transfers and a key component of social spending in many countries, are expected to rise even further (OECD, 2023). This financial strain is exacerbated in environments where informal and rural employment are prevalent, resulting in a smaller tax base and a more significant portion of the population depending on non-contributory or partially contributory schemes (Guven et al., 2021). While much research has examined the effects of pensions on affected households (e.g., Case and Deaton, 1998; Duflo, 2000, 2001; Joubert, 2015), the aggregate impact of increasing pensions transfers remains unexplored. This question is particularly relevant in contexts of high informality, where the generosity of non-contributory schemes can reduce the premium of formal sector jobs and favor informality (Ulyssea, 2020). Conversely, by providing older workers with a higher and more stable source of income than average labor earnings, higher amount of transfers might increase formal jobs through local multiplier effects (Corbi et al., 2019; Egger et al., 2022; Gerard et al., 2021). The net formalization effect is more likely to be positive if the constraint on formal job creation stems more from labor demand than from labor supply.

This paper studies how expanding pension payments, primarily from non-contributory schemes, affects labor market outcomes at the municipal level. I empirically evaluate this question in the context of Brazil, a developing economy with pervasive informality that allocates 10% of its GDP to its pension system. Using two comprehensive administrative datasets on the universe of pensioners and formal employment from 2007 to 2018, I leverage the quasi-random establishment of new pension agency offices to identify the causal effects of increased pension registration on local labor markets. The availability of an office induces an increase in pension take-up at the municipal level, mainly expanding the number of new non-contributory pensions. At the same time, I document an increase in formal private employment, particularly within the services sector and larger firms. Despite the rise in employment, wages remain constant. Combined, these results suggest positive local multipliers for the pension transfers and indicate that a lack of labor demand was the main factor preventing an increase in formal employment. This underscores the importance of considering aggregate effects when examining the labor market impacts of pension systems.

In 2011, the Brazilian National Pensions System Institute (INSS) launched an expansion plan for its agency network. The goal was to establish 651 new offices within the next four years, aiming for 1,843 units across 1,685 of the 5,570 Brazilian cities. I accessed the list of targeted municipalities and combined it with the data on agency opening dates. The key feature of the expansion policy is that, as described by the INSS, the target municipalities constitute the cities that, according to the institute's metrics, should have had an agency by 2010 but did not. My empirical strategy relies on the premise that by reducing commuting and informational costs, establishing a new agency induces an increase in pension take-up at the municipal level. I exploit the variation in treatment timing to identify the causal effect of the associated increase in pension payments on labor market dynamics.

A fast-growing econometric literature documents that using conventional two-way fixed effect (TWFE) regressions might produce misleading estimates of treatment effect (see De Chaisemartin and d'Haultfoeuille (2023) and Roth et al. (2023) for a comprehensive review of the new available methods). The central issue is that in the presence of heterogeneous effects across groups and time periods, the aggregation method and the implied weighting scheme are not irrelevant to the correctness of the estimates. In settings with staggered treatment adoption and heterogeneous cohort sizes, as is the case here, Sun and Abraham (2021) discuss that TWFE can be particularly misleading. To address these concerns, my main empirical strategy relies on the estimator proposed by Callaway and Sant'Anna (2021). The method allows me to compare municipalities treated earlier on with those treated in the future, or never treated, using a less stringent parallel trend assumption (De Chaisemartin and d'Haultfoeuille, 2023).

I show that establishing an INSS facility increases pension take-up by up to 6% after four years. The differential trend in the growth of the number of benefits triggered by the arrival of a new office allows me to causally identify the effect of increased pension income on local labor market outcomes. My results suggest that an overall 4.3% increase in pension payments after the agency's arrival leads to a rise in full-time private formal employment of roughly the same magnitude. I combine these two estimates to recover an implied cost per job of 7,200 US dollars or 2.7 minimum wages, strikingly similar to what other papers found studying transfers to Brazilian municipalities (Corbi et al., 2019; Gerard et al., 2021).

The observed change in formal employment is concentrated in the service sector, where employment increases by 5.1% after treatment. Importantly, I observed no increase in public employment, which can be explained by the relatively small size of the average agency in our sample (2.7 employees). This is particularly relevant to guarantee that a new agency only affects the local labor market through increased pension transfers. Additionally, I document that large firms are the ones that accommodate the new employment flow. The number of large firms increases by nearly 5% and the share of employment at the top decile of firms distribution by 1.14 percentage points. Finally, we don't observe any significant effect on average wages.

My results are robust across various specifications and methods. Although different differences-in-differences strategies require distinct identifying assumptions, the results hold under alternative choices of the control group and across estimations methods. The same patter of effects emerge when using a set of predetermined covariates.

This paper contributes to the long-standing literature on the aggregate impact of fiscal stimulus in economics. Recently, there has been a surge in studies assessing the size of fiscal multipliers using cross-sectional variation in government expenditures (e.g., Nakamura and Steinsson, 2014; Serrato and Wingender, 2016; Auerbach et al., 2019). For a comprehensive review, see Chodorow-Reich (2019), as well as Farhi and Werning (2016) for an in-depth theoretical discussion. This literature has predominantly focused on developed countries and on expenditures allocated to local governments rather than direct income transfers to households. Few studies have extended this literature to developing countries or direct income transfers. Kraay (2012, 2014) show that the multiplier effect of different sources of external lending (World Bank and official creditors) to low-income countries is small, around 0.5. Corbi et al. (2019) use discontinuities in allocating federal transfers to estimate a local multiplier of 2 for Brazilian municipalities. Regarding the multiplier effect of direct income transfers, Egger et al. (2022) study a randomized one-off transfer and find a significant impact on consumption with an implied multiplier of 2.5. Gerard et al. (2021) estimate the multiplier effect of a large-scale expansion of the well-known *Programa Bolsa Família* (PBF) in Brazil. They find that municipalities more affected by the policy experience an increase in formal employment in the private sector, with an implied multiplier of 2.8 and a cost-per-job of 2.1 minimum wages My paper aims to extend this literature in three main dimensions.

First, I contribute to this literature by studying the most common type of transfer globally, which is likely to represent an increasing share of total government expenditures in the developing world. Pensions are particularly interesting for this setting not only due to their pervasiveness in most nations but also because they connect with one of the central challenges in economics: how to organize intergenerational transfers that can sustain a minimum level of welfare across life. An extensive literature explores the impact of pensions on household wellbeing and labor supply in developing countries (e.g., Case and Deaton, 1998; Duflo, 2000, 2003; de Carvalho Filho, 2008; Ardington et al., 2009; Kaushal, 2014; Galiani et al., 2016; Ceni, 2017; Becerra, 2017; Huang and Zhang, 2021). To the best of my knowledge, this is the first paper assessing pensions' multiplier effects.

The nature of pensions as a transfer is also very different from conditional cash transfers (and both clearly differ from intra-government transfers as they directly channel resources to households). Compared to cash transfers, pensions target older people, provide a larger average benefit amount—at least in the Brazilian case—and are entitlements. Once eligible, an individual is guaranteed to receive the pension after registering, and there is no risk of it being discontinued. These differences lead to potentially different responses regarding labor decisions and consumption. The entitled nature of pensions can potentially lead to different labor market responses as any reduction or reallocation of labor supply is entirely driven by income effects. On the consumption side, older people are assumed to typically have a smaller propensity to save. Also, because pension income is stable, recipients can use it as loan collateral, which might induce a more robust consumption response. However, because the amount received is bigger and assuming the marginal propensity to consume declines with income, the total increase in aggregate consumption might be smaller. The fact that my implied cost-per-job is statistically indistinguishable to the one found by Gerard et al. (2021) in the same context around the same period but for conditional cash transfers suggests that the recipients might use those payments similarly.

The second dimension to which this paper adds is by providing evidence on how fiscal stimuli through increased pension transfers interact with labor market allocation in a context of high informality. In the developing world, more than half of total employment is in the informal sector (Bonnet et al., 2019). Informality accounts for nearly 55% of the total jobs in target municipalities at baseline, and non-contributory and partially contributory pensions constitute 76% of total benefits in the sample. A significant literature seeks to understand the interaction of different social insurance schemes on individual choices across formal and informal sectors (Joubert, 2015; Joubert and Kanth, 2022; Canelas and Nino Zarazua, 2022; Cruces and Bérgolo, 2013; McKiernan, 2021; Gerard and Gonzaga, 2021; Finamor, 2022; Delalibera et al., 2023). Non-contributory pension schemes are known to reduce the formal job premium, particularly for low-wage occupations, which can, in turn, push marginal workers towards the informal sector (Ulyssea, 2020). By making the existence of such non-contributory schemes more salient and their access less costly, the opening of the agency could lead marginal workers to opt for the informal sector. However, our findings suggest that demand effects dominate, at the aggregate level, individual disincentives to formalize associated with a more generous system.

The push towards formalization may be significantly influenced by non-homotheticity in demand, which prompts a higher share of formal consumption as income increases, as posited by Bachas et al. (2023). My results provide suggestive evidence in favor of this channel. Using household survey budgets, I document that a quasi-random increase in pension payments around the pension age eligibility cut-off increases consumption with an implied marginal propensity to consume of 0.95, entirely driven by the rise in formal store consumption.

