Pension Reforms: Evidence from Brazilian Administrative Data

Juan Rios* Gustavo Gonzaga† Gabriel Lemos‡

November 21, 2025

Abstract

We study the welfare and fiscal effects of a major reform to the Brazilian Early Retirement Pension schedule and use these estimates to characterize how the schedule should be redesigned. The Brazilian pension schedule is determined by a "points" formula (age + years of contribution), and the 2015 reform introduced a threshold at which benefits rise discretely and the post-threshold schedule becomes flatter, corresponding to changes in the benefit level at the threshold (b_L) and in the slope of the post-threshold schedule (b_S). Using administrative data on the universe of early retirement claims matched to formal labor earnings, we find that the reform induced large behavioral responses: substantial bunching at the threshold reflects both postponement of claiming from below the threshold and anticipation from above it. We also quantify fiscal externalities by tracking payroll and income tax revenues for affected workers. The implied Marginal Value of Public Funds (MVPF) of the reform is about 0.21. We then use the welfare-weighted MVPF framework to evaluate marginal changes in the two policy parameters underlying the reform and to infer the optimal local budget-neutral reform to the schedule.

1 Introduction

Retirement pensions account for a substantial share of government spending. This is particularly evident in Brazil, where public expenditures on pensions reached 10.4% of GDP in 2019 (OECD, 2023a), compared to an OECD average of 7.7% in the same year (OECD, 2023c).

*PUC-Rio. Email: juanrios@econ.puc-rio.br. Applicant.
†PUC-Rio. Email: gonzaga@econ.puc-rio.br.
‡MIT Sloan. Email: gtlemos@mit.edu.

Yet Brazil remains relatively early in its demographic transition, with a much smaller share of elderly individuals than high-income countries.[1] The growing fiscal burden of pension programs has increased pressure on policymakers to reform their design, making the evaluation of alternative arrangements a central policy issue.

In this paper, we use the MVPF framework (Hendren and Sprung-Keyser, 2020) to conduct a welfare analysis of a reform to the Brazilian Early Retirement Pension schedule, which accounts for 35.4% of all old-age pension benefits. Our analysis quantifies the reform's fiscal and welfare effects and draws implications for optimal pension design in ageing countries. In addition to estimating the MVPF of the 2015 reform, we take a further step and evaluate how the benefit schedule should be adjusted going forward. Using the welfare-weighted MVPF framework of Bergstrom et al. (2024), we quantify the welfare consequences of marginal changes in each of these policy parameters and examine which policy parameters should be adjusted in a locally optimal budget-neutral reform.

The 2015 reform introduced a pension schedule defined as a function of "points" ($p =$ age + number of years of contribution at the time of claiming). This design eliminates the explicit age thresholds that often anchor retirement decisions in other countries. The reform substantially increased pension benefits once claimants reached a threshold number of points $\bar{p}$ (85 for women and 95 for men). In terms of our later optimal-reform framework, the policy operated through two margins. First, it raised the benefit level at the threshold, which we denote by b_L: claimants reaching $\bar{p}$ experienced a discrete increase in benefits relative to those with $p < \bar{p}$. Second, by introducing a plateau region for $p \geq \bar{p}$, the reform reduced the slope of the post-threshold schedule, which we summarize by the parameter b_S. These changes increased the incentive to defer claiming until reaching $\bar{p}$ while reducing the marginal return to delaying claiming beyond the threshold.

Using administrative data on pension claims and formal employment in Brazil, we estimate the MVPF of this reform in three steps. First, we estimate how the reform affected the distribution of claimants' points. Second, we identify its impact on the average benefit at each point, capturing selection effects among claimants. These two estimates allow us to separate the mechanical and behavioral components of the reform's impact on total pension payments. Assuming that these benefits provide no insurance value,[2] the mechanical effect represents the willingness to pay for the reform. In the third step, we estimate the reform's impact on payroll and income tax revenues using administrative data on formal labor income. These estimates capture the main fiscal externalities generated by the policy, allowing us to

[1] In 2022, the ratio of individuals aged 65 and older to working-age individuals (20–64) was 15.8 in Brazil, compared to 31.3 across OECD countries (OECD, 2023b).

[2] Retirement benefits may insure older workers against negative labor market shocks.

construct the MVPF of the reform.

Beyond evaluating the 2015 reform, we use our estimated behavioral responses and fiscal externalities to characterize how the pension schedule should be redesigned. We decompose claiming behavior around the threshold into postponement responses to the benefit level at the threshold (b_L) and anticipation responses to the slope of the post-threshold schedule (b_S). Combining these estimates with tax-revenue responses, we can compute welfare-weighted MVPFs for marginal changes in each parameter. This mapping allows us to identify whether the optimal local reform should raise the benefit level at the threshold, flatten or steepen the post-threshold plateau, or reallocate generosity across both margins. The final section of the paper applies this framework to characterize the optimal direction of reform for the Brazilian Early Retirement Pension schedule.

2 Background and the 2015 Reform

In Brazil, pension benefits are determined by formal workers' contribution wages, age, and number of years of contribution at the time of claiming. Importantly, the claiming decision is entirely independent of the retirement decision (leaving formal employment). Workers may claim their pension while continuing to work, as the system imposes no earnings test or reduction in benefits for ongoing employment.

Figure 1 plots the stylized replacement rates for women and men as a function of points. Before 2015, the replacement rate increased almost linearly with points.[3] The 2015 reform sharply raised the replacement rate to 100% once claimants reached a threshold of points $\bar{p}$ set at 85 for women and 95 for men. At those thresholds, the replacement rates were approximately 70% and 80%, respectively. Beyond the threshold, the replacement rate no longer increased with points, creating a plateau region.[4]

3 Data

We use four complementary data sources in our analysis. First, to investigate the effects on benefit payments, we rely on novel administrative records from SUIBE (the Single System for Social Security Benefit Information), which contain the universe of formal-sector early

[3]The actual replacement rate, known as *Fator Previdenciário*, was determined by a formula depending on age, years of contribution, and life expectancy. The empirical averages of these replacement rates conditional on points closely follow this stylized schedule.

