Labor supply responses to a Very Large Income Shock: Evidence from Daughters of Brazilian Military Officers.

Marcelo Barbosa Ferreira*

October 28, 2025

Updated version

Abstract

This paper provides estimates of income effects on labor supply by studying a large, permanent, and unconditional monthly cash flow. I leverage a unique setting in Brazil where adult daughters of military personnel receive a lifetime pension after their parent's death. This transfer is exceptionally large, the median pension payment is 282% of the national median earnings, and truly unconditional, available regardless of their employment, earnings, or marital status. Using administrative pension records matched to linked employer-employee data, I find that this transfer induces a 10% decrease on the extensive margin of labor supply after seven years, with no effect on the intensive margin, and a marginal propensity to earn (MPE) of -0.08 after seven years. Labor supply adjusts gradually, driven by a large temporary increase in the hazard of separation and a persistent moderate reduction in the job finding rate. The marginal propensity to earn is greater for larger pension payments, consistent with adjustment costs, and larger for older daughters, consistent with binding liquidity constraints and desire for early retirement. Using a sufficient statistics model, our MPE estimate of -0.08 implies that the redistributive gains from a UBI in Brazil would outweigh its efficiency costs, resulting in net welfare gains, as long as the uncompensated elasticity of labor supply is less than 0.38, consistent with microeconomic evidence.

Keywords: Universal Basic Income, Welfare Analysis, Labor Supply.

JEL Classification: JEL H55, J13, J22, M13.

*Princeton University. E-mail: marcelof@princeton.edu

1 Introduction

The relationship between unearned income and individual economic behavior, particularly labor supply, is a cornerstone of public and labor economics. This relationship is a key parameter for one of the most significant and enduring policy debates: the equity-efficiency trade-off. Understanding the income effect, how individuals adjust their work effort when given money unconditionally, is critical for several reasons. First, it is directly related to the efficiency costs, or deadweight loss, of transfer programs. If transfers cause a large withdrawal from the labor force, the tax base shrinks, making redistribution more costly. Second, the income effect is a key component of the Slutsky equation, which connects the uncompensated labor supply elasticity to the compensated elasticity. This compensated elasticity, which captures the pure substitution effect, is a key parameter for calculating the deadweight loss of taxation. Third, the income effect provides a theoretical link between the uncompensated elasticity and the Frisch elasticity, the response to temporary wage changes keeping marginal utility of wealth fixed, which is a key parameter for calibrating macroeconomic models of business cycle fluctuations.

Additionally, in recent years, the debate has been reinvigorated by the focus and interest on Universal Basic Income (UBI). Proponents argue that UBI can reduce inequality and administrative burdens, while opponents raise the specter of catastrophic fiscal costs and deadweight loss driven by the income effect. Therefore, the income effect, often measured as the Marginal Propensity to Earn (MPE), acts as a crucial sufficient statistic for evaluating the welfare trade-offs of UBI and other large-scale redistribution systems.

Resolving this debate demands robust empirical estimates of the income effect on labor supply. However, identifying this effect cleanly is notoriously difficult. A naive regression of labor earnings on unearned income is severely biased by selection and omitted variables, individuals with low earning potential or a strong preference for leisure may be more likely to seek out transfers, while unobserved factors like can influence both income streams.

To circumvent this endogeneity, the early empirical literature has followed two primary tracks. The first was the development of structural models that impose theoretical assumptions to identify parameters (MaCurdy, 1981), which used a life-cycle framework to estimate the intertemporal substitution, the Frisch elasticity. The second involved large-scale social experiments, most notably the Negative Income Tax (NIT) experiments (Robins, 1985), conducted in the United States during the 1970s. The more recent literature leverages quasi-experimental methods to exploit natural experiments that provide plausibly exogenous shocks to unearned income. The most prominent strand of this literature studies lottery winners. Beginning with the seminal work of (Imbens et al., 2001), these studies compare the outcomes of individuals who win large prizes to those who win smaller ones or do not win at all, offering a compelling source of exogenous variation.

Although these quasi-experimental studies have provided invaluable insights, several open questions remain, suggesting a need for further research to inform policy. A significant challenge for many studies is the presence of adjustment costs. UBI pilot programs, as discussed in (Hoynes and Rothstein, 2019), for instance, are often explicitly temporary, and studies of existing programs (Jones and Marinescu, 2022) like the Alaska Permanent Fund analyze transfers that are relatively small. Similarly, lottery studies often rely heavily on the statistical power of a large number of relatively small prizes (Cesarini et al., 2017) (Picchio et al., 2018). If individuals face fixed costs to changing their labor, they may not respond at all to small or temporary income shocks, even if they would respond to a large, permanent one. A more critical limitation, particularly for lottery and inheritance studies, is that they measure the effect of a one-time wealth shock, not a permanent income flow. To map the observed behavioral change to the MPE from a permanent income stream, researchers must annuitize this wealth shock by assuming an arbitrary discount or interest rate. This methodological choice is far from trivial, as analysis in (Golosov et al., 2024) demonstrates, MPE estimates can change dramatically up to a factor of 4 depending on whether one assumes an annuitization rate of 2.5% versus 10% for example. Theoretically, this conversion from a wealth effect to an income effect depends on the ratio between the marginal propensity to consume out of wealth MPC^w and the marginal propensity to consume out of a perpetual income flow MPC^y, a ratio that might not be easy to estimate with tight bounds. Finally, the specific context of the natural experiment can present challenges. The MPE is theoretically related to the curvature of the utility function, directly proportional to the risk aversion in (Chetty, 2006) when the utility function is separable in consumption and leisure. If, for example, individuals who play the lottery have systematically lower risk preferences, the resulting income effect might be systematically lower than the general population. Similarly, studies of widows' pensions analyze a population facing a potentially large drop in income, which might lead to high curvature of the utility function and therefore exaggerated income effects (Giupponi, 2024) due to adjustment frictions in consumption.