My third main contribution lies in providing a causal identification of how increased disposable income for a particular demographic group influences employment allocation across sectors and firms of different sizes. Income effects significantly shape consumption patterns and affect structural change (Matsuyama, 2019; Buera and Kaboski, 2009; Comin et al., 2021; Fan et al., 2023; Gollin and Kaboski, 2023). Specifically, I relate to research exploring the rise of the service sector, such as Buera and Kaboski (2012) and Fan et al. (2023). A closely related work is Hackmann et al. (2023), which demonstrates how subsidies for long-term care insurance for older adults in Germany stimulate demand for labor-intensive care activities and increase employment in those sectors. Demand constraints are also likely to prevent more productive sectors from expanding (Murphy et al., 1989; Goldberg and Reed, 2023). The fact that wages did not increase even though employment grew suggests that the labor supply for formal jobs is relatively elastic and the lack of labor demand was limiting formal employment growth in the target municipalities. This also relates to the old development tradition of growth with perfect elastic labor supply *à la* Lewis (1954). If the utilization rate of production factors is slack, a surge in demand due to higher public expenditure can be met without price increases (Michaillat and Saez, 2015, 2019). The existence of slackness in factor utilization is also what Egger et al. (2022) argue to rationalize the large multiplier effects in their cash transfer experiment.

The paper proceeds as follows: Section 2 discusses the institutional context and data. Section 3 presents the empirical strategy. Section 4 discusses the results for the aggregate level impact of the pension expansion. Section 5 shows how pensions' impact at the household level can rationalize what is found at the aggregate level. Section 6 concludes.

2 Context and data

2.1 Pensions system in Brazil

The Brazilian pension system combines contributory and non-contributory schemes with distinct rules for urban and rural areas. In the contributory scheme, there are specific regulations for certain public service careers. The system also provides coverage for workers who are permanently incapacitated because of illness or accident. Additionally, spouses or children can access survivor pensions if their partner or parent has died. The duration of these benefits varies based on the age of the dependent and the number of months the deceased insured person had contributed.

Before the implementation of the online platform "INSS Digital" in late 2017, the process of obtaining an INSS pension in Brazil involved several steps, requiring physical presence and substantial paperwork. Applicants had to visit an INSS agency to initiate their application. The procedure began with scheduling an appointment, which could be done either by phone through the INSS service line or by visiting the agency directly. During the appointment, applicants had to present various documents, including identification, proof of contribution periods, and other relevant paperwork depending on the type of pension sought (e.g., retirement by age, disability, survivor's pension).

This process often involved multiple visits to the agency, as initial appointments were used to verify the documentation and provide a formal list of any additional documents needed. Applicants were required to fill out detailed forms and submit them along with their documents for processing. The INSS officials would then review the submission, a process that could take several weeks. During this period, the potential beneficiaries could track the status of their request by revisiting the agency or through follow-up phone calls.

For rural workers or those seeking disability pensions, additional documents, such as medical reports or certifications from rural workers' unions, were necessary, adding to the complexity and length of the process. The in-person nature of the process, coupled with the bureaucratic requirements, often led to significant barriers to accessing the system.

In this paper, I use the cost reduction associated with establishing a new social security facility in a given municipality as a source of variation in pension take-up. Here, cost can be broadly understood to include monetary, time, and informational costs associated with obtaining benefits. In what follows, I detail the policy that expanded the agency network, allowing me to exploit quasi-experimental variation in the timing and location of new openings.

2.2 Expansion Plan - PEX

According to pension system regulations, the decision to open an agency should be based on the municipality's total population, potential demand for benefits, and the distance to an existing facility. In practice, however, these criteria are weighted in ways that make it difficult to predict which municipalities will be selected. Therefore, simply comparing cities with or without an agency would still be affected by selection bias. To address this issue, I ideally compare a municipality that was due to receive an agency but did not, with one that was targeted and did receive an agency.

In the country's 2012-2015 multi-year strategic plan (PPA 2012-2015) launched in 2011, the INSS set the goal to deliver 651 new agencies over the next four years. With the expansion, the aim was to have an agency in every municipality with population above

20.000, spreading the network to 1,685 Brazilian cities. The INSS gave me access to the list of targeted municipalities and the opening dates for each agency. From these data, I observed that of the 651 proposed, 28 agencies were on the plan but were already delivered when it was launched, 288 were built between 2012 and 2016, while 335 were not delivered.

On the demand side, to qualify for receiving an agency, selected municipalities had to donate a plot of land meeting specific criteria: a minimum area of 1,000 square meters, a flat surface, and a central, easily accessible location. The donation process involved several steps: the mayor proposed the land, INSS engineers conducted a preliminary inspection to ensure it met construction requirements, the mayor submitted a land donation bill to the municipal council, and upon the bill's approval, the donation proposal was sent to INSS for authorization. Once authorized, the bidding process for the construction work began, followed by the construction itself. On the supply side, the central government followed the queue of requests to set priorities among the listed municipalities. Each year, they would initiate new construction projects depending on the budget available to the INSS agency.

2.3 Data

These are the main data sets used in this research:

- **RAIS** The RAIS (Relação Anual de Informações Sociais) is an employer-employee administrative dataset collected annually by the Brazilian Ministry of Labor. It provides comprehensive data on formal employment across all sectors of the Brazilian economy.The RAIS dataset includes detailed information on employees, such as their demographic characteristics, job tenure, wages, and occupation. It also covers firm-level details like industry classification and location.

- **INSS payment sheets**: administrative records maintained by the INSS, updated using the same type of information (social security numbers) as the widely used RAIS. These records provide detailed information on the total payments and benefits disbursed in each category of pension. The payment sheets cover all Brazilian municipalities and encompass yearly data from the year 2007 to 2018.

- **CENSO 2010** The 2010 Census dataset in Brazil offers detailed information at the household level. It allows me to recover municipal-level covariates such as informality levels, age profiles, income, and education.

- **POF** The POF (Pesquisa de Orçamentos Familiares) is a household budget survey conducted periodically by the Brazilian Institute of Geography and Statistics (IBGE). This survey provides detailed information on household expenditure, income, and consumption patterns. The POF dataset includes data on various aspects of household spending, such as food, housing, transportation, education, and healthcare. It also collects information on income sources and household assets, enabling a comprehensive analysis of living standards and economic well-being.

3 Empirical Strategy

The primary goal of this research is to identify the employment response caused by an increase pension earnings within a local labor market. A simple correlation between total pension payments and the level and composition of employment might suffer from significant endogeneity. For instance, a sluggish labor market with fewer employment opportunities in dynamic sectors could lead to increased pension take-up (reverse causality). Conversely, more densely populated local labor markets might simultaneously offer more public services and a larger diversification in terms of economic activities, leading to omitted variable bias. To address endogeneity concerns, I use quasi-experimental variation from the establishment of pension system facilities in specific municipalities.

Empirical setting: I observe pension data from 2007 to 2018 [1]. The sample consists of the 623 units targeted by the policy that were not treated before 2012. We define each cohort G by the year the agency arrived in that unit. If a target municipality was never treated or treated after our period of analysis, we set $G = \infty$. Figure 1 summarizes the distribution of treatment dates in our sample.

The descriptive statistics presented in Table 1 provide a comparison of baseline characteristics between municipalities that were target by the policy and treated at different times or never treated. Specifically, we report eventually treated and never treated groups' mean of total population, the share of the population receiving pension benefits, the share of the population employed as private employees, average wages, the total number of establishments, and the number of establishments with more than 50 employees. In order to compare municipalities treated at different years, the last column of Table 1 reports the p-values of F-tests on the joint significance of the cohort dummies in a linear regression for each dependent variable. The non-significant p-values indicate that there are no systematic differences in baseline characteristics between the municipalities treated at different times at a 5% confidence level. This similarity at the baseline supports the assumption that any future observed differences in outcome trends can be attributed to the timing of the treatment rather than pre-existing differences between the municipalities.

Accordingly, the two maps displayed in Figure 2 and Figure 3 provide information on the geographical distribution of the INSS agencies across the country. Figure 1 highlights the municipalities targeted by the policy. Those are the municipalities that define my sample. The second map, Figure 3, displays, among the targeted municipalities, the treatment timing of each city. As we can see, the policy was implemented across the country. Of the 26 states in Brazil, 23 had municipalities included in the intended list of beneficiaries. Additionally, we observe that there is no systematic geographical correlation between the targeted municipalities and whether they ultimately received an agency.

Staggered differences-in-differences approach: I use the different timing of the opening (or non-opening) of an agency in municipalities targeted by the PEX to identify the causal effect of expanding pension system coverage on labor market outcomes. The implied exclusion restriction imposes that the only way through each an agency affects the labor market is through the reduction in pensions register costs and therefore increase

[1] Importantly, in the end of 2017 the implementation of *INSS Digital* affected how to access the benefit and also in 2019 a major pension system reform changed eligibility criteria and incentives to work

in benefits. One potential concern is that the establishment of an agency might itself mobilize labor resources, directly influencing the labor market. In the next section, I show that this is not likely to be the case.