[4]The schedule resumes increasing slightly beyond the plateau for very high point levels, although few individuals claim benefits in that range.

retirement claims in Brazil from 2012 to 2019 (approximately 2.7 million workers).[5] The dataset provides detailed information on benefit amounts, claiming dates, replacement rates, and contribution histories.

Second, we merge SUIBE with RAIS (the Social Information Annual Report), an employer–employee administrative dataset from the Ministry of Labor that covers the universe of formal workers and firms in Brazil. Because SUIBE reports only cumulative years of contribution at the time of claiming, RAIS is used to reconstruct workers' contribution trajectories before and after claiming. RAIS also records formal earnings and tax-relevant wage information, which allows us to measure payroll and income tax payments over time and quantify the fiscal externalities generated by behavioral responses to the reform.

Third, we use microdata from the Pesquisa de Orçamentos Familiares (POF), Brazil's nationally representative household expenditure survey. POF provides detailed consumption information for the general population and for households receiving different types of government transfers. We use these data to approximate consumption levels of early retirees relative to the broader population, which is needed for our welfare calculations and for normalizing marginal utilities in our welfare-weighted MVPF estimates.

Finally, we draw on ELSI (Estudo Longitudinal da Saúde dos Idosos Brasileiros), a nationally representative longitudinal survey of older adults in Brazil. ELSI includes information on age, years of contribution to the pension system, and benefit receipt. Because the Brazilian points formula depends on age and contribution years, these variables allow us to map individuals in ELSI to point levels and estimate average consumption by points among pension beneficiaries—an important input for estimating welfare weights.

4 Reduced-Form Evidence

Figure 2a plots the frequencies of points at claiming relative to the cutoff ($p-\bar{p}$) before (green) and after the reform (orange). This distribution was smooth around $\bar{p}$ before the reform, but there is substantial bunching at $\bar{p}$ after the reform, indicating significant claiming responses to the reform. Figure 2b plots the claiming hazards at each relative point, i.e., the probability of a worker having claimed at $p-\bar{p}$ given that she has not claimed up to that point. After the reform, these claiming hazards (i) remained stable for workers with fewer than $\bar{p}-6$ points, (ii) declined for individuals with points in the range $(\bar{p}-6, \bar{p})$ indicating claiming deferral, and (iii) increased for individuals claiming with more than $\bar{p}$ points. This rise is sharp immediately above the threshold and more gradual thereafter. The sharp increase

[5] SUIBE also includes data for age-based retirement claims, which were not affected by the reform.

reflects bunching from both sides of $\bar{p}$, while the gradual rise in claiming hazards at higher point levels indicates anticipation of claiming beyond the threshold due to the reduction in incentives to defer in the plateau region.

5 Research Design

We construct the MVPF in three steps. First, we estimate the impact of the reform on the hazard of claiming retirement benefits. Workers with normalized points $p - \bar{p} < -6$ (i.e., those who claimed at least three years before reaching the threshold) form the control group, as their claiming hazards were unaffected by the reform. Treatment groups are defined over intervals of normalized points $(p - \bar{p}) \in \tilde{p}$, where $\tilde{p} = [-6, -3), [-3, 0), [0, 2), [2, 7), [7, \infty)$. The identifying assumption is that, absent the reform, claiming hazard rates for all treatment groups would have evolved in parallel to those of the control group. We estimate the following difference-in-differences specification:

$$CH_{ipt} = \sum_{k \neq 14Q4} \beta_{k,\tilde{p}} \, 1\{(p - \bar{p}) \in \tilde{p}\} \times 1\{t = k\} + \delta_t + \gamma_{p-\bar{p}} + \zeta X_{ip} + \varepsilon_{ipt}, \tag{1}$$

where CH_{ipt} denotes the claiming hazard of individual i with normalized points $p - \bar{p}$ in period t. The coefficients $\beta_{k,\tilde{p}}$ measure the causal impact of the reform on the claiming hazard for individuals in group $\tilde{p}$ during period k. The terms δ_t and $\gamma_{p-\bar{p}}$ are time and normalized-points fixed effects, respectively, controlling for common temporal shocks and systematic differences across groups. The vector X_{ip} includes individual covariates (gender, microregion, schooling, race, and birth year) and ε_{ipt} is an idiosyncratic error term.

Figure 3a plots the raw claiming hazards for different $\tilde{p}$ groups. These hazards were roughly parallel to those of the control group ($\tilde{p} = [-15, -6)$) before the reform. After the reform, however, points below the threshold exhibit lower claiming hazards than the control group, whereas points above the threshold show higher claiming hazards. This variation enables us to estimate the difference-in-differences coefficients precisely for each quarter following the reform and to construct counterfactual claiming hazards by points and quarter. Integrating these hazards over points yields the counterfactual density of claimants. Figure 3b compares this counterfactual density with the observed post-reform distribution. Under our identifying assumptions, the figure indicates that the reform had a large and statistically meaningful effect on the distribution of points. The plot also reveals the composition of the bunching around the threshold: while some individuals deferred claiming in response to the higher benefits moving from below the threshold, others anticipated claiming above it, likely due to

the reduction in marginal incentives to postpone once the threshold was reached.

We cannot directly compute counterfactual benefit outlays from the counterfactual points distribution because the reform may have altered the average benefit paid at each point. For instance, the region just above the threshold may have attracted higher-wage claimants if wages and knowledge about the reform are correlated. To account for these composition effects, we construct the observed pre-reform average benefit $\bar{b}^o_{p,t}$ as the average benefit that would have been paid to individuals claiming with p points in quarter t under the pre-reform schedule, but using post-reform choices (wages, age, and years of contribution). Our goal is to estimate the counterfactual observed pre-reform benefits that would have prevailed in the absence of the reform. Because points below $\bar{p} - 6$ were unaffected by the reform, they provide a natural control group. Under the identifying assumption that trends in $\bar{b}^o_{p,t}$ for control and treatment points would have remained parallel in the absence of the reform, the coefficient $\beta^{pre}_{k,\tilde{p}}$ in the following specification captures the reform's effect on the average pre-reform benefits:

$$\bar{b}^o_{p,t} = \delta^{pre}_p + \gamma^{pre}_t + \sum_{k \neq -1} \beta^{pre}_{k,\tilde{p}} 1(p \in \tilde{p}) * 1(t = k) + \varepsilon^{pre}_{p,t} \tag{2}$$

Here, δ^{pre}_p and γ^{pre}_t are point-level and time fixed effects, respectively, $\varepsilon^{pre}_{p,t}$ is the error term, and $1(p \in \tilde{p})$ identifies whether point level p falls within a treatment region.