This paper provides new estimates of the income effect by analyzing a unique institutional setting in Brazil. We exploit a public pension program for adult daughters of military personnel. Daughters of military personnel are entitled to receive their parent's pension for their remaining lifetime after the parent's death. This transfer is truly unconditional: it is not means-tested and does not depend on the daughter's labor force status, marital status, or other sources of income. This setting provides a near ideal natural experiment because the transfer is permanent, unconditional, and, most distinctively, exceptionally large. In our analysis sample, the median monthly pension is 5522 BRL, which is 87% higher than the sample's own median monthly earnings and 282% the national median earnings. This allows us to estimate the MPE from a permanent income flow rather than a wealth shock, providing a clean estimate that is robust to assumptions about annuitization rates. To conduct this analysis, we link administrative pension records to employer-employee linked data, matching individuals by full name and partially masked tax ID. Our identification strategy employs an event-study design centered on the pension start date, using a sample of military daughters who were active in the formal sector before the pension payments began. My sample has age, educational attainment, and formal employment similar to those of developed countries.

I find a precisely estimated Marginal Propensity to Earn (MPE) of -0.08 after seven years, similar to (Imbens et al., 2001), lower than (Picchio et al., 2018) and (Cesarini et al., 2017), and substantially lower than (Golosov et al., 2024). This indicates that for every dollar of permanent, unconditional pension income received, formal labor earnings decline by only 8 cents in $t+7$. This adjustment in labor supply is gradual and driven entirely by the extensive margin (exiting formal employment), with no effect on the intensive margin (hours worked) for those who remain employed. The dynamics of this adjustment are driven by two forces: a sharp, but temporary, spike in job separations immediately after the pension starts, combined with a persistent, long-run reduction in the job-finding rate. My results can be reconciled with (Golosov et al., 2024), (Cesarini et al., 2017), and (Picchio et al., 2018) if we consider a 10% annuitization rate. The lottery in (Imbens et al., 2001) is unique, as at that time a lump sum option was not available and individuals had to take the 20 year annuity.

This modest average MPE masks significant heterogeneity that provides crucial insight into the underlying economic mechanisms. First, and unlike in (Cesarini et al., 2017), the MPE is highly non-linear by transfer size: the response is near zero for small pensions but becomes significantly more negative as the pension size increases. This is a core prediction of models with adjustment costs, where individuals only overcome the fixed costs of quitting once the transfer is large enough. Second, MPE is significantly stronger for older daughters. This contradicts standard dynamic life-cycle models (where the lifetime wealth shock is largest for the young) and instead provides strong evidence of binding liquidity constraints, younger daughters cannot borrow against their future pension and must continue working.

I also examine other outcomes like education, fertility, and entrepreneurship. I find no effects on female fertility, which is consistent with (Bulman et al., 2022), although I cannot rule out the small effects found in (Yonzan et al., 2024). Given that studies on child benefits (Riphahn and Wiynck, 2017; Malkova, 2018; González, 2013) do find fertility increases, my result suggests those findings are driven by a substitution effect, not the income effect I test. I find no effect on entrepreneurship, in contrast to (Huang et al., 2023; Bellon et al., 2021;

Bermejo et al., 2025). This null finding is likely explained by two issues: my data omits solo entrepreneurs, and the monthly income stream in my setting provides less liquidity and collateral than a large lottery prize. Lastly, using a different identification strategy by leveraging the pension's equal-split rule to generate exogenous variation in payments, and comparing daughters with brothers to those with sisters in Census data, I find no effect of higher future pensions on educational attainment, a result that aligns with the wealth-effect findings of (Cesarini et al., 2016) but contrasts with the liquidity-driven results of (Akee et al., 2010).

Finally, we formalize these results within a sufficient statistics welfare framework. In this model, the redistributive gains from a UBI program depend on the initial level of earnings inequality and the ratio between the ratio of income effects and substitution effects. The efficiency cost, or deadweight loss, is determined by the marginal tax rate and the magnitude of the income effect and the substitution effect. Using our central MPE estimate of -0.08 , our calibrations show that an increase in redistribution in Brazil generates a net welfare gain as long as the compensated labor supply elasticity is less than 0.46. The results from my model are in contrast to (Daruich and Fernandez, 2024), where a fall in parental investments in child skills leads to large negative welfare consequences, however I can rule out empirically moderate negative effects in early life educational investments.

The remainder of this paper is organized as follows. Section 2 details the institutional history and rules of the Brazilian military pension system. Section 3 outlines the theoretical frameworks used to interpret the empirical results. Section 4 describes the administrative data sources and the sample construction. Section 5 presents the main empirical findings on labor supply, dynamics, and heterogeneity, as well as the results for entrepreneurship, education, and fertility. Section 6 develops the welfare analysis and calibrates the gains and losses of Universal Basic Income. Section 7 concludes.

2 Pension system

In Brazil, as in most places, widows of military personnel receive a pension upon the death of the serviceman. Uniquely in Brazil, adult daughters of military personnel inherit a share of their military parent's pension in perpetuity, regardless of their labor-market status or marital status. Payments begin after the death of the military member; however, the timing and fraction paid to daughters depend on whether a widow is alive at that time and on the blood relationship between the widow and each daughter, and the number of daughters. While the widow is alive, a portion of the pension is reserved to her, and only daughters from other marriages receive a share; when the widow dies, all living daughters become eligible to split the benefit. The benefit level is tied to the current pay of a service member of the deceased parent's rank, so it moves mechanically with the military wage grid rather than being fixed at the nominal wage at the time of death. The total amount allocated to dependents is then divided among eligible daughters (and eligible sons below 21 or disabled), with step-changes over time as siblings age out, die, or as the widow passes away. In some special cases, other family members may receive benefits as defined in the statute; these cases are rare and historically account for less than 5% of all payments.