Regarding the exogeneity, I exploit the fact that selected municipalities were targeted due to similarities in multiple characteristics, but their ability to provide the necessary space for a new agency varied. Additionally, the central government faced different budgetary constraints each year. This combination of factors introduced variation in the treatment assignment. My identification strategy leverages this quasi-randomness to accurately assess the impact of expanding pension coverage on labor market outcomes. More precisely, we assume that when conditioned on municipality-fixed characteristics, that accounts for bureaucratic capacity and space availability, and time-fixed effects accounting for differential central government fiscal constraints, the arrival of an agency in a certain locality in a given year is exogenous.

The goal of our empirical method is to estimate for each cohort, indexed by the treatment day g, what is the effect of being treated for e periods. Following the notation of Callaway (2023), that is:

$$ATT(g, g+e) = \mathbb{E}[Y_{g+e}(g) - Y_{g+e}(0) | G_g = 1] \tag{1}$$

$Y_{g+e}(g)$ is the outcome at date $g+e$ for units that were treated in time g (implying $G_g = 1$) and $Y_{g+e}(0)$ the counterfactual level of the variable of interest at the same date in case those units remained untreated. This quantity is identified under alternative versions of a parallel trend assumption (more on this bellow). The aggregate effect at the length of exposure e across different cohorts is given b:,

$$ATT_{ES}(e) = \sum_{g \in [2012,2018]} \delta^{ES}(g, e) ATT(g, g+e) \tag{2}$$

where δ^{ES} set the underlying weight scheeme. The question is how to aggregate those quantities across cohorts in a meaningful way. Conventionally, the two-way fixed-effects estimator provides an alternative widely used by the applied literature. [2] However, even absent endogeneity concerns, estimators such as two-way fixed effects, may still suffer from sign reversal problems due to negative weights when aggregating individual effects (Borusyak and Hull, 2024). In the specific case of two-way fixed effects (TWFE), this issue is driven by heterogeneity in the treatment effect, which introduces bias in comparisons made between newly treated and already treated units (De Chaisemartin and d'Haultfoeuille, 2020; Goodman-Bacon, 2021; Borusyak et al., 2023). In particular, for the type of event-study design that we presented, Sun and Abraham (2021) decomposes the underlying weights and show that in the case of heterogeneous treatment effects,

[2] In this setting, it implies: $y_{it} = \gamma_i + \lambda_t + \sum_{e \neq -1} \beta_e \times D^e_{it} + \varepsilon_{it}$. Where y_{it} is the outcome of interest for municipality i at year t. γ_i is a municipality fixed effect, λ_t a year fixed effect, and D^e_{it} is a dummy variable indicating if i has received an agency for exactly e periods at date t, where $e = t - g$, and g tracks for each unit the moment of treatment, i.e the cohort indicator labeling. Our main outcomes of interest are the logarithm of total pension benefits and payments in a given municipality, and the logarithm of different employment measures.

the coefficients for a given $g + e$ is contaminated by the effect from different periods[3]. Moreover, they show the problem can be particularly severe when the sizes of cohorts are more heterogeneous, as in our case.

In settings where a type of parallel trend assumption holds, a growing body of literature on staggered difference-in-differences models addresses the treatment effect heterogeneity problem by suggesting alternative weighting schemes (De Chaisemartin and d'Haultfoeuille, 2020; Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Borusyak et al., 2023) [4].

My results are robust to different estimation methods. My preferred estimator is from Callaway and Sant'Anna (2021). Their method addresses the issue of forbidden comparisons while maintaining a less stringent parallel trend assumption (Marcus and Sant'Anna, 2021; De Chaisemartin and d'Haultfoeuille, 2023). It also easily accommodates covariates. Additionally, Callaway (2023) demonstrate how these measures compare with different strategies suggested in the literature. Specifically, the Callaway and Sant'Anna (2021) estimator can be expressed as an imputation estimator, similar to the approaches by Borusyak et al. (2021) and Gardner (2022). Moreover, the Callaway and Sant'Anna (2021) never-treated estimator can be mapped into regression approaches, such as those developed by Sun and Abraham (2021) and Wooldridge (2021).

Under the assumptions that (i) treatment is absorbing: once a unit becomes treated, it remains treated for the rest of the period; (ii) there is no anticipation, meaning that for periods before the arrival of the agency, the observed outcome would be the same regardless of whether the unit will be treated in the future; and (iii) a parallel trends assumption (PTA) for the staggered case depending on how one defines the comparison group. If the never-treated group is set as the fixed reference, then the PTA can be stated as: for all $g, t = 2008, \ldots, 2018$, such that $t \geq g$

$$\mathbb{E}[Y_t(0) - Y_{t-1}(0) \mid G_g = 1] = \mathbb{E}[Y_t(0) - Y_{t-1}(0) \mid G = \infty]$$

In words, in the absence of treatment, the average outcome in treated cohorts and never treated cohorts, would have evolved similarly after the treatment date.

And the ATT can be computed by:

[3]One important remark made by Callaway (2023) is that for β_e to estimate the correct event-study $ATT(e)$, it does not require the condition that $ATT(g,t)$ is constant across t, unlike the TWFE static regression. Instead, the homogeneity restriction in this case requires $ATT(g, g+e)$ to be constant across groups.

[4]If we shift from a model-based approach of potential outcomes to a design-based approach for treatment assignment and further assume that, for the target municipalities, the adoption date of the policy is effectively random, Athey and Imbens (2021) shows that standard DID estimators are an unbiased weighted average of a particular causal effect. Moreover, under this quasi-random assignment, Roth and Sant'Anna (2023) suggest a more efficient method for estimating the causal effect of the intervention with a more *intuitive* interpretation of the weight scheme. In the setting I am studying here it seems potentially to strong to go all the way to assume the randomness in treatment timing required by those papers. Still, in the robustness checks, we show that our results hold using Roth and Sant'Anna (2023) estimators.

$$\widehat{ATT}^{CS}_{\text{NT}}(g, g+e) = \frac{1}{n}\sum_{i=1}^{n}\frac{\mathbf{1}\{G_i = g\}}{\hat{P}(G_i = g)}(Y_{i,g+e} - Y_{i,g-1}) - \frac{1}{n}\sum_{i=1}^{n}\frac{\mathbf{1}\{G_i = \infty\}}{\hat{P}(G_i = \infty)}(Y_{i,g+e} - Y_{i,g-1}) \tag{3}$$

The estimator computes the difference between the average path of the outcome variable for units at group G from $g-1$ to e and the same measure for the same period for units that are never treated.

Analogously, we can use as a control group units that are not-yet-treated by that date. This would require the following PTA: for all $g, d, t = 2008, \ldots, 2018$, such that $t \geq g$ and $d > t$

$$\mathbb{E}[Y_t(0) - Y_{t-1}(0) \mid G_g = 1] = \mathbb{E}[Y_t(0) - Y_{t-1}(0) \mid D_d = 0]$$

Where D_d switches to 1, when that unit becomes treated in a date $d > t$. This implies the following estimator:

$$\widehat{ATT}^{CS}_{\text{NYT}}(g, g+e) = \frac{1}{n}\sum_{i=1}^{n}\frac{\mathbf{1}\{G_i = g\}}{\hat{P}(G_i = g)}(Y_{i,g+e} - Y_{i,g-1}) - \frac{1}{n}\sum_{i=1}^{n}\frac{\mathbf{1}\{D_{i,g+e} = 0\}}{\hat{P}(D_{g+e} = 0)}(Y_{i,g+e} - Y_{i,g-1}) \tag{4}$$

Regarding the identification, we know that the treatment is absorbing, as no agency closures were observed during the analysis period. The no-anticipation assumption might be at risk if news about the opening of an agency prompted people to start organizing their documents and traveling to different cities to claim benefits before the agency in their locality began operations. However, even if individuals were granted benefits in a different municipality, the data records the payment at the level of the beneficiary's place of residence. By setting the treatment date to periods before the actual opening, I demonstrate that there is no evidence of anticipation effects. The PTA follows from our premise that, when conditioning on municipality fixed effects and time fixed effects, the arrival of an agency is exogenous to the trends in the outcomes of interest.

Although it is impossible to test the validity of the parallel trends assumption (PTA) since we do not observe treated units in the absence of treatment after treatment has started, common practice relies on examining pre-trends. In the next section, I show that for all outcomes of interest, there were no differential trends before the treatment date across groups. For the PTA violation to occur, it would require that an unobservable, time-varying municipal characteristic determines the specific year a city received treatment and that this characteristic systematically correlates with labor market dynamics only after the treatment date. It is plausible that political preferences, mayoral effort, and state capacity influence a municipality's ability to find the land, enter the queue, and manage the construction process. However, for these factors to invalidate our results, the effort and capacities associated with the delivery of the agency must not be captured by our fixed effects (i.e, they vary over time) and must determine labor market-related policies only after the agency's arrival, not playing any role in the outcome variables before or even during the process of setting up the building. Finally, standard errors are clustered by municipality, the level of treatment assignment.

4 Main Results

This section discusses the results of the main empirical strategy. The task here is to show how an increase in pension payments changes labor market outcomes in treated municipalities. To test this broader question, several hypotheses must hold true.