Figures 4a and 4b plot the estimated coefficients $\hat{\beta}^{pre}_{k,\tilde{p}}$ for regions immediately below ($\tilde{p} = [-3, -1]$) and above ($\tilde{p} = [0, 3]$) the threshold, respectively. We find no significant pre-trend differences between these groups and the control region. After the reform, however, the estimates show a negative effect on the average pre-reform benefit below the threshold and a positive effect above it, indicating positive selection of bunchers in terms of benefit payments. We construct the average counterfactual benefit at each point as the difference between the observed post-reform average and the estimated selection effect. Combining these counterfactual averages with the counterfactual point density, we estimate counterfactual benefit payments in the absence of the reform. This, in turn, allows us to compute the additional pension expenditures attributable to the reform.

Finally, we need to identify what share of the additional pension expenditures was mechanical. To do so, we estimate the amount of benefits that would have been paid if the reform had occurred but individuals had not adjusted their claiming behavior. This requires both the counterfactual distribution of points and the average benefits under the reformed schedule at each point, conditional on no behavioral response. To construct these counterfactual benefit averages, we estimate the same specification as above but using as the dependent variable the average benefits at each relative point $p - \bar{p}$ and quarter t under the

reformed schedule, $\bar{b}'^{o}_{p,t}$:

$$\bar{b}'^{o}_{p,t} = \delta_p^{post} + \gamma_t^{post} + \sum_{k \neq -1} \beta_{k,\tilde{p}}^{post} 1(p \in \tilde{p}) * 1(t = k) + \varepsilon_{p,t}^{post}, \tag{3}$$

All coefficients are defined as before. Under the identifying assumption that trends in these post-reform benefits would have remained parallel across treatment and control point groups in the absence of the reform, the estimator $\beta_{k,\tilde{p}}^{post}$ captures the effect of the reform on the average post-reform benefits.

Figure 5 shows the estimated coefficients for these effects. Once again, we find no differential pre-trends before the reform and a significant change afterward, consistent with the patterns discussed above. These estimates allow us to compute the counterfactual average benefits that would have been paid under the post-reform schedule in the absence of behavioral responses.

6 Measuring the MVPF, Welfare Effects and Other Fiscal Externalities

The MVPF is defined as the ratio of the willingness to pay for a reform (WTP) to its costs net of its fiscal externalities (C):

$$MVPF = \frac{WTP}{C}.$$

Under the assumption that these pension benefits have no insurance value,[6] the WTP for the reform corresponds to the mechanical increase in benefits described above. Considering only the fiscal externalities associated with benefit payments, we can recover the total cost of the reform. Figure 6a plots the mechanical effects (red line) and the fiscal externalities (blue lines). This preliminary calculation suggests an MVPF of approximately 0.21.

There are, of course, other potential sources of fiscal externalities. For instance, higher pension benefits may induce individuals to retire earlier, thereby reducing government revenues from income and payroll taxes. We exploit the fact that workers earning less than 1.25 times the minimum wage received minimum-wage benefits both before and after the reform, owing to the existence of a benefit floor. Using administrative data on formal employment and earnings in Brazil, we construct the payroll and income taxes paid by these claimants over time. To estimate the effect of the reform, we implement a difference-in-differences design

[6]Kolsrud et al. (2024) estimate that most of the welfare gains from pensions in Sweden arise from redistributional rather than insurance effects.

that compares the growth in tax collections between claimants whose contribution wage was above and below 1.25 times the minimum wage:[7]

$$\Delta\tau_{i,t} = \delta_i^{tax} + \gamma_t^{tax} + \sum_{k\neq -1} \beta_k^{tax} 1(w_{it} > 1.25MW) \times 1(t = k) + \varepsilon_{it}^{tax}$$

Here, $\Delta\tau_{i,t}$ denotes the growth in tax payments for individual i at time t; δ_i^{tax} and γ_t^{tax} are individual and time fixed effects, respectively; and β_k^{tax} captures the effect of the reform on tax revenues for higher-wage claimants. Figure 6b presents the results, which suggest a negative effect of the reform on claimants' tax payments, implying that the MVPF is lower than our baseline estimate.

We also compute welfare effects of redistribution toward the elderly by normalizing the marginal utility of consumption to the national average (e.g., Gruber (1997)). To do so, we apply a standard consumption-based welfare approximation that combines consumption measures with an assumed coefficient of relative risk aversion of 2. Two empirical patterns allow us to implement this calculation. First, Figure 7a shows that average consumption is similar for claimants below and above the threshold, both before and after the reform, suggesting that the reform generated little redistribution within the beneficiary pool. Second, Figure 7b indicates that the average pension beneficiary (labeled "INSS" in the figure) consumes about 7.5% less than the average Brazilian. This latter fact implies that, in this setting, welfare effects are smaller than mechanical effects (as indicated in Figure 6a) because the marginal utility of consumption for early retirees is only moderately higher than the population average.

7 Disentangling Responses and the Optimal Reform

The analysis above characterizes the welfare and fiscal effects of the 2015 reform, treating the existing schedule as given. In this section, we take a first step toward an optimal reform analysis of the pension schedule. Bergstrom et al. (2024) show that the optimal local budget-neutral reform consists of increasing (decreasing) the policy parameter with the highest (lowest) WMVPF.

We parameterize the benefit schedule around the threshold $\bar{p}$ using two policy levers. The first lever, b_L, captures the increase in benefits at the threshold; an increase in b_L strengthens incentives to defer claiming for $p < \bar{p}$ in order to reach the higher benefit level at the threshold.