The closest comparable benefit around the world is found in Turkey: under the national survivor pension (Ölüm ayligi), an unmarried, widowed, or divorced daughter of any age may receive a benefit only while she is not in covered employment and not receiving another pension in her own right, and payments cease upon (re)marriage, at which point a one-time marriage grant is paid. Hence, unlike Brazil's lifetime individual entitlement, which is unconditional with respect to the daughter's marital and labor market status and is tied to the current pay of the deceased parent's rank, Turkey's rule is explicitly conditional on dependency status and functions as a benefit that can be suspended as the daughter's circumstances change.[1]

[1] Social Security around the World, accessed August 26, 2025.

According to (de Oliveira, 2020), the origins of the pensions of the daughters date back to 1795, when unmarried daughters of naval officers of the Portuguese Empire would receive a pension if the widow was already deceased when the military father died. The policy logic in the eighteenth century—when women's economic opportunities were limited—was arguably to insure women's subsistence after the death of the male breadwinner. However, what is striking is that the scope expanded through the nineteenth and twentieth centuries in Brazil even as women's labor-force participation increased and as alternative social protections developed. A watershed in modern legislation is the 2000 reform, which restricted the new eligibility for daughter's pensions to cases in which the service member joined the armed forces before December 28, 2000; daughters linked to personnel who enlisted after that date are not eligible under the old rules. For this paper, the central legal framework is Law 3765[2], enacted in 1960, which regulates the payment of pensions to widows and daughters of military personnel who joined before the 2000 cut-off and codifies how the pension is split across dependents. A detailed treatment of the legislative history, the pension formulas and the military wage structure is provided by (de Oliveira, 2020).

The daughter receives the pension regardless of her marital or employment status, and the amount she receives depends on (i) whether a widow is alive, (ii) whether that widow is also her mother (the "blood relationship" in the statute), and (iii) the number of eligible dependents at each point in time. If the widow is already dead, all living daughters share the pension equally, each receiving the current salary of the deceased parent's rank divided by the number of eligible daughters and sons (sons are eligible only until age 21 or for life if disabled). While the widow is alive, only daughters from other marriages receive any payment, and each such daughter receives half of what she would receive if the widow were deceased, that is, half of the rank salary divided by the number of eligible children. As a result, three common events mechanically reallocate payments to daughters over time: (a) a brother turning 21 (which reduces the number of eligible children and raises each remaining child's share), (b)

[2]Law 3765, original text, accessed on 04/25/2024

the death of a sister (which similarly increases the shares of surviving eligible daughters), and (c) the death of the widow, which triggers a larger, discrete jump whereby all living daughters—including those who previously received nothing while the widow lived—become eligible to divide the full children's portion among themselves. Formally, the rules can be summarized as follows:

$$P_{i,t} = a_t \left(\frac{1 - b_i}{2} \frac{W_t}{N_t} \right) + (1 - a_t) \frac{W_t}{N_t}$$

where $P_{i,t}$ is the pension received by daughter i at time t, a_t is an indicator equal to 1 if the widow is alive, b_i is an indicator equal to 1 if the widow is also the mother of the daughter i' (so that daughters of the widow do not receive a share while the widow is alive), W_t is the current salary of a service member of the deceased parent's rank and qualification (reflecting contemporaneous military pay scales), and N_t is the number of eligible sons and living daughters at time t. When $a_t = 1$ (the widow is alive), each daughter receives $\frac{W_t}{2 N_t}$ (if $b_i = 0$, that is, the widow is not their mother) or zero (if $b_i = 1$, the widow is their mother). When $a_t = 0$ (the widow is dead), each daughter receives $\frac{W_t}{N_t}$.

Direct survey evidence on beneficiaries' knowledge is scarce, but all three military branches publish step-by-step guides and FAQs to apply to and maintain these pensions, and the topic is widely discussed in Brazilian media. Given the scale of the program—approximately 160,000 daughters on the payroll—and the stakes involved, average payments reportedly above mean family income in Brazil, it is reasonable to expect that current and prospective beneficiaries are generally well informed.

3 Theoretical Framework

This section outlines the theoretical frameworks used to interpret changes in earnings following the receipt of a daughter's survivor pension based on different assumptions about information and credit constraints, and also on the advantages and shortcomings of this paper compared to the previous literature studying income effects on labor supply.

3.1 Static Labor Supply

A static labor supply model is the appropriate framework if two key conditions hold: (i) agents are liquidity constrained, meaning they cannot borrow against future pension income to increase current consumption, and (ii) daughters and parents do not form a single economic unit (e.g., they live in separate households and do not pool income), meaning the pension represents a new, external source of funds.

In this framework, the agent solves a one-period utility maximization problem:

$$V_t(w, y) = \underset{l_t, c_t}{\text{Max}} \quad u(c_t, l_t) \quad \text{s.t.} \quad c_t \leq w_t\, l_t \,+\, y_y$$

where u is a concave and separable[3] ($u_{cl} = 0$) utility function,w is the wage rate, c the flow of consumption, l the labor supply, y the unearned income.

Let λ_t be the shadow value of money in utility, or marginal utility of income, then the first order conditions for the maximization problem are:

$$[c_t] \quad u_c(c_t, l_t) = \lambda_t$$

$$[l_t] \quad u_l(c_t, l_t) = -\lambda_t\, w_t$$

[3] A discussion on separability is left for the welfare analysis section.

Since $u_{cc} < 0$ and $u_{ll} < 0$, a rise in unearned income y increases consumption. This, in turn, decreases the marginal utility of income, λ_t. To restore the first-order condition for labor, the marginal disutility of working $|u_l|$, must also decrease, which implies a decrease in the labor supply l_t. This is the standard income effect, which predicts a shift from work to leisure. The magnitudes of u_{cc} and u_{ll} determine the strength of this effect.