To clearly set the timing of the events, all the data is from December of the respective year. A unit is treated in a given year if it receives an agency before or in October of that year[5] Our exogenous variation in pensioners' income comes from the increase in the total number of pensioners in a given municipality after an INSS agency office opens. Therefore, the arrival of the facility must increase the number of benefits and total payments. This is precisely what the plot (a) of Figure 4 illustrates. The event study plot shows the increase in the logarithm of total payments after the opening of an agency using the Callaway and Sant'Anna (2021) estimator.

First, it is important to highlight the absence of pre-trends. Municipalities treated in different time periods were trending similarly for more than five years before the office arrival. In the post-treatment period, we observe a sharp and significant increase in pension payments channeled to those municipalities. After five years, the increase reaches up to 6%. The aggregate rise in payments for the period after the agency arrival is 4.3%, as summarized in Table 2, column (1), which displays my preferred specification using the Callaway and Sant'Anna (2021) estimator and excluding the 5% outliers in terms of 2010 total population. [6]. Columns (2) and (3) of Table 2 show that the result is consistent when using Never treated units as a control group and also including as control the baseline characteristics discussed in Table 1 (i.e mean of total population, the share of the population receiving pension benefits, the share of the population employed as private employees, average wages, the total number of establishments, and the number of establishments with more than 50 employees). For the median municipality in the never treated group, the total average annual payments from 2012 to 2018 were R$ 45 million, in 2016 values. This implies that if this municipality were treated, an additional R$ 1.9 million would have been transferred per year. In terms of number benefits, I find a 4% increase, which translates in extra 266 benefits for the same municipality. The increase in number of benefits is reassuring that the observed hike in total payments is due to extensive margin response, that is, increase in take-up of eligible individuals.

Given the documented increase in benefits, I am interested in understanding how it translates into labor market outcomes. Plot (b) of the same figure reproduces the event study for the logarithm of total private formal jobs. The second row of column (1) in Table 2 shows that the average effect of the treatment on employment over the treatment period

[5]In the sample, all the agencies were opened between January and October.

[6]The results are highly consistent when using the whole sample. The restricted sample excludes municipalities smaller than 10,000 and larger than 100,000, which lie beyond the population limits stated in the expansion policy goal.

is 4.2%. Again, the implied response of the outcome variable to the agency arrival is consistent across different specifications of the Callaway and Sant'Anna (2021) estimator. The percentage increase translates to an additional 64.5 jobs for the median municipality in the never treated group. These results are aligned with the growing literature that finds a positive employment multiplier effect from government transfers (Corbi et al., 2019) in general and cash transfers specifically (Gerard et al., 2021; Egger et al., 2022).

Following Gerard et al. (2021) methodology, our estimates suggest a cost-per-job of 2.74 minimum wages, very similar to the 2.1 they find in the *Bolsa Família* context. In terms of monetary amount, using 2016 dollars, my estimate suggests a cost of US$ 7,200, which lies between the US$ 5,600 of Gerard et al. (2021) and the US$ 8,000 of Corbi et al. (2019). I compute the standard deviations using the same bootstrap procedure suggested by Gerard et al. (2021).

Understanding the context is crucial to grasp the meaning of these coefficients. In 2011, before the expansion program, target municipalities had on average 2 beneficiaries for every formal employment worker. This high ratio is due to Brazil's non-contributory schemes, which allow rural and informal workers, as well as those who spend much of their working lives unemployed, to retire with a pension. Nationally, the ratio was 1.7 beneficiaries per formal worker at the same year. This highlights the pension system's role as the main source of stable income in Brazil's smaller municipalities, where informality, seasonal rural work, home production, and unemployment are particularly widespread.

Importantly, we observe no significant changes in the average wage, with a non-significant coefficient of -0.002 in the main specification, as shown by the third row of Table 2 and plot (c) of Figure 4. I interpret these results as suggestive evidence of an infinitely elastic labor supply to the formal sector. This could be due to a large contingent of involuntary unemployed workers or informal workers willing to formalize if offered the opportunity. Although, given the lack of data on unemployment and informality at the municipal level, I cannot fully test each of these channels, the main message that these results convey remains the same: the main factor preventing a higher level of employment in the formal sector was labor demand rather than labor supply.

Now, I turn to analyze the effects of the increased purchasing power of pension recipients on labor allocation across sectors and firms. First, as shown by plot (d) in Figure 4, employment growth was concentrated in service sector employment. While manufacturing and agriculture show no significant increase in employment, with coefficients of 0.01 and -0.006, respectively, the fourth row of Table 2 suggests the service sector drives the overall employment growth with a 5.09% increase. These figures are consistent with the non-tradable sector being more responsive to local income surges and with non-homothetic preferences disproportionately increasing service sector demand following an income hike.

The average size of firms and the number of establishments were not significantly affected by the treatment, which raises the question of where the employment increase is being directed. However, the null effect on the total number and average size of firms is consistent with the increase in employment favoring the skewness of the distribution of firm size . As the plots (e) and (f) of Figure 4 show, the income surge triggered by the arrival of the pension agency reallocates employment to larger firms.

I use two different measures to demonstrate this effect. First, the total number of firms with more than 50 employees increases. This effect results from two mechanisms: the growth of firms that were already above the threshold and smaller firms expanding to surpass the 50-employee cut-off. The effect is smaller for firms above 20 employees and quite similar for firms above 70 employees, suggesting that the increase in employment is primarily driven by the growth of already medium-sized companies. On average, during the treatment period, the never-treated units had 7.4 firms with more than 50 employees. My estimates suggest that this number would have increased by 0.4 firms. In terms of employment share, the largest 10% of firms in the never-treated group accounted for 80% of total employment on average. Treatment exposure would have increased this concentration to 81.14%.

4.1 Robustness Checks

One core assumption of my identification strategy is that the arrival of a pension office in a municipality does not directly impact employment. My main results focus on private employment, demonstrating job creation outside the public sector. It could still be possible that the surge in demand explaining the multiplier effects in private employment is driven by an increase in public employment rather than pensioners' income. However, this is unlikely since INSS records indicate that the average agency created during the period employs only 2.7 people, who could be reallocated within the municipal public service or provided by the central government. Consistent with this, Table 4 in Appendix B shows no significant effect of agency arrival on public employment.

Another potential issue with my empirical strategy could be that information about the agency's arrival might push individuals to register for pensions at existing agencies in different municipalities. If that were the case, the relevant treatment would be the announcement of the agency rather than its actual opening. Regardless of where registration occurs, payments are assigned to the municipality of residence of each recipient in the administrative records of pensions. This allows me to verify if pension payments increase before the agency's opening. Figure 5 in Appendix B presents the event study for the logarithm of total payments when treatment trimming is set to one year before the actual event. The figure shows that this is not the case. The estimated impact of agency arrival on payments at the new event time is both small and non-significant. Payments start to grow after the agency's arrival during time 1 in the plot. This further confirms that the relevant treatment measures the agency's opening and not its announcement.

Finally, the choice of differences-in-differences estimator should be informed by the suitability of the identifying assumptions for each context. As discussed in Section 3, I believe the assumptions of the Callaway and Sant'Anna (2021) estimator are appropriate for my setting. However, the robustness of my estimates across various methods, as shown in Table 5, reinforces the validity of the main findings. The consistency of results, regardless of the estimator used, indicates that imposing different, albeit related, identifying assumptions produce similar results.

5 Household level evidence

This section investigates the causal effects of pension income on household total income and consumption behavior. For pensions to generate an aggregate demand response substantial enough to impact employment levels, it must be shown that they increase household income,leading to a rise in consumption.

I begin by documenting that, at the national level, average pension earnings are higher and more stable than informal and more stable than formal earnings. Table 6 illustrates this using PNAD survey data from 2007 to 2015, showing that average pension income is nearly 40% higher than informal earnings.[7]

The literature extensively documents that informal sector income is both lower and more volatile compared to formal sector earnings in Brazil (e.g., Gomes et al., 2020; Engbom and Moser, 2022). Using the sample of treated municipalities during the period, we find that the number of pension benefits are more stable than formal private jobs and that average pension earnings are more stable than average formal earnings. Table 7 presents, for the target municipalities, the average coefficient of variation for the total number of benefits and private formal contracts, as well as the average coefficient of variation for average earnings across years.

The fact that pensions provide agents with a source of income that is higher than informal sector earnings and more stable than both informal and formal sector wages rationalizes the significant demand effect observed at the aggregate level. To provide causal evidence of this mechanism, I leverage discontinuities in pension eligibility to recover the causal impact of the benefits on income and consumption from household budget surveys.

Rural workers become eligible for the rural non-contributory pension benefit at 60 years of age. Workers who have transitioned between formal and informal jobs and have accumulated only 15 years of contributions can retire if they are old enough. That means 60 years old if female, and at 65 if male. Under this partial contributory system, the benefit is calculated based on their average contributions. Finally, informal workers who contribute for less than 15 years can retire at 65 years old to receive a minimum wage benefit (BPC).[8] Combined, these three schemes represent 76% of the total retirement benefits in the target municipalities at baseline (2011).