[7]To ensure comparability, we exclude claimants with contribution wages above five times the minimum wage from the sample.

The second lever, b_S, measures the slope of the post-threshold schedule; a decrease in b_S reduces the marginal return to delaying claiming once the threshold has been reached. It is useful to consider the two counterfactual reforms in Figure 8. The first, in panel a, only increased b_L, inducing *postponement*: individuals who would have claimed with $p < \bar{p}$ in the absence of the reform wait until they accumulate enough points to claim exactly at $\bar{p}$ or slightly above it. The second, in panel b, only decreased b_S, inducing *anticipation*: individuals who would have claimed with $p > \bar{p}$ in the absence of the reform bring their claims forward to $\bar{p}$ or just above it, because the flatter plateau reduces the value of deferring further. The actual reform described in Figure 1 was a combination of these two counterfactual reforms. To describe the optimal reform $(\Delta b_L^*, \Delta b_S^*)$, we need to estimate the welfare-weighted MVPF for each of these counterfactual reforms.

Figure 3b, comparing the actual and counterfactual densities of claimants by points at the end of the sample period, provides one key input for distinguishing behavioral responses to b_L from those to b_S. Observed bunching at the threshold therefore reflects a combination of postponed claims from below and anticipated claims from above. The difference between the observed and counterfactual densities on each side of $\bar{p}$ allows us to separate these two types of behavioral responses as indicated in the figure. It is clear that most of these responses come from the left of the graph, indicating that there are more individuals postponing in response to Δb_L than anticipating in response to Δb_S. To measure the impact of each of these groups on benefit outlays, we need to adjust these responses for selection as described above.

To translate these density changes into welfare-weighted MVPFs for b_L and b_S, we need to further decompose the reform's effects into mechanical and behavioral components. The mechanical share of the total effect can differ across these parameters because each margin generates a distinct pattern of fiscal externalities. For anticipation responses driven by $\Delta b_S < 0$, the reduction in post-threshold marginal benefits lowers government expenditures since claimants receive lower benefits at a given point. At the same time, earlier claiming increases the number of benefit-payment periods and thus raises outlays.[8] The postponement responses induced by $\Delta b_L > 0$ create a different trade-off in terms of policy costs. On the one hand, claimants who postpone receive higher benefits once they eventually claim, increasing average benefit levels at a given point. On the other hand, postponement mechanically reduces the number of benefit-payment periods because individuals spend more time in the no-benefit state before claiming. The net fiscal impact of b_L therefore depends on the balance between higher benefit levels and fewer payment periods, as well as on the induced changes

[8]All these responses also generate fiscal externalities by altering labor supply and tax payments in the periods before and after claiming.

in payroll and income tax revenues during the additional working years.

The decomposition by b_L and b_S will reveal whether the low aggregate WMVPF is driven primarily by the anticipation margin associated with b_S, by the postponement margin associated with b_L, or by the interaction between the two. This, in turn, will inform whether an optimal reform should increase the benefit level and reduce the slope (as in the current reform), or increase the slope and reduce benefits, thereby reducing incentives to claim beyond the threshold while increasing marginal benefits for those who reached it.

8 Conclusion

This paper provides new evidence on the welfare and fiscal consequences of a major reform to the Brazilian Early Retirement Pension schedule. Using rich administrative data on pension claims and formal labor income, we show that the 2015 reform generated substantial bunching at the points threshold, driven by both claiming postponement and anticipation. Our empirical strategy disentangles mechanical and behavioral components of the resulting increase in benefit payments and quantifies the associated fiscal externalities through changes in payroll and income tax revenues. Under the assumption that early retirement benefits provide negligible insurance value, our preliminary estimates point to an MVPF well below one, indicating that, at the margin, the resources devoted to this reform could have generated larger welfare gains if allocated to alternative policies.

We then move beyond a retrospective evaluation of the 2015 reform and use the welfare-weighted MVPF framework to study how the pension schedule should be redesigned. By decomposing behavioral responses into postponement reactions to the benefit level at the threshold (b_L) and anticipation reactions to the slope of the post-threshold schedule (b_S), we characterize the fiscal and welfare implications of marginal changes in each policy parameter. This decomposition will allow us to determine whether an optimal local reform should tilt the schedule toward higher benefits at the threshold, a flatter or steeper post-threshold region, or some combination of both. More broadly, our approach illustrates how sufficient-statistics methods can be used to evaluate and redesign pension systems in ageing economies, linking detailed administrative data to the normative question of how governments should allocate scarce fiscal resources across competing social insurance programs.

References

Bergstrom, K., Dodds, W., and Rios, J. (2024). Optimal policy reforms. Technical report, Texto para discussão.

Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *The American Economic Review*, 87(1):192–205.

Hendren, N. and Sprung-Keyser, B. (2020). A unified welfare analysis of government policies. *The Quarterly journal of economics*, 135(3):1209–1318.

Kolsrud, J., Landais, C., Reck, D., and Spinnewijn, J. (2024). Retirement consumption and pension design. *American Economic Review*, 114(1):89–133.

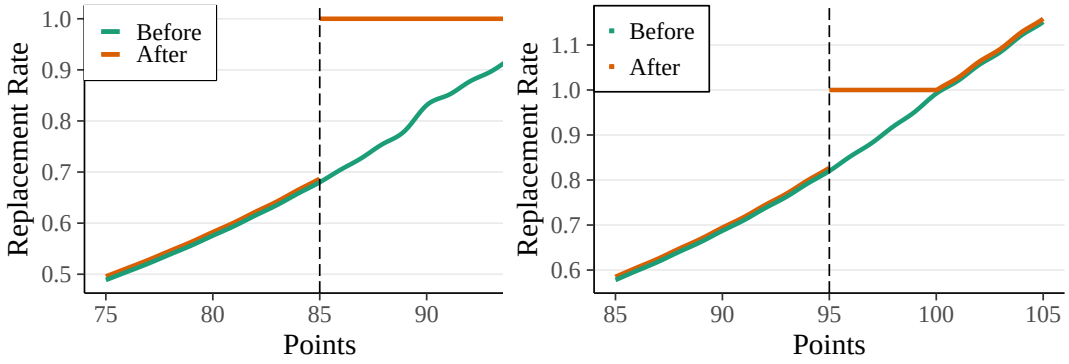
OECD (2023a). *OECD Economic Surveys: Brazil 2023*. OECD Publishing.