A key parameter for welfare analysis is the marginal propensity to earn (MPE), $\frac{w_t \partial l_t}{\partial y_t}$, which we can estimate by measuring the change in earnings normalized by the change in unearned income, like the start of pension payments. The marginal propensity to earn also mediates the relation between the uncompensated and compensated elasticity of labor supply:

$$\varepsilon^u = \varepsilon^c + \text{MPE}$$

3.2 Curvature of the utility function

Following (Chetty, 2006), under the separability of consumption and leisure, we can rewrite the marginal propensity to earn as a function of the curvature of utility function and the compensated labor supply elasticity:

$$\text{MPE} = \frac{w_t\, l_t}{y_t \,+\, w_t\, l_t} \times \varepsilon^c \times \frac{u_{cc}}{u_c} \quad \Leftrightarrow \quad \varepsilon^I \equiv \frac{\partial l_t}{\partial y_t} \frac{w_t l_t + y_t}{l_t} = \varepsilon^c \frac{u_{cc}}{u_c}$$

For welfare analysis, the curvature of the utility function is a key parameter. By estimating the income elasticity of labor supply, the compensated elasticity of labor supply, and the empirical distribution of consumption, it is possible to compute the gains from redistribution.

3.3 Dynamic Labor Supply

If liquidity constraints are not important, the static labor supply model is inadequate to analyze this pension policy. The proper framework should depend on how we model the uncertainty of death and the access to capital markets.

3.3.1 Complete Markets (Arrow-Debreu Securities)

If agents have access to complete markets, they can trade claims on future uncertain events. A daughter could, for example, sell her claim to the future survivor pension (which has an uncertain start date) in exchange for a smaller, certain flow of income today.

A testable Implication in this model is that the actual death of the parent and the start of the pension payments are not "news". The consumption and labor supply paths would already have been smoothed. Therefore, we should observe no change in labor supply at the time of the parent's death, beyond any potential preference shocks (e.g., from grief), which shall be unrelated to the size of pension payment.

3.3.2 Transversality condition with clairvoyance

Now, assume markets are incomplete (agents cannot trade the pension claim) but agents can borrow and save at interest rate r, and they have "clairvoyance"—they know the exact date d of their parent's death. The agent maximizes intertemporal utility subject to a single lifetime budget constraint:

$$V(w, y) = \underset{l_t}{\text{Max}} \quad \sum \beta^t u(c_t, l_t) \quad \text{s.t.,} \sum (1+r)^{-t} \Big(c_t - w l_t - y \mathbb{I}(t > d) \Big) = 0$$

Testable Implication: With no uncertainty, the agent perfectly smooths consumption over their lifetime. The start of the pension at date d is fully anticipated. Therefore, we should observe no discrete jump in consumption or labor supply at date d.

3.3.3 Transversality condition with death uncertainty

This is the most realistic dynamic case. Agents can borrow and save but face uncertainty about the parent's date of death d. The probability of parental death in the next period, p, is positive and known (and is likely to increase with the parent's age).In this model, the agent's expected lifetime wealth is always lower when the parent is alive than when the parent is deceased, because pension payments are discounted by time and mortality risk.

When the parent is alive ($t < d$), expected pension wealth is

$$W_{alive} = E[\sum_{s=t}^{\infty}(1+r)^{-(s-t)}y\mathbb{I}(s > d)]$$

. When the parent is dead ($t \geq d$), pension wealth is

$$W_{dead} = \sum_{s=t}^{\infty}(1+r)^{-(s-t)}y = y/r$$

The size of the shock in expected wealth depends not only on flow of pension but also on the parent's life expectancy, which is decreasing by the daughters age. To compute the 12 month risk premium of the parent, I use as inputs: (i) the distribution of parent-child age gap for military households from Census 2000 and Census 2010, in figure 11, (ii) the 12 month mortality risk from national actuarial tables by age in Brazil, presented in figure 13, to construct the parental life expectancy by daughter age in 16. The size of the shock as a function of the age of the daughter can be seen in figure 3.

A testable implication is that in this model the labor supply response and the drop in earnings upon the start of the pension should be: (i) Increasing in the pension transfer size (y). (ii) Decreasing in the daughter's age (since older daughters have older parents with higher p, making the wealth shock smaller).

3.4 Alternative Interpretation: Pooled Household Income

A key threat to interpreting the earnings drop as a pure income effect (MPE) is the possibility that the pension does not represent new income but rather a simple relabeling of existing, pooled household income.

3.4.1 Co-residence

If daughters and parents live together, the parent's income (and later the daughter's pension) may have been part of a single household budget all along. In this case, the parent's death does not change the total resources available to the daughter, only the legal owner. However, Census data (Figure 18) shows that the number of adult daughters per military household is low[4] and similar to that of civilian households. This suggests co-residence is not the norm, and we can largely rule out this interpretation.

3.4.2 Intervivo transfers

A more subtle threat is the presence of inter-vivos transfers (transfers from parent to daughter while alive). When parents pension income is higher than daughters earnings, he may wish to equalize consumption in his dynasty. If a parent was already transferring, say, 3,000 BRL per month to their daughter, and the daughter's new pension is 7,500 BRL, the true change in her income is only 4,500 BRL, not 7,500 BRL. Using 7,500 BRL instead of 4,500 BRL as the denominator would cause us to underestimate the true MPE, for example incorrectly estimating -0.06 when the true parameter was -0.10.

However, a parent with multiple children would likely prioritize transfers to those not eligible for the pension (sons) or those with the highest marginal utility of income (non-working daughters). Data from the pension rolls shows 70% of recipient daughters have at least one sister. Labor force participation for women in Brazil is around 50% or less (Figure

[4] Considering that female fertility in Brazil before 1980s was between 4 and 6 children.

19).

Assuming an equal probability of having a brother as the probability of having a sister, and labor supply among sisters being uncorrelated, a simple calculation suggests that only a small fraction (roughly 15%) of recipient daughters lack a brother or a non-working sister.