Empirical strategy household level analysis Assignment to the treatment is not independent of consumption and labor market decisions. Factors that are unobservable, such as effort, experience, job quality, and family needs, influence the decision to retire. Merely comparing individuals based on similar observable characteristics who differ in their treatment status could introduce selection bias into the estimates. However, the variation introduced by age-based eligibility rules in the likelihood of receiving a pension,

[7]PNAD was discontinued after 2015; thus, the values are computed using information from 2007 to 2015.

[8]In theory, the *Benefício de Prestação Continuada* (BPC) is a means-tested benefit available for older individuals in households with per capita income less than one-quarter of the minimum wage. However, the means test is much less stringent than the rules for receiving a cash transfer. The definition of family income for accessing the program excludes i) other social assistance programs such as the PBF; ii) other BPC transfers or pensions; iii) transitory benefits; and iv) health-related expenses.

allows me to use age as an instrumental variable (IV) to discern the causal effect of pension transfers on household income and consumption decisions.[9]

The POF 2008-2009 household budget survey combines information on the level and source of income and consumption for each type of good and store. Since my data is at the household level, I use the maximum age within the family as the eligibility criterion.

The exogeneity condition—or the condition that there is no manipulation of the running variable—is naturally met because of the immutable nature of age. Hence, our primary identifying assumption is that the maximum age of a given household only influences consumption decisions around the cutoff through pension eligibility.[10] To provide further evidence in favor of the identification strategy, Table 9 shows that for a variety of variables, households above and below each age cutoff are similar.

In the two-stage least squares (2SLS) method employed here, the first stage assesses the effect of age's exogenous variation on treatment status, denoted as $Pension = 1$ if pension income is greater than zero for an individual i. This distinguishes the treated from the control units. The fitted value for treatment status from the first stage is then used in the second stage to determine the program's causal effect by comparing those distinguished only by the exogenous component of treatment assignment.

Given the survey nature of the data, the sample size is a constraint to obtaining precise estimates. Therefore, I define the following $AgeGap$ variable to be able to run the regression, combining the two age eligibility cutoffs to increase power.

$AgeGap =$ Household maximum age $- 60$ if Household maximum age $\in [58, 62]$

$AgeGap =$ Household maximum age $- 65$ if Household maximum age $\in [63, 67]$

$AgeGap = \infty$ otherwise.

For households for which $AgeGap$ is $\in \{-2, -1, 1, 2\}$, we run the following 2SLS equation:[11]

The first stage equation is:

$$Pension_i = \gamma_0 + \gamma\mathbb{1}(gap > 0)i + G_i + \mathbb{X}i'\Gamma + \eta_i$$

where G_i is a group fixed effect indicating whether the unit is around the 60 years-old age-eligibility cutoff or the 65 years-old cutoff, ensuring that each household above a given threshold is compared with the correct control group below that threshold. $\mathbb{X}_i$ is a vector

[9]This empirical context aligns with a fuzzy regression discontinuity design, which can be linked to an instrumental variable estimation around the cutoff (Abadie and Cattaneo, 2018). Given the discrete nature of our running variable, the IV terminology seems more precise.

[10]Formally, the IV strategy requires a monotonicity assumption. Here, it implies that an individual above the age cutoff cannot be less likely to receive the treatment—a condition that is clearly met.

[11]I exclude $AgeGap = 0$ because depending on where individuals live, it can take longer to access the benefit, and it is not clear that just being eligible will make one a recipient of the transfer during the same year.

of covariates of predetermined characteristics.[12] Accordingly, this implies the following second stage equation:

$$Y_i = \delta_0 + \delta \widehat{Pension}_i + \mathbb{X}i'\Delta + G_i + \epsilon i$$

where Y_i is the outcome of interest for household i and δ is the estimated causal effect of the pensions.

Results household level analysis Table 9 shows that the likelihood of receiving a pension effectively hikes when the maximum age within a household crosses an age eligibility cutoff. This results in an increase in pension income, which translates to an increase, albeit of smaller magnitude, in total household income. Consistent with the literature on pensions, particularly the findings on rural pensions in Brazil by de Carvalho Filho (2008), the income effects associated with a higher guaranteed income lead to a reduction in other sources of earnings within the household.

The rise in disposable income leads to an increase in consumption, primarily driven by a surge in consumption at formal stores, as defined by Bachas et al. (2023). The estimates suggest a marginal propensity to consume of 0.95, consistent with the observation that the non-contributory and partially contributory systems primarily benefit liquidity-constrained individuals with no access to credit markets.

To put these results into perspective, the mean income for the control group is 14,132 BRL, and the mean expenditures for the control group are 11,915 BRL. This implies an income increase of 36% for the average household in the control group and a corresponding increase in consumption of nearly the same magnitude (40%).

At the aggregate level, for the median municipality, the agency implies an additional 1.9 million BRL in income, or estimates that 5,091 out of 9,297.47 BRL becomes disposable income, which is 54% of the total, amounting to 1.04 million BRL. Of this, 95% translates into consumption, nearly 1 million BRL. This represents a consumption shock of 3.1% of the total formal sector wage bill of the median municipality and 5.4% of the private formal sector wage mass.

6 Discussion

This contributes to understanding local multipliers and the interaction between pension systems and development. To the best of my knowledge, it is the first attempt to document the local multiplier effect of the world's most common type of transfer. In developing countries, pensions raise average incomes and stabilize earnings. After expanding pension coverage, I show that private formal employment rises, especially in the service sector and mainly within larger firms. The absence of a significant effect on the

[12] It includes the predicted value of household income excluding pension benefits following Gadenne et al. (forthcoming). This is important because people retiring after 65 are usually from poorer households to begin with, so we must compare households that, if not for age, would have similar income streams.

average wage suggests a elastic labor supply and slackness in the use of labor. This aligns with a demand-constrained development path where workers are willing to work at current wage levels, but firms lack the capacity to absorb them into the formal sector due to low aggregate demand. The results indicate that only larger firms can leverage increased demand to create more jobs, consistent with a positive returns-to-scale economic model.

The household-level analysis confirms that the increase in pension income is associated with a rise in family expenditures, especially at formal stores. Unlike fiscal multipliers of inter-government transfers, my results are based on family-level consumption decisions. The main takeaway is that transferring resources to liquidity-constrained households triggers an aggregate demand response that benefits the whole economy. This suggests that a lack of demand is an important constraint for job creation and firm growth. It also highlights that understanding the consumption behavior of the beneficiaries of each type of fiscal expenditure is crucial for assessing its aggregate-level impact.

Raising formalization rates is vital for managing demographic shifts, as only the formal sector typically funds the pension system. Ensuring stable income for retirees prevents old-age poverty and boosts the formal sector by alleviating demand constraints in areas where income would otherwise be lower and unstable. Sustainable inter-generational transfers require concurrent productive growth. The findings suggest that a lack of demand can also limit the growth of more productive larger firms.

This paper argues that viewing pensions as a cash transfer that boosts local income and earning stability can provide policymakers with valuable insights into optimally designing pension schemes and preparing for demographic transitions. As developing countries age, this issue will likely become more relevant. In this paper, I do not address the financing of the pension system. Future research should combine the general equilibrium demand effects discussed here with the potential aggregate distortions of raising revenues, particularly if this implies higher tax rates on formal jobs. Specifically for Brazil, the results suggest that efforts to make pension systems more fiscally sustainable must balance the aggregate importance of guaranteeing income to poorer elderly households.

Finally, demographic changes may necessitate the production of different goods due to varying consumption patterns across age groups. An interesting avenue for research is to explore how to combine pension design with policies that enable productivity gains in sectors more heavily consumed by retirees.

Appendix

A: Main Figures and Tables

Figure 1: Agency arrival date

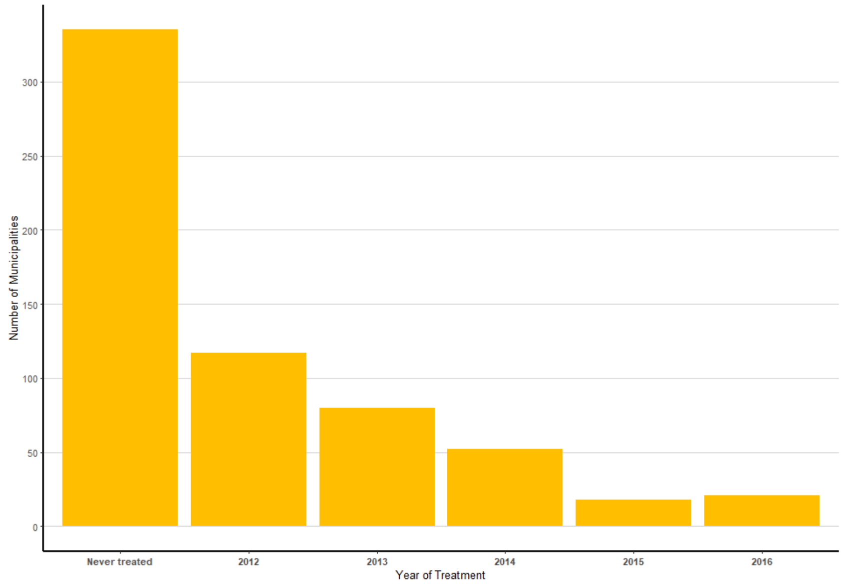

Note: The plot shows the size of each cohort, defined by the year when each municipality received an agency.