OECD (2023b). *Pensions at a Glance 2023: OECD and G20 Indicators*. OECD Publishing.

OECD (2023c). Sizing up welfare states: How do oecd countries compare?

Figures

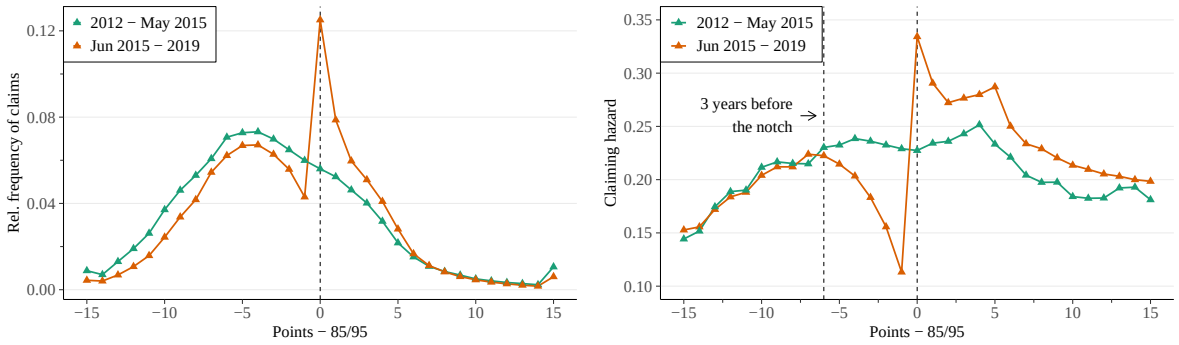
Figure 1: Replacement Rate Before and After the 2015 Reform.



(a) Women (b) Men

Notes: Panel (a) plots the average replacement rate for women as a function of points (age + years of contribution) before and after the 2015 Pension Reform. Panel (b) presents the same comparison for men. The sample includes all formal private-sector early retirement claims in Brazil between 2012 and 2019. Replacement rates are calculated using the statutory rules in place at the time of claiming, based on workers' recorded contribution history. The 2015 reform increased benefits discretely at the point thresholds of $\bar{p} = 85$ (women) and $\bar{p} = 95$ (men) and introduced a plateau region for $p \geq \bar{p}$, with full replacement rates after the threshold. Replacement rates before the reform correspond to the linear Fator Previdenciário schedule. Replacement rates after the reform reflect the post-2015 rules, which raised benefit levels at $\bar{p}$ and flatten the schedule beyond the threshold.

Figure 2: Reduced Form Evidence

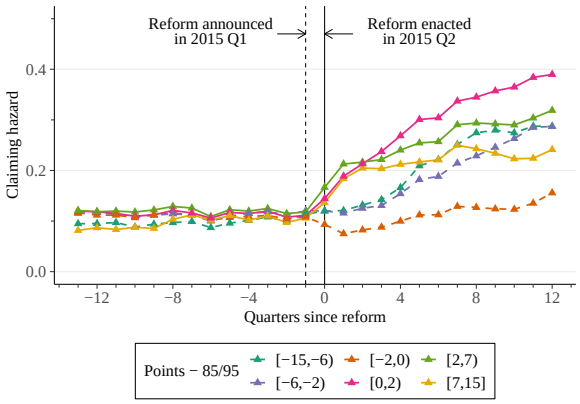


(a) Claiming Frequencies

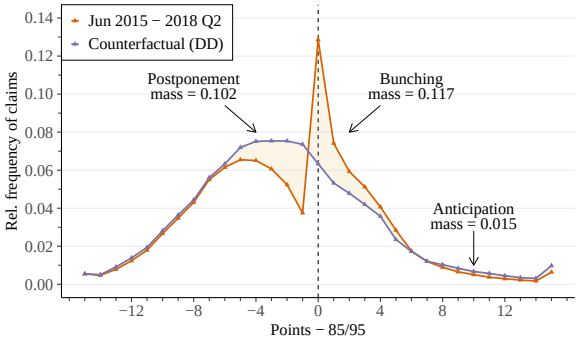
(b) Claiming Hazard

Notes: The figure documents retirement claiming dynamics relative to the cutoff ($p - \bar{p}$). Panel (a) plots the distribution of points at claiming relative to the threshold $\bar{p}$ for the pre-reform (2012–May 2015) and post-reform (June 2015–2019) periods.Before the reform, the distribution is smooth around the threshold, whereas after the reform we observe a sharp spike around $\bar{p}$, indicating substantial bunching. Panel (b) reports the claiming hazard at each normalized point level, defined as the probability of claiming at a given point conditional on not having claimed earlier. Claiming hazards were stable across groups before the reform but diverge sharply afterward: claiming hazards decrease for individuals approaching $\bar{p}$ from below, consistent with postponement of retirement claiming for individuals approaching $\bar{p}$, and increase just above the threshold, consistent with anticipation of claiming for individuals that were already past the threshold during the Reform, due to reduction in incentives. Together, these patterns provide reduced-form evidence of strong behavioral responses to the 2015 reform.