Therefore, since rational inter-vivos transfers would likely be directed to other children, the potential downward bias in our MPE estimates from this source is likely very small.

3.5 Income Flows and Wealth Shocks

A central question in the literature is how to map the labor supply effects of a one-time wealth shock (e.g., a lottery prize) to those of a permanent income flow (e.g., a pension). To explore this, consider an infinitely-lived agent in a standard dynamic model:

$$V(A_t, y) = \max_{c_t, l_t, A_{t+1}} \quad u(c_t, l_t) + \beta V(A_{t+1}, y)$$

$$\text{s.t.} \quad c_t + A_{t+1} = (1+r)A_t + w_t l_t + y$$

where A_t is the stock of wealth, y is the perpetual flow of unearned income, β is the discount factor, and r is the interest rate.

The first-order conditions yield the standard result for labor supply:

$$[c] \quad u_c(c,l) = \lambda \qquad [A'] \quad (1-\beta)\,(1+r)\,\frac{\partial V(A', y)}{\partial A'} = \lambda \qquad [l] \quad u_l(c,l) + \lambda\, w = 0$$

To see how the labor supply l_t responds to permanent changes in income (y) versus wealth (A), we differentiate the FOC. Assuming separability:

$$[y] \quad u_{ll}\frac{\partial l}{\partial y} + w\frac{\partial \lambda}{\partial y} = 0 \quad \Longrightarrow \quad \frac{\partial l}{\partial y} = -\frac{w}{u_{ll}}\frac{\partial \lambda}{\partial y}$$

$$[A] \quad u_{ll}\frac{\partial l}{\partial A} + w\frac{\partial \lambda}{\partial A} = 0 \quad \Longrightarrow \quad \frac{\partial l}{\partial A} = -\frac{w}{u_{ll}}\frac{\partial \lambda}{\partial A}$$

The relative labor supply response to income versus wealth is therefore equal to the relative change in the marginal utility of wealth:

$$\frac{\partial l/\partial y}{\partial l/\partial A} = \frac{\partial \lambda/\partial y}{\partial \lambda/\partial A}$$

Since $\lambda = u_c(c)$ and $u_{cc} < 0$, this ratio is equivalent to the ratio of the marginal propensities to consume (MPC) out of income and wealth:

$$\frac{\partial l/\partial y}{\partial l/\partial A} = \frac{u_{cc}(\partial c/\partial y)}{u_{cc}(\partial c/\partial A)} = \frac{\partial c/\partial y}{\partial c/\partial A}$$

This result is intuitive. Both income and wealth affect labor supply through consumption (by changing λ). Therefore, the relative effect of income vs. wealth on labor supply must equal their relative effect on consumption.

To map the marginal propensity to earn (MPE) from wealth ($w\frac{\partial l}{\partial A}$) to the MPE from income ($w\frac{\partial l}{\partial y}$), one must know this MPC ratio.

$$\text{MPE(income)} = \text{MPE(wealth)} \times \left(\frac{\partial c/\partial y}{\partial c/\partial A}\right)$$

In a simple permanent income model, agents consume all of a permanent income flow ($\partial c/\partial y = 1$) but only consume the "interest" from a wealth stock ($\partial c/\partial A = r$). In this case, the ratio is $1/r$. Studies using wealth shocks (like lotteries) to estimate the MPE must therefore "annuitize" the wealth shock by an assumed rate r.

This approach is highly sensitive to the choice of r. For example, the difference between assuming an annuitization rate of 4% versus 6%—a plausible range—will change the final MPE(income) estimate by 50% (since 1/0.04 is 50% larger than 1/0.06).

Furthermore, if agents are liquidity-constrained or behave myopically, their MPC out of a wealth shock ($\partial c/\partial A$) might be much larger than r. If agents consume the lottery prize "too fast," the true ratio $(\partial c/\partial y)/(\partial c/\partial A)$ is much smaller than $1/r$. Using the simple $1/r$ annuitization would then cause the researcher to severely overestimate the true MPE(income).

Given that the mapping from wealth shocks to income flows is fraught with uncertainty and sensitive to assumptions about r and consumption behavior, it is preferable to estimate the MPE(income) directly. Since this paper estimates the MPE from the start of a pension, an income flow, our estimates are directly useful for the welfare analysis of social programs that provide income flows, without requiring assumptions about annuitization rates.

3.6 Adjustment costs

A key challenge in estimating the income effect is that labor supply adjustments are not frictionless. In a classical model without adjustment costs, any change in unearned income, no matter how small, should theoretically trigger a corresponding (though perhaps small) change in labor supply.

However, in reality, agents face significant adjustment costs when changing their labor supply. These costs can be monetary (e.g., search costs for a new job with fewer hours) or non-monetary (e.g., the inertia of a stable career, or the fixed costs of working like childcare and commuting). If these adjustment costs are present, small income shocks will not be sufficient to trigger an observable change in labor supply.

Therefore, to identify the MPE, the income shock must: (i) The shock must be substantial enough to make the gain in leisure valuable enough to overcome the fixed costs of adjustment, and (ii) The shock must be permanent (or at least long-lasting) to justify a permanent change in labor supply. Agents are unlikely to incur the fixed costs of quitting a job in response to a temporary windfall, which they would instead smooth via consumption.

Unlike UBI pilot programs (Hoynes and Rothstein, 2019), a permanent shock, such as the pension studied in this paper, fundamentally alters the agent's lifetime budget constraint and justifies a permanent labor supply response.