Figure 2: Geographical distribution of targeted municipalities

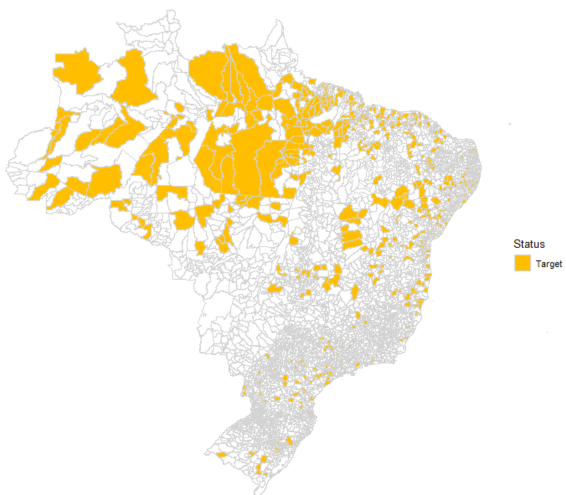

Note: The map displays the municipalities that were target by the policy.

Figure 3: Treatment roll-out

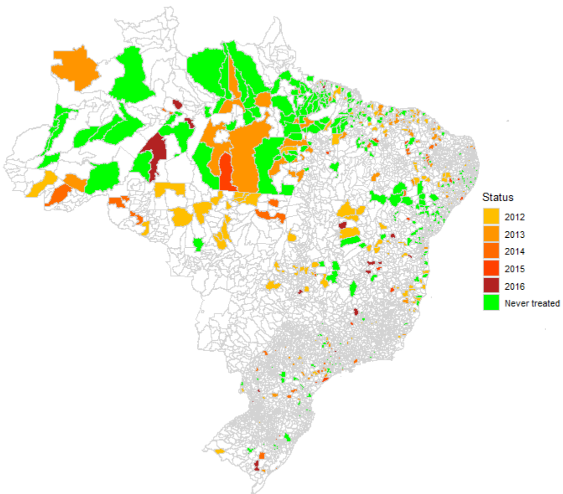

Note: The map displays the municipalities that were target by the policy according to their year of treatment.

Figure 4: Impact of Agency Arrival - Event Study

(a) Total Pension Payments

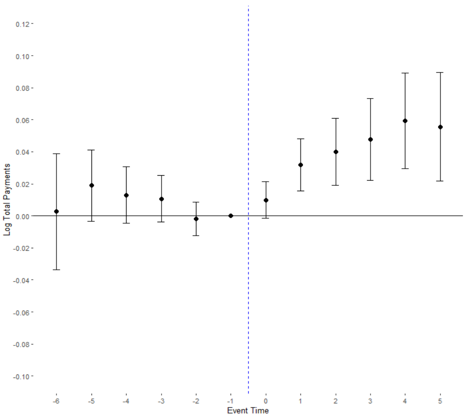

(b) Total Private-sector Employment

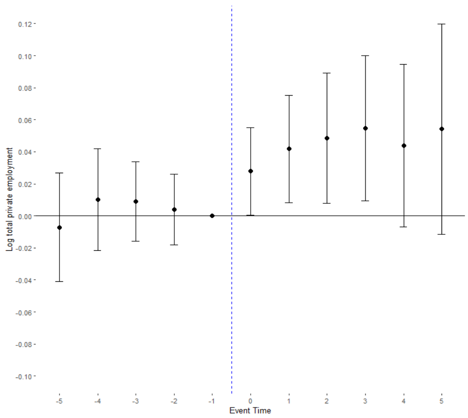

(c) Average Wages

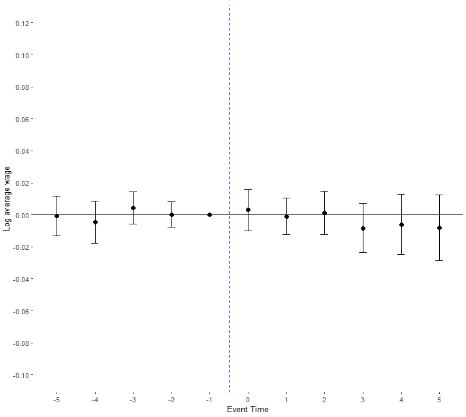

(d) Service-sector Employment

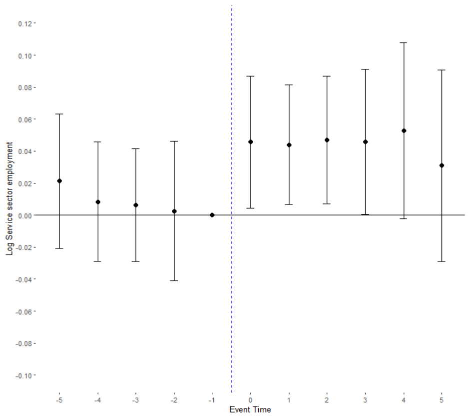

(e) Number of Firms with more than 50 Employees

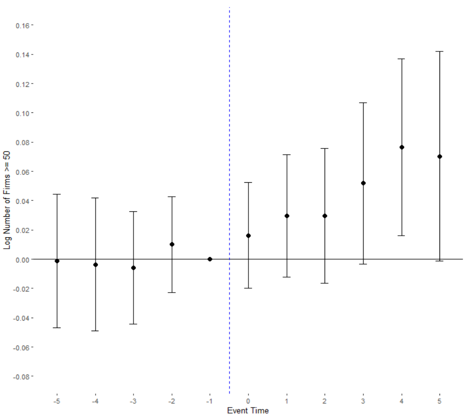

(f) Employment Share at Top 10% Biggest Firms

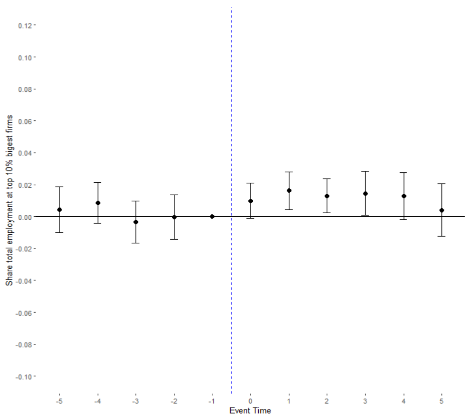

Note: The plots depict the impact of the opening of an INSS agency on various economic indicators in a municipality. The first plot shows the impact on the log of total amount of pension payments (a). The second plot illustrates the effect on the log of total private-sector formal employment (b). The third plot shows the impact on the log of average wages (c). The fourth plot depicts the effect on the log of total service-sector employment (d). The fifth plot shows the impact on the log of number of firms with more than 50 employees (e). The sixth plot illustrates the effect on the employment share at the top 10% biggest firms (f), with the y-axis representing the the share. In all plots, the x-axis indicates the event time in years, with the agency's opening normalized at time zero. The points represent the estimated effects, and the vertical lines indicate the 95% confidence intervals for these estimates. Standard errors are clustered at the municipal level.

Table 1: Descriptive Statistics

Variable	**All Sample**	**Eventually Treated**	**Never Treated**	**p-value (F test)**
Population	31864.19	31507.53	32396.62	0.41
Share Beneficiaries	0.14	0.15	0.13	0.08
Share Private Formal Employment	0.08	0.07	0.09	0.13
Average Wage	958.42	935.41	992.76	0.34
Number of Establishments	291.87	287.60	298.24	0.97
Number of Establishments (50+)	7.17	6.22	8.59	0.45
Number of Observations	623	288	335	

Note: This table shows descriptive statistics for baseline characteristics across different samples. The first column presents the average value of each variable for the whole sample, the second column shows the average for those treated between 2012 and 2016, and the third column provides the average for those never treated. The final column reports the p-value of a joint F-test of the coefficients from a linear regression of each variable against a set of cohort dummies, testing whether the timing of treatment explains any differences in baseline information across cohorts. Population data is sourced from the 2010 census, while other variables are baseline values calculated for 2011.

Table 2: Summary Main Results

	(1)	(2)	(6)
	CS Not Yet Treated	CS Never Treated	Baseline Controls
Total Pension Payments	0.0429***	0.0426***	0.0432***
	(0.0111)	(0.0114)	(0.0118)
Total Private Employment	0.0423**	0.0424**	0.0395**
	(0.0193)	(0.0193)	(0.0201)
Average wage	-0.0024	-0.003	-0.0056
	(0.0061)	(0.0065)	(0.0072)
Service Sector Employment	0.0509***	0.0526***	0.0490**
	(0.0215)	(0.0213)	(0.0249)
Share of Employment at the 10% biggest firms	0.0114**	0.0127**	0.0141**
	(0.0053)	(0.0057)	(0.0062)
Number of Firms with 50+ employees	0.0493**	0.0455**	0.0455**
	(0.0214)	(0.0218)	(0.022)
Observations	6501	6501	6501
Periods interact with baseline controls			✓

Note: The table depicts the Average Treatment Effects on the treated of the agency arrival using different specifications of the Callaway and Sant'Anna (2021) estimator. Each column shows the results for the logarithm of the variables in the first column, but for the share of employment at the top decile that is kept as a share. The sample excludes the 5% outliers in terms of total population as measured by the 2010 Census. Column 1 represents Callaway and Sant'Anna (2021) not yet treated, Column 2 represents Callaway and Sant'Anna (2021) using only never treated units as comparison group, Column 3 runs the same model as (1) controlling for baseline mean of total population, the share of the population receiving pension benefits, the share of the population employed as private employees, average wages, the total number of establishments, and the number of establishments with more than 50 employees. Standard errors are clustered at the municipal level.Significance levels are denoted as: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 3: Cost per Job in Terms of Minimum Wages (MW) and Dollars in 2016

	Point Estimate	Lower Bound (95% CI)	Upper Bound (95% CI)
MW (2016)	2.74	2.02	3.52
Dollars (2016)	7200	5300	9300

Note: The table provides estimates of the cost per job created by a specific intervention, measured in 2016 minimum wages (MW) and US dollars. Confidence intervals follow Gerard et al. (2021) and are based in 1,000 bootstrapped replications of the point estimates.