Figure 3: Effects on Claiming



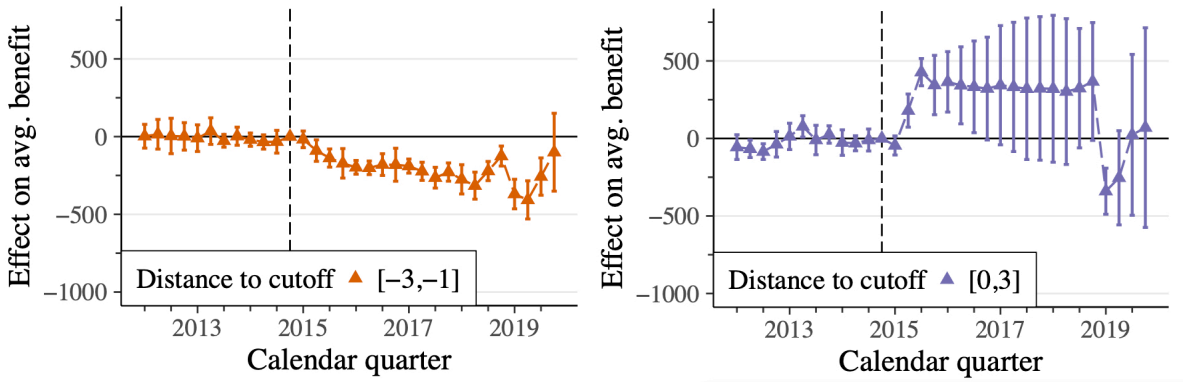
(a) Claiming Hazards for Different Points



(b) Actual and Counterfactual Claiming Densities

Notes: The figure documents the 2015 Reform's effect on claiming hazards. Panel (a) plots the evolution of the claiming hazard for individuals at different normalized point levels relative to the threshold $\bar{p}$. Claiming hazards are defined as the probability of claiming in a given quarter conditional on not having claimed previously. The vertical lines indicate the announcement (2015Q1) and implementation (2015Q2) of the reform. The sample includes all formal private-sector early retirement claims from the SUIBE administrative records between 2012 and 2019. Individuals are grouped into six point ranges based on their distance to the threshold at the time of the reform: $[-15,-6)$, $[-6,-2)$, $[-2,0)$, $[0,2)$, $[2,7)$, and $[7,15]$. Individuals with $p-\bar{p}<-6$ serve as the control group, as their hazard rates remain unchanged across the reform period. Panel (b) reports estimates based on the event–study difference-in-differences specification described in equation 3. For each normalized point group $(p-\bar{p})\in\tilde{p}$, we estimate a linear probability model for the quarterly claiming hazard, including point fixed effects, quarter fixed effects, and controls for gender, age at the beginning of the year, education, and average earnings rank between 2010 and 2012. The regression sample contains one quarterly observation for each individual who had not yet claimed, resulting in a balanced panel of hazard-eligible individuals from 2012Q1 to 2019Q4. The coefficients $\beta_{k,\tilde{p}}$ trace the relative change in the hazard for each point group $\tilde{p}$, using individuals with $p-\bar{p}<-6$ as the control group. To build the counterfactual claiming density, we predict hazards from the same model while shutting down all post-reform treatment effects (i.e., imposing $\beta_{k,\tilde{p}}=0$ for all $k\geq$ 2015-Q2). Aggregating these predicted hazards yields the claiming hazard distribution that would have prevailed at each point had the behavioral responses not happened. The comparison between the actual and counterfactual distributions allows us to decompose the total excess mass at $\bar{p}$ into postponement mass (from below the threshold), anticipation mass (from above), and the resulting bunching mass at the cutoff.