4 Data

The project links three core administrative data sets: (i) Ministry of Defense pension records (2023-2025) identifying widows/daughters, payment start dates and amounts, parent rank, and masked CPFs; (ii) RAIS employer-employee data (2002-2016) covering all formal contracts with detailed information on wages, hours, occupation, education, race, age, and maternity leave spells; and (iii) CNPJ corporate registry listing owners/partners of multi-owner firms with masked CPFs and entry dates (2024). Together, these sources provide comprehensive data on pension payments, formal labor market outcomes, and firm ownership needed to study earnings, labor force participation, fertility, and entrepreneurship. In addition, I utilize the Brazilian Decennial Census for the years 2000 and 2010 to examine early life outcomes, such as private school attendance, college enrollment, and employment status, providing valuable context on the beneficiaries' backgrounds. Individuals are linked across datasets using a partially masked tax ID (CPF6) and fuzzy name matching. The final sample includes 30,400 daughters who started pensions after Jan 1, 2002, and were formally employed on Jan 1, 2002. This baseline condition allows observation of age (missing in pension data) and minimizes selection bias. The sample consists of older adults (median age 48) who are positively selected on education and earnings. The analysis covers 2002-2016.

4.1 Data sources

The first dataset is the public roster of *military survivor pensions* paid to widows, daughters, and a small set of other eligible dependents. For January 2020–December 2022 it reports, at the payment record level, the paid amount, recipient's full name, a partially masked individual tax identifier (CPF), the benefit start date, the deceased parent's military rank, and the recipient type (e.g., daughter, widow). In most records, the parent's name and the masked identifier are also present, allowing beneficiaries to be grouped by family ties. These files are released under Brazil's transparency program and can be downloaded in monthly spreadsheets from a government website *Portal da Transparência*.[5] The dataset does not contain information on pensions that stopped being paid before January 2020, due to the aging of sons or the death of a daughter or widow, and does not contain information on the age of the recipient.

The military survivor pensions in Brazil are generous in both value and scope. In 2020, gross expenditures were about BRL 30 billion (current prices), roughly one quarter of contemporaneous defense expenditures, and close to 1% of federal spending. A comparison with other social insurance programs in Brazil is presented in 9 and 10. The administrative roster identifies 258,103 distinct beneficiaries who received at least one payment between January 2023 and July 2025; daughters account for 159,077 (about 60%), and widows comprise most of the remaining. The mean monthly benefit is BRL 8,808 (median BRL 6,347). For comparison, the average wages in Brazil are BRL 2,483 in 2010 and the average wage for college graduates is BRL 6,575 in 2010. Among surviving daughters specifically, the mean monthly pension is BRL 7,225 (291% of the national average wage). Because daughters receive the benefit for life, monthly amounts of this magnitude imply large lifetime transfers. Figure 24 shows the overall distribution of monthly payments.

Two observables that govern the monthly amount received by a daughter are: (i) the number of eligible beneficiaries tied to the same service member (the statutory split across dependents) and (ii) the deceased parent's military rank at death, which pins down the relevant pay grid. Residual dispersion likely reflects within-rank pay heterogeneity (e.g., certifications, tenure, type of operational unit, or location). Using the parent's rank (available for about 85% of records)[6], the composition of daughters by parental rank is: lower enlisted (private/corporal) 990 (5.5%), non-commissioned officers (sergeants) 5,264 (29.5%), junior officers (lieutenant/captain) 7,778 (43.6%), senior officers (major/colonel) 2,202 (12.3%), and flag officers (general/admiral/air marshal) 1,607 (9.0%). Using the reported parent name,

[5] Portal da Transparência — *Servidores e Pensionistas*, downloads.
[6] Rank coverage is approximately 85% of the sample.

I reconstruct family ties: 39% of beneficiaries are sole recipients, and 24% share with exactly one other dependent. Among daughters, 21% are the sole recipients and 27% share with exactly one other, implying that most daughters split the benefit with two or more relatives. Although family size varies somewhat with parental rank, the rank–family-size gradient appears second-order; see Figure 22.

The second dataset is an employer-employee administrative data set for Brazil's *formal* labor market (RAIS), available annually from 1984 onward[7]. RAIS lists every active formal job contract in the reference year and includes worker characteristics (education, race, sex, age, name, unique tax code identifier), establishment characteristics (industry codes, public or private status, municipality and state, firm size, unique tax code identifier) and contract information (occupation codes, admission and termination dates, contracted weekly hours, and earnings components). RAIS also records leave episodes, including *maternity leave*, which in Brazil typically lasts 4 to 6 months in Brazil (Machado et al., 2024), enabling a fertility proxy. Microdata and documentation are provided by the Ministry of Labor. The dataset does not observe informal or self-employment. Some recent papers that also use the same dataset, the Brazilian Administrative Labor Market Data, are (Carvalho et al., 2018), (Van Doornik et al., 2022), and (Machado et al., 2024).

The third dataset is Brazil's national corporate registry (*CNPJ*), maintained by the Internal Revenue Service (RFB). It enumerates all legal entities and, for those that report the owner roster (*Quadro de Sócios e Administradores*, QSA), typically multi-owner legal forms, lists each owner (natural person or legal entity) and their unique tax identification. The information available in the record is a unique tax identification code of the entity (CNPJ), the type of organization (government, nonprofit, for-profit, etc.), the start date, the end date (if applicable), the name of the organization, the full name of the owner, the date the owner joined the organization, and the tax code of the owner. The ownership share of each

[7] Although employer-employee data have been available since 1984, the full name of each individual has only been provided by the Ministry of Labor since 2002. Since the full name is required for the matching between datasets, this limits the usefulness of data prior to 2001.

individual or organization and the list of owners who have left a firm's ownership list are not available in this dataset. This dataset can be downloaded by anyone from the Brazilian IRS, and its publication also came as a consequence of the push for more transparency in the Brazilian government. Since this dataset is very new, only a few studies have used it, for example (Dias and Rocha, 2021) and (Colonelli et al., 2024).