B: Robustness Checks

Figure 5: Total Pension Payments - No antecipation

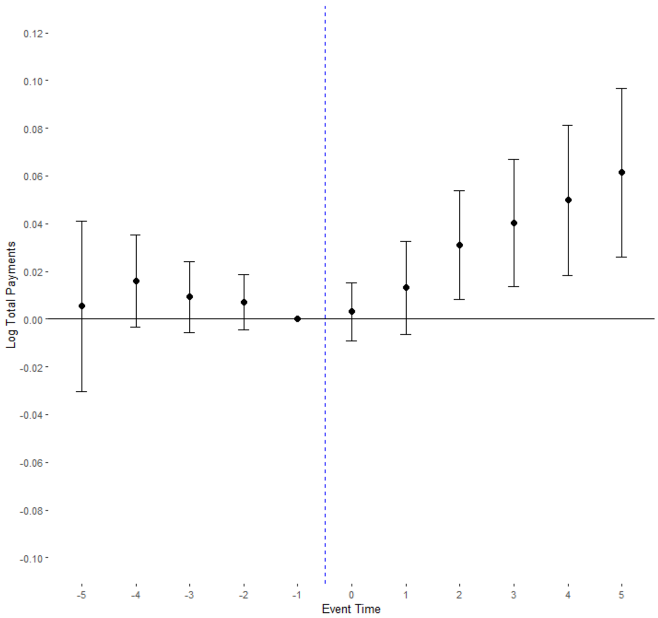

Note: The plots depict the impact of the opening of an INSS agency on the log of total amount of pension payments. The x-axis indicates the event time in years, with the agency's opening set to one period before it happened. The points represent the estimated effects, and the vertical lines indicate the 95% confidence intervals for these estimates. Standard errors are clustered at the municipal level.

Table 4: Public Employment

	(2)	(3)	(4)
Public Employment	0.0283	0.0294	0.0218
	(0.0499)	(0.0523)	(0.0535)

Note: The table depicts the Average Treatment Effects on the treated of the agency arrival using different specifications of the Callaway and Sant'Anna (2021) estimator. Each column shows the results using an inverse hyperbolic sine transformation due to some municipalities reporting 0 public employment in some years. The sample excludes the 5% outliers in terms of total population as measured by the 2010 Census. Column 1 represents Callaway and Sant'Anna (2021) not yet treated, Column 2 represents Callaway and Sant'Anna (2021) using only never treated units as comparison group, Column 3 runs the same model as (1) controlling for baseline mean of total population, the share of the population receiving pension benefits, the share of the population employed as private employees, average wages, the total number of establishments, and the number of establishments with more than 50 employees. Standard errors are clustered at the municipal level. Significance levels are denoted as: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 5: Robustness Checks

	(1)	(2)	(3)	(4)
	SA	RS	TWFE	CS - All sample
Total Payments	0.0425***	0.0598***	0.0352***	0.0426***
	(0.0114)	(0.0091)	(0.0097)	(0.0106)
Total Private Employment	0.0451**	0.0502***	0.0347**	0.041**
	(0.0197)	(0.0182)	(0.0168)	(0.0184)
Average wage	-0.003	-0.0079	-0.0001	-0.0018
	(0.006)	(0.0053)	(0.0060)	(0.0059)
Service Sector Employment	0.0552***	0.0558***	0.0437**	0.0498***
	(0.0219)	(0.0201)	(0.0186)	(0.0204)
Employment at the 10% biggest firms	0.0122**	0.0012	0.0053	0.0107**
	(0.0053)	(0.0045)	(0.0051)	(0.0051)
# Firms with 50+ employees	0.0493**	0.0489**	0.0358*	0.0475**
	(0.0213)	(0.0209)	(0.0192)	(0.0201)
Observations	6501	6501	6501	6853

Note: The table depicts the Average Treatment Effects on the treated of the agency arrival using different methods and samples. Each column shows the results for the logarithm of the variables in the first column, but for the share of employment at the top decile that is kept as a share. The sample used in the three columns excludes the 5% outliers in terms of total population as measured by the 2010 Census. Column 1 shows the estimates using Sun and Abraham (2021), Column 2 presents the results when using Roth and Sant'Anna (2023), Column 3 is for TWFE, running a sparse version of the TWFE discussed in footnote 2, where post-treatment periods are combined under one *post* dummy. Column 4 includes all samples without excluding outliers. Standard errors are clustered at the municipal level.Significance levels are denoted as: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

C: Household level analysis

Table 6: Average earnings across different income sources

Formal Private Sector	Informal Sector	Pensions
1385	827	1226

Note: The table displays the average earnings according to each income source using PNAD survey from 2007 to 2015. Values are in terms of 2015 prices..

Table 7: Average Coefficient of Variation across Target Municipalities

	Formal Private Sector	Pensions
Number of Earners	0.253	0.173
Average Earnings	0.337	0.222

Notes: The table displays, for the target municipalities, the average coefficient of variation (CV) of the total number of benefits and private formal contracts and the average coefficient of variation of the average earnings across years using RAIS and pensions payment sheets. To compute the average earnings we deflate all the values to 2016 prices.

Table 8: Descriptive Statistics for Different Age Groups Using POF Data

	Age group 1		Age group 2	
	58-59	61-62	63-64	66-67
Years of Schooling	4.486	3.723	3.133	3.035
Illiteracy	0.250	0.321	0.344	0.407
Non-white	0.604	0.599	0.600	0.606
Household Size	3.257	3.214	3.184	3.016
Number of Rooms	2.001	1.987	1.939	1.900
Number of Bathrooms	1.144	1.150	1.111	1.144
Connection to Water	0.752	0.769	0.743	0.763
Connection to Sewage	0.352	0.356	0.330	0.318

Notes: The table displays descriptive statistics for different age groups using POF data. Columns show mean values for the indicated age ranges. The first row indicates the number of years of schooling of the oldest member of the household. The second row shows the share of illiterate oldest members, and the third row shows the share of non-white members. The following rows describe household conditions, including household size, number of rooms, number of bathrooms, and whether the household is connected to water and sewage systems.

Table 9: Results - Household level

	P(Pension = 1)	Pension Income	Total Income	Expenditure	Expenditure in Formal Stores
Eligibility = 1	0.15***				
	(0.019)				
$P(\widehat{Pension} = 1)$		9297.47***	5091.21**	4877.44*	4855.22*
		(1298.34)	(2105.42)	(2937.71)	(2679.07)
Observations	5244	5244	5244	5244	5244
F-test	150.8				

Notes: The table reports 2SLS estimation results. The dependent variable for the first column is P(Pension = 1), for the second column is Pension Income, for the third column is Total Income, for the fourth column is Expenditure, and for the fifth column is Expenditure in Formal Stores. Standard errors are reported in parentheses. The number of observations for each regression is 5,244. The F-test statistic is reported for the first stage regression in the first column. Significance levels are denoted as: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

References

A. Abadie and M. D. Cattaneo. Econometric methods for program evaluation. *Annual Review of Economics*, 10:465–503, 2018.

C. Ardington, A. Case, and V. Hosegood. Labor supply responses to large social transfers: Longitudinal evidence from south africa. *American economic journal: Applied economics*, 1(1):22–48, 2009.

S. Athey and G. W. Imbens. Design-based analysis in difference-in-differences settings with staggered adoption. *Journal of Econometrics*, 225(2):254–277, 2021.

A. J. Auerbach, Y. Gorodnichenko, and D. Murphy. Local fiscal multipliers and fiscal spillovers in the united states. Technical report, National Bureau of Economic Research, 2019.

P. Bachas, L. Gadenne, and A. Jensen. Informality, Consumption Taxes, and Redistribution. *The Review of Economic Studies*, page rdad095, 09 2023. ISSN 0034-6527. doi: 10.1093/restud/rdad095. URL https://doi.org/10.1093/restud/rdad095.

O. Becerra. Pension incentives and formal-sector labor supply: Evidence from colombia. *Documento CEDE*, (2017-14), 2017.

F. Bonnet, J. Vanek, and M. Chen. Women and men in the informal economy: A statistical brief. *International Labour Office, Geneva*, 20, 2019.