Figure 4: Effect on Average Benefits under the Original Schedule

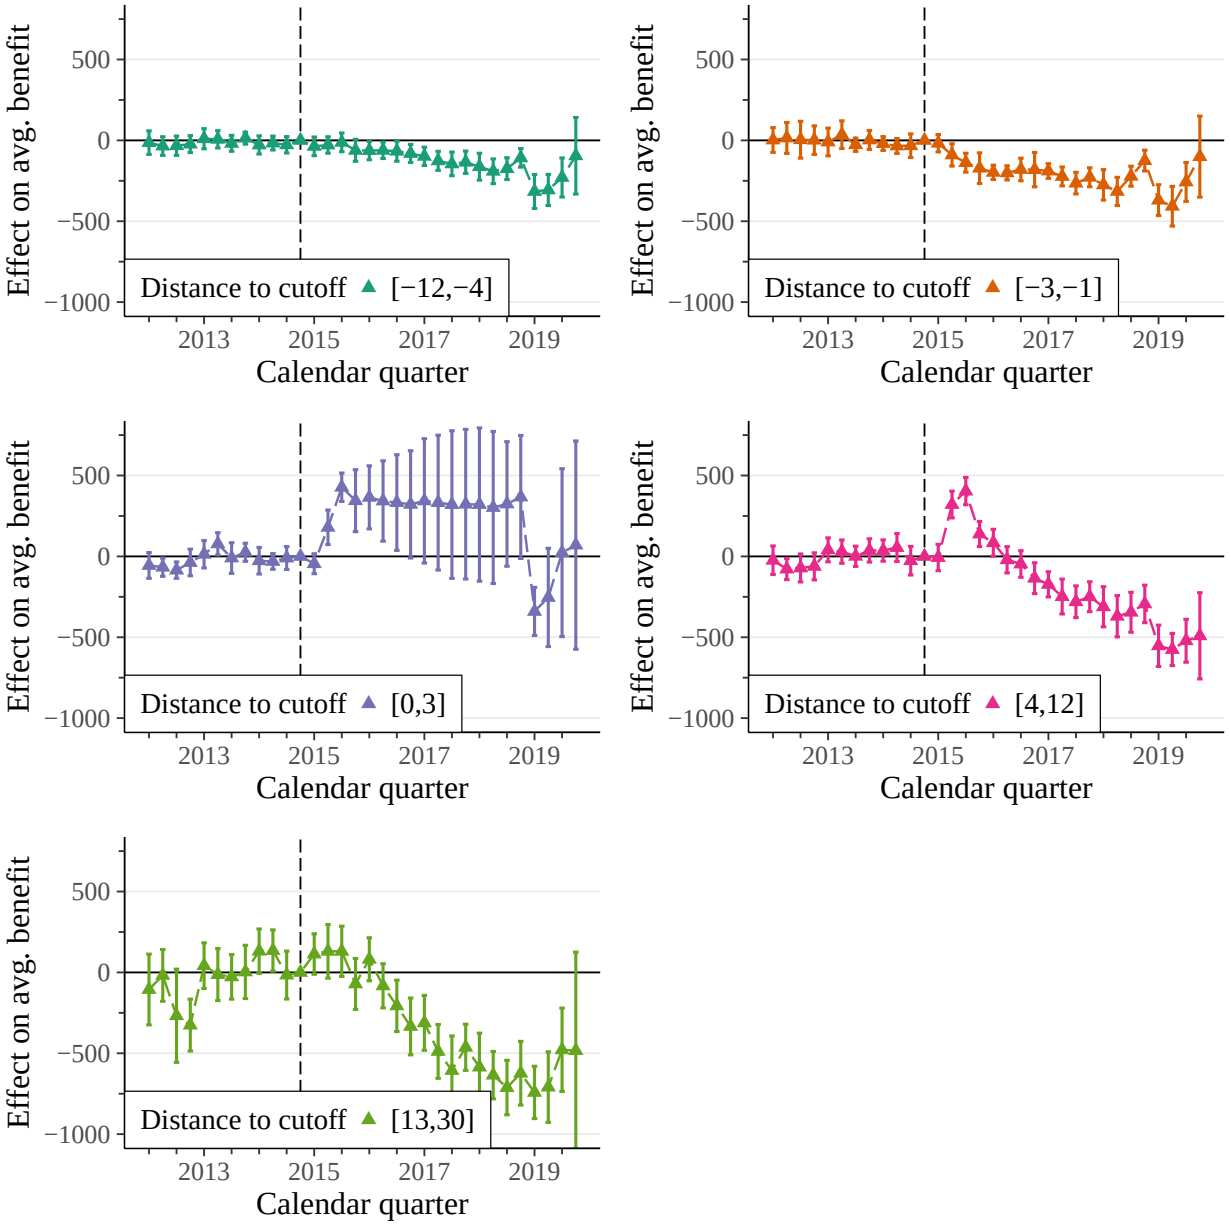


(a) Below $\bar{p}$ (b) Above $\bar{p}$

Notes: The figure displays the estimated impact of the 2015 reform on average pension benefits computed under the pre-reform benefit schedule. The estimates are obtained from the baseline specification 2, which relates realized pension benefits to an indicator for being affected by the reform, interacted with normalized point groups. The regression sample includes all formal private-sector early claimants in the SUIBE records between 2012 and 2019 whose benefits can be fully reconstructed under the statutory pre-2015 rules. We include fixed effects for gender, cohort, calendar quarter of claiming, and contribution density, as well as controls for education, age, years of contribution, andthe average earnings between 2010 and 2012. Standard errors are clustered at the point-group level.

The plotted coefficients represent the average difference between the actual benefit level received after the reform and the counterfactual benefit level that would have been received under the original (pre-2015) schedule, conditional on claiming at each normalized point $(p - \bar{p})$. In panel (a), negative estimates for individuals just below the threshold reflect the fact that the reform induced postponement: these individuals would have claimed earlier (at a lower point value) and therefore qualified for a lower benefit under the original schedule. On the other hand, in panel (b), positive estimates for those just above the threshold reflect early claiming driven by anticipation. The magnitude and asymmetry of these effects quantify how behavioral adjustment around the point threshold is mapped into differences in benefit entitlements.

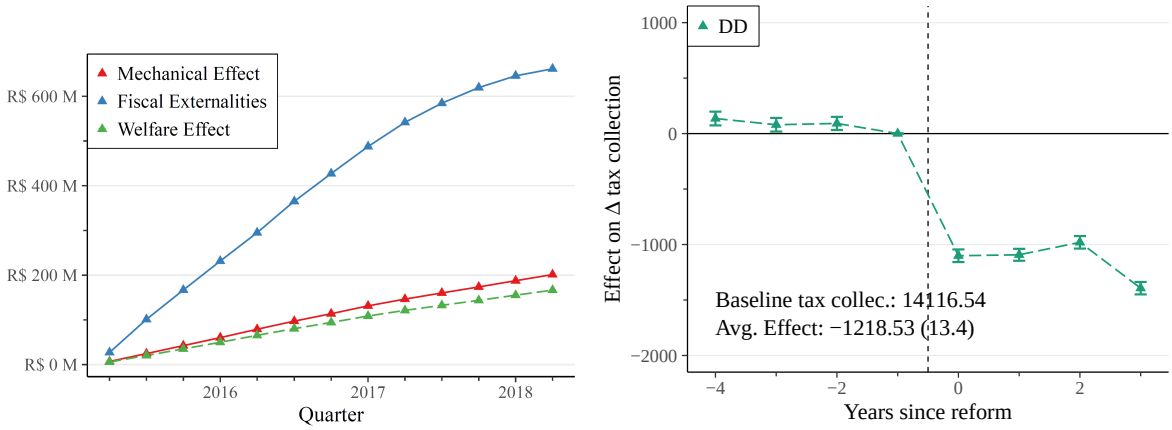
Figure 5: Effect on Average Benefits under the Reformed Schedule



Notes: The figure plots the estimated effect of the 2015 reform on average pension benefits evaluated under the post-reform schedule for five ranges of distance to the threshold $\bar{p}$ ($[-12,-4]$, $[-3,-1]$, $[0,3]$, $[4,12]$, and $[13,30]$). Estimates come from the same specification used in Figure 4, but with the dependent variable equal to the average benefit paid under the new rules at each point–quarter cell. The model includes point and time fixed effects, and points with $(p-\bar{p}) < -6$ act as the control group. Under the identifying assumption of parallel trends between treatment and control points, the coefficients capture the causal effect of the reform on average benefits under the new schedule.

No differential pre-trends are observed prior to the reform. After the 2015 Reform, average benefits fall for point groups just below the threshold (e.g., $[-3,-1]$), reflecting that post-reform claimants tend to have lower earnings histories than pre-reform claimants at the same point level. On the other hand, benefits rise for groups just above the threshold (e.g., $[0,3]$ and $[4,12]$), consistent with positive selection of individuals into these ranges.Effects are smaller for more distant point groups (e.g., $[-12,-4]$ and $[13,30]$), as expected if behavioral responses get weaker with distance from the cutoff.