The fourth dataset used is the Brazilian Decennial Census for 2000 and 2010. The Decennial Census collects basic information on all Brazilian residents every 10 years, with a detailed questionnaire applied to 11% of households. The microdata for the decennial census can be downloaded at Puc-Rio Data Zoom with a dictionary and data cleaning available at Data Zoom. The census lists occupation, education, labor market participation, household characteristics, and family ties for all individuals in the detailed questionary, approximately 20 million individuals for each year.

4.2 Matching

I link the military survivor pensions dataset and the national corporate registry dataset with the employer-employee linked dataset by combining a *stable numeric token* with *human–readable name*. Public data provides only digits 4–9 of the CPF (henceforth *CPF6*), so I first require an exact match on {CPF6, first given name} to form a small candidate set. Then I compare the full names and accept a link only when the names are essentially the same, up to minor typos and abbreviations or marriage-consistent changes common in Brazil (paternal surname preserved, husband surname added at the end). This conservative approach leverages Brazilian naming conventions to recognize the same person across files while avoiding forced links; full algorithmic details and validation appear in Appendix A.

4.3 Final Sample

The final analysis sample consists of daughters of military personnel who (i) were receiving a pension between January 2023 and July 2025, (ii) began receiving payments after January 1, 2002, and (iii) were formally employed on January 1, 2002. Table 1 details the progressive application of these filters, resulting in a final sample of 30,400 individuals. This final sampling step is restrictive. This sharp drop is not due to incomplete matching but rather reflects the historically low female labor force participation in Brazil (see Figure 19).

The final restriction is essential for the identification strategy. Age is a critical control variable for life-cycle earnings, but it is not available in the pension administrative data. However, age is recorded in the employer-employee labor market data. By conditioning the sample on individuals employed on the fixed baseline date of January 1, 2002, I can observe a baseline age for every individual. This strategy avoids selection bias induced by the pension itself, as Table 1 shows the median pension start year for this sample is 2015, well after the baseline. This paper's estimations use the period from 2002 to 2016, excluding daughters who were "always treated" (i.e., those whose payments began before 2002).

Table 2 and Figure 23 describe the characteristics of this final sample. The daughters are primarily older adults, with ages in their 40s to 60s (median age 48). This is an expected consequence of sampling individuals who have already experienced parental loss. To validate this age profile, I simulate the age at parental loss using actuarial tables and the empirical age gap between parents and children (Figure 15). This simulation predicts a median age of 51 at parental death, which closely matches the median age in my final sample.

This sample is also positively selected. As shown in Table 2, daughters in the sample are more educated (66% with college) and have higher earnings than the average Brazilian woman. This is consistent with the fact that military careers in Brazil are well-compensated, leading to positive selection among parents, which in turn positively influences the human capital formation of their children.

5 Results

In this section, I discuss the identification strategy, the effects of pension payments on labor supply, earnings, fertility, and entrepreneurship. A quick summary is that labor force participation declines slightly after payments start, the marginal propensity to earn is very low and increasing with payment size and age of recipient, separations raise sharply and return to baseline but reemployment remains lower for a longer window, marginal propensity to earn is larger for individuals with higher wages in 01/01/2002, but similar across individuals with and without a college degree. Additionally, neither business ownership, fertility, or educational attainment are affected by pension payments, ruling out moderate size effects.

5.1 Identification

I use an event study design to identify the effect of an unconditional permanent income stream on formal labor market outcomes (formal employment, weekly hours, monthly earnings, probability of being hired and probability of being fired each month), fertility [8], and entrepreneurship.

For outcome y_{it} (for example earnings), define the relative time $k = t - T_i$ and estimate

$$y_{it} = \sum_{k \neq -1} \beta_k \, \mathbf{1}[t - T_i = k] \; + \; \alpha_i \; + \; \alpha_t \; + \; \alpha_a \; + \; \varepsilon_{it}, \tag{1}$$

The coefficient β_k is the average effect k years from onset relative to the omitted year $k = -1$. Individual fixed effects α_i absorb time-invariant differences, calendar-year effects α_t capture aggregate shocks, α_a is a fixed effect for age, and errors are clustered at the individual level. All age fixed effects are not identified, but relative age fixed effects, setting an age dummy as baseline, are identified. The goal of including age fixed effects is to remove non linear effects of aging in the event study, since individuals age as relative time increases

[8] Maternity leave is taken as proxy for fertility.

in an event study.

To scale effects by the size of the transfer, aiming to recover the marginal propensity to earn when earnings is the outcome of interest, I interact the event dummies with the monthly pension amount assigned to the daughter in the registry:

$$y_{it} = \sum_{k \neq -1} \theta_k \left(\mathbf{1}[t - T_i = k] \times \text{Payment}_i \right) \; + \; \alpha_i \; + \; \alpha_t \; + \; \alpha_a \; + \; \varepsilon_{it}. \tag{2}$$

The coefficients θ_k trace the *marginal propensity to earn* (MPE) at each event time.

5.2 Participation, Hours, and Earnings

Figures 26, 27, and 28 present some of the main event study results, illustrating the dynamic impact of pension receipt on formal employment, weekly hours, and monthly earnings, respectively. A consistent pattern emerges across all three outcomes.

First, the figures provide strong evidence for the validity of the research design. In the years preceding the pension, from t-5 to t-1, the point estimates are stable and statistically indistinguishable from zero, confirming the absence of pre-existing trends between the treatment and control groups. However, a small "anticipation" effect is visible in the 12 months immediately before the pension start date (t-1). This is likely attributable to administrative delays as seen in figure 21, since the pension is triggered by a parent's death, some individuals may experience the death and begin adjusting their labor supply before the first payment is officially recorded.

Following the start of the pension, all three outcomes exhibit a gradual and persistent decline. Figure 26 shows that the formal employment rate falls steadily, stabilizing approximately 6 percentage points below the baseline four years after the event, a substantial effect relative to the sample mean of 63%. Similarly, Figure 27 shows that weekly hours, which include zeros for non-employed individuals, decline by approximately 2.5 hours after five years, from a mean of 22.9 hours. This reduction is driven almost entirely by exits from employment (the extensive margin) rather than reductions in hours for those who remain employed.