K. Borusyak and P. Hull. Negative weights are no concern in design-based specifications. *AEA Papers and Proceedings*, 114:597–600, May 2024. doi: 10.1257/pandp.20241046. URL https://www.aeaweb.org/articles?id=10.1257/pandp.20241046.

K. Borusyak, X. Jaravel, and J. Spiess. Revisiting event study designs: Robust and efficient estimation. *arXiv preprint arXiv:2108.12419*, 2021.

K. Borusyak, X. Jaravel, and J. Spiess. Causal inference with two-way fixed effects models: An application to trade policy. *American Economic Journal: Economic Policy*, 15(1): 1–32, 2023.

F. J. Buera and J. P. Kaboski. Can Traditional Theories of Structural Change Fit the Data? *Journal of the European Economic Association*, 7(2-3):469–477, 05 2009. ISSN 1542-4766. doi: 10.1162/JEEA.2009.7.2-3.469. URL https://doi.org/10.1162/JEEA.2009.7.2-3.469.

F. J. Buera and J. P. Kaboski. The rise of the service economy. *American Economic Review*, 102(6):2540–69, 2012.

B. Callaway. Difference-in-differences for policy evaluation. *Handbook of Labor, Human Resources and Population Economics*, pages 1–61, 2023.

B. Callaway and P. H. Sant'Anna. Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230, 2021.

C. Canelas and M. Nino Zarazua. Social protection and the informal economy: What do we know? 2022.

A. Case and A. Deaton. Large cash transfers to the elderly in south africa. *The Economic Journal*, 108(450):1330–1361, 1998.

R. Ceni. Pension schemes and labor supply in the formal and informal sector. *IZA Journal of Labor Policy*, 6(1):1–29, 2017.

G. Chodorow-Reich. Geographic cross-sectional fiscal spending multipliers: What have we learned? *American Economic Journal: Economic Policy*, 11(2):1–34, 2019.

D. Comin, D. Lashkari, and M. Mestieri. Structural change with long-run income and price effects. *Econometrica*, 89(1):311–374, 2021. doi: https://doi.org/10.3982/ECTA16317. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA16317.

R. Corbi, E. Papaioannou, and P. Surico. Regional transfer multipliers. *The Review of Economic Studies*, 86(5):1901–1934, 2019.

G. Cruces and M. Bérgolo. Informality and contributory and non-contributory programmes. recent reforms of the social-protection system in uruguay. *Development Policy Review*, 31(5):531–551, 2013.

I. E. de Carvalho Filho. Old-age benefits and retirement decisions of rural elderly in brazil. *Journal of Development Economics*, 86(1):129–146, 2008.

C. De Chaisemartin and X. d'Haultfoeuille. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–96, 2020.

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

B. R. Delalibera, P. C. Ferreira, and R. M. Parente. Social security reforms, retirement and sectoral decisions. 2023.

E. Duflo. Child health and household resources in south africa: evidence from the old age pension program. *American Economic Review*, 90(2):393–398, 2000.

E. Duflo. Schooling and labor market consequences of school construction in indonesia: Evidence from an unusual policy experiment. *American economic review*, 91(4):795–813, 2001.

E. Duflo. Grandmothers and granddaughters: old-age pensions and intrahousehold allocation in south africa. *The World Bank Economic Review*, 17(1):1–25, 2003.

D. Egger, J. Haushofer, E. Miguel, P. Niehaus, and M. Walker. General equilibrium effects of cash transfers: experimental evidence from kenya. *Econometrica*, 90(6):2603–2643, 2022.

N. Engbom and C. Moser. Earnings inequality and the minimum wage: Evidence from brazil. *American Economic Review*, 112(12):3803–47, 2022.

T. Fan, M. Peters, and F. Zilibotti. Growing like india—the unequal effects of service-led growth. *Econometrica*, 91(4):1457–1494, 2023. doi: https://doi.org/10.3982/ECTA20964. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA20964.

E. Farhi and I. Werning. Fiscal multipliers: Liquidity traps and currency unions. In *Handbook of macroeconomics*, volume 2, pages 2417–2492. Elsevier, 2016.

L. Finamor. Labor market informality, risk, and public insurance. 2022.

L. Gadenne, S. Norris, M. Singhal, and S. Sukhtankar. In-kind transfers as insurance. *American Economic Review*, forthcoming. URL https://example-url.com. Forthcoming.

S. Galiani, P. Gertler, and R. Bando. Non-contributory pensions. *Labour economics*, 38: 47–58, 2016.

J. Gardner. Two-stage differences in differences. *arXiv preprint arXiv:2207.05943*, 2022.

F. Gerard and G. Gonzaga. Informal labor and the efficiency cost of social programs: Evidence from unemployment insurance in brazil. *American Economic Journal: Economic Policy*, 13(3):167–206, August 2021. doi: 10.1257/pol.20180072. URL https://www.aeaweb.org/articles?id=10.1257/pol.20180072.

F. Gerard, J. Naritomi, and J. Silva. Cash transfers and formal labor markets: Evidence from brazil. *CEPR Discussion Paper No*, (DP16286), 2021.

P. K. Goldberg and T. Reed. Presidential address: Demand-side constraints in development. the role of market size, trade, and (in) equality. *Econometrica*, 91(6):1915–1950, 2023.

D. Gollin and J. P. Kaboski. New views of structural transformation: insights from recent literature. *Oxford Development Studies*, pages 1–23, 2023.

D. B. Gomes, F. S. Iachan, and C. Santos. Labor earnings dynamics in a developing economy with a large informal sector. *Journal of Economic Dynamics and Control*, 113: 103854, 2020. ISSN 0165-1889. doi: https://doi.org/10.1016/j.jedc.2020.103854. URL https://www.sciencedirect.com/science/article/pii/S0165188920300245.

A. Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277, 2021.

M. Guven, H. Jain, and C. Joubert. Social protection for the informal economy. *World Bank, Washington,*, (DC), 2021.

M. Hackmann, J. Heining, R. Klimke, M. Polyakova, and H. Seibert. Health insurance as economic stimulus? evidence from long-term care jobs. September 2023. Working Paper.

W. Huang and C. Zhang. The power of social pensions: Evidence from china's new rural pension scheme. *American Economic Journal: Applied Economics*, 13(2):179–205, 2021.

C. Joubert. Pension design with a large informal labor market: Evidence from chile. *International Economic Review*, 56(2):673–694, 2015.

C. Joubert and P. Kanth. Life cycle savings in a high-informality setting. 2022.

N. Kaushal. How public pension affects elderly labor supply and well-being: Evidence from india. *World Development*, 56:214–225, 2014.

A. Kraay. How large is the government spending multiplier? evidence from world bank lending. *The Quarterly Journal of Economics*, 127(2):829–887, 2012.

A. Kraay. Government spending multipliers in developing countries: evidence from lending by official creditors. *American Economic Journal: Macroeconomics*, 6(4):170–208, 2014.

W. A. Lewis. Economic development with unlimited supplies of labour. *The Manchester School*, 22:139–191, 1954.

M. Marcus and P. H. Sant'Anna. The role of parallel trends in event study settings: An application to environmental economics. *Journal of the Association of Environmental and Resource Economists*, 8(2):235–275, 2021.

K. Matsuyama. Engel's law in the global economy: Demand-induced patterns of structural change, innovation, and trade. *Econometrica*, 87(2):497–528, 2019. doi: https://doi.org/10.3982/ECTA13765. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA13765.

K. McKiernan. Social security reform in the presence of informality. *Review of Economic Dynamics*, 40:228–251, 2021.

P. Michaillat and E. Saez. Aggregate demand, idle time, and unemployment. *The Quarterly Journal of Economics*, 130:507–569, 2015.

P. Michaillat and E. Saez. Optimal public expenditure with inefficient unemployment. *The Review of Economic Studies*, 86(3):1301–1331, 2019.

K. M. Murphy, A. Shleifer, and R. Vishny. Income distribution, market size, and industrialization. *The Quarterly Journal of Economics*, 104(3):537–564, 1989.

E. Nakamura and J. Steinsson. Fiscal stimulus in a monetary union: Evidence from us regions. *American Economic Review*, 104(3):753–792, 2014.

OECD. *Pensions at a Glance 2023*. 2023. doi: https://doi.org/https://doi.org/10.1787/678055dd-en. URL https://www.oecd-ilibrary.org/content/publication/678055dd-en.

J. Roth and P. H. C. Sant'Anna. Efficient estimation for staggered rollout designs. *Journal of Political Economy Microeconomics*, 225(2):254–277, 2023.

J. Roth, P. H. Sant'Anna, A. Bilinski, and J. Poe. What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2):2218–2244, 2023.

J. C. S. Serrato and P. Wingender. Estimating local fiscal multipliers. Technical report, National Bureau of Economic Research, 2016.

L. Sun and S. Abraham. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199, 2021.

G. Ulyssea. Informality: Causes and consequences for development. *Annual Review of Economics*, 2020.

United Nations. World population ageing 2019: Highlights, 2019. URL https://www.un.org/en/development/desa/population/publications/pdf/ageing/WorldPopulationAgeing2019-Highlights.pdf. United Nations.

J. M. Wooldridge. Two-way fixed effects, the two-way mundlak regression, and difference-in-differences estimators. *SSRN Electronic Journal*, 2021.