Figure 6: Inputs for MVPF

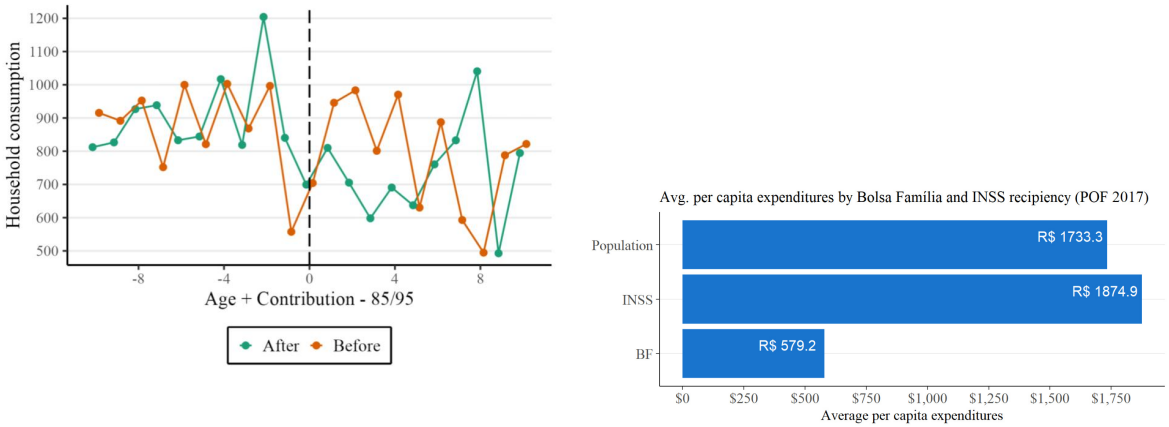


(a) Mechanical and Behavioral Effects

(b) Effect on Tax Collection Growth

Notes: Panel (a) plots the mechanical effect of the reform on pension expenditures, the fiscal externalities from changes in labor supply and earnings, and the implied welfare effect used in the MVPF calculation. The mechanical component is obtained by combining the counterfactual point density from Figure 3b with the counterfactual benefit averages constructed in Figures 4 and 5. Fiscal externalities are computed by applying the estimated semi-elasticity of taxable earnings with respect to claiming behavior to the reform-induced changes in claiming. The welfare effect corresponds to the sufficient-statistic expression in Section 5, which aggregates mechanical gains and fiscal externalities over time.
Panel (b) reports estimates from a difference-in-differences event-study specification for quarterly tax collections. The regression sample includes all formal-sector workers observed in the SUIBE tax records between 2011 and 2019. The dependent variable is the deviation of tax collections from their pre-reform trend, and the model includes state, year, and quarter fixed effects, as well as controls for age, gender, and education. For each event time k relative to the reform implementation quarter, we estimate a coefficient that measures the change in tax collection compared to the reference period ($k = -1$). Standard errors are clustered at the state level.

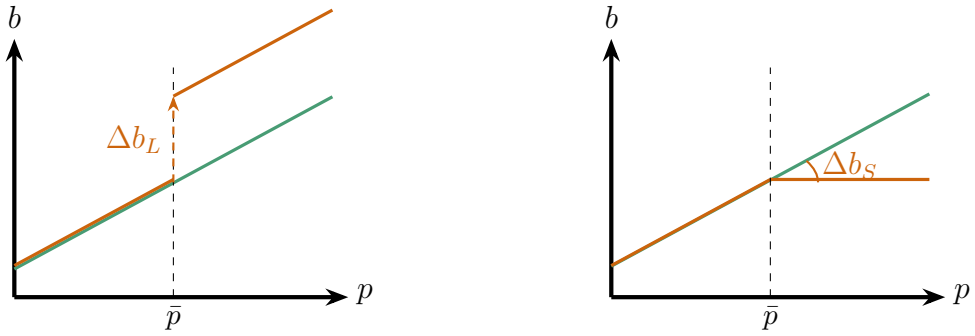
Figure 7: Consumption Measures



(a) Average Consumption by Points (ELSI) (b) Average Consumption by Program Status

Notes: The figure provides descriptive patterns of consumption across heterogeneous groups. Panel (a) plots average household consumption by points using data from the 2019 wave of ELSI, a nationally representative survey of older Brazilians. Consumption is constructed from self-reported expenditures across durable and non-durable categories following the harmonized ELSI methodology. Each plotted value corresponds to the mean consumption level among individuals falling into a given point interval around the claiming threshold $\bar{p}$. The sample is restricted to individuals aged 55 and above with complete information on household composition and labor-market histories.
Panel (b) reports average consumption by program-status group using microdata from the 2017–2018 POF (Pesquisa de Orçamentos Familiares). Consumption is measured as total household expenditures net of imputed rents and standardized by household size using the adult-equivalence scale adopted in the survey. Individuals are classified according to their program status at the time of the survey (non-beneficiary, early claimant, normal claimant, or late claimant). Within each group, we compute the mean level of household consumption. These descriptive patterns provide external validation for the direction of consumption differences across groups documented in the administrative estimates presented in the main text.

Figure 8: Two Counterfactual Reforms



(a) Pure Benefit Level Reform

(b) Pure Benefit Slope Reform

Notes: The figure illustrates the two counterfactual policy levers used to characterize local reforms of the pension schedule around the threshold $\bar{p}$. Panel (a) depicts a pure benefit-level reform, represented by an increase in b_L, which raises the benefit paid at $\bar{p}$ while keeping the slope of the schedule unchanged. Such a shift strengthens the incentive to postpone claiming for individuals with $p < \bar{p}$, leading them to wait until they reach the threshold. The vertical distance Δb_L in the figure captures the size of this level adjustment. Panel (b) shows a pure benefit-slope reform, represented by a decrease in b_S, which flattens the post-threshold schedule while leaving the benefit at $\bar{p}$ fixed. This reduces the marginal return to delaying claiming once the threshold is reached, inducing individuals who would have claimed with $p > \bar{p}$ to bring their claims forward toward $\bar{p}$. The gap Δb_S captures the slope change relative to the original rules.