This decline in employment and hours translates directly into lower earnings. As shown in Figure 28, monthly formal labor earnings gradually decrease by approximately $180 (USD PPP) after seven years, relative to a sample mean of $2,250. While the magnitudes of these effects on labor supply are modest, the narrow confidence intervals in all three figures indicate that they are precisely estimated.

5.3 Admissions and Separations

The overall decline in formal employment is driven by changes in both the inflow and outflow rates. Figures 29 and 30 decompose this extensive margin adjustment by plotting the effect on monthly hiring rates (inflows) and involuntary job separations (outflows), respectively.

Figure 30 shows a sharp, immediate, and short-lived spike in involuntary job separations (firings) precisely when the pension payments begin. This increase in the outflow rate is large, representing a temporary 30% jump over the baseline firing rate of 2.2%. However, this effect dissipates quickly, returning to the baseline by t+3.

In contrast, the effect on inflows from non-employment is more persistent. Figure 29 shows that the hiring rate, which has a baseline mean of 2.1% per month, drops immediately after the pension starts. This drop remains statistically and economically significant for at least six years, representing a persistent 10% to 20% reduction in the job finding probability.

Together, these results show that the labor supply adjustment is dynamic. The initial, rapid decline in employment is partially driven by a temporary spike in job separations, while the long-run, stable reduction in the employment rate is maintained by a persistent decrease in the job-finding rate.

5.4 Marginal propensity to earn

The main finding of this paper is the Marginal Propensity to Earn (MPE), estimated by interacting the event study dummies with the size of the pension payment. Figure 4 plots the dynamic coefficients of this estimation.

The plot confirms the validity of the identification strategy, as the coefficients in the pre-treatment period are stable and statistically indistinguishable from zero. Following the pension start at $t = 0$, there is a gradual labor supply adjustment, with earnings declining steadily over time. Seven years after the pension begins, the MPE stabilizes at approximately -0.08. This indicates that for every dollar of permanent, unearned pension income, daughters reduce their formal labor earnings by 8 cents. The confidence intervals remain narrow even seven years out, demonstrating that this small effect is precisely estimated.

Figures 1, 25a, and 25b places this result in the context of the existing literature. The MPE of -0.08 is small compared to estimates from some lottery studies. However, this paper's estimate is derived from a direct income flow, unlike lottery studies[9] that must annuitize a wealth shock—a calculation that is highly sensitive to the chosen interest rate. As Figure 1 shows, my estimate is substantially lower than lottery-based MPEs that use low annuitization rates (e.g., 2.5%). Yet, the results can be reconciled: my finding is very similar to the MPEs from those same lottery studies when a higher, 10% annuitization rate is applied. This suggests that the MPE is small and that differences across studies may be driven by methodological choices rather than underlying economic behavior.

However, my small MPE is hard to reconcile with (Giupponi, 2024), (Böheim and Topf, 2020), and (Rocha, 2021). A possible reason is that adjustment costs in consumption can make agents too sensible for income losses, which would imply a high marginal propensity to earn, as discussed previously, more details about the link between the marginal propensity to earn and the curvature of the utility function are in (Chetty, 2006).

[9] The exception is (Imbens et al., 2001)

5.5 Heterogeneity

The modest average MPE masks significant heterogeneity across different subgroups. I explore how the labor supply response varies by baseline education, baseline wage, age at first payment, and the size of the pension itself, and how we can rule out the dynamic labor supply model as preferred framework in favor of the static labor supply model in this setting.

The response to the pension is similar across education levels but differs starkly by baseline income. As shown in Figure 31, there is no statistically significant difference in the MPE between daughters with a college degree and those without. The confidence intervals for the two groups largely overlap, suggesting education at baseline is not a key margin of heterogeneity. In contrast, Figure 32 shows a clear divergence based on baseline earnings. Both groups (above and below the median wage of $958 USD PPP) reduce their earnings, but the effect is substantially stronger and more persistent for those with above-median baseline earnings. A possibility is that higher-income individuals, who may have more flexibility or be closer to their target income, have a larger labor supply elasticity with respect to income.

A more complex pattern emerges when analyzing the MPE by the daughter's age at the start of the pension. Figure 5 plots the MPE as a flexible function of age, revealing that the labor supply response is modest and close to zero for the youngest daughters in the sample (e.g., in their 20s and 30s). However, the MPE becomes increasingly negative as age increases, reaching nearly -0.10 for daughters in their 70s.

This finding presents an apparent puzzle. Standard dynamic models would predict the opposite, as the lifetime wealth shock from the pension is largest for the youngest daughters (who have the longest time horizon to receive payments). This puzzle is resolved by considering the role of liquidity constraints. A simple static labor supply model, in which agents cannot borrow against future income, predicts this exact pattern. Younger daughters, while ”wealthier” on paper, cannot access their future pension wealth to finance current consumption and are thus constrained to continue working. Older daughters, conversely, receive the income immediately and can adjust their labor supply at a point where they are already closer to retirement, making the income effect more immediately salient.

Finally, Figure 6 demonstrates that the MPE is non-linear with respect to the size of the pension payment. The MPE is close to zero for very small pension amounts. As the monthly payment increases, the MPE becomes increasingly negative, eventually flattening out at approximately -0.05 for large pension amounts (e.g., $> \$10,000$ USD PPP).

This non-linear relationship is highly consistent with labor supply models that feature adjustment costs. In such models, individuals will not change their labor supply (e.g., quit a job or reduce hours) in response to a small income shock because the utility gain does not outweigh the fixed costs of adjustment (such as search costs, loss of firm-specific human capital, or disrupting a career). Only when the income shock is sufficiently large to overcome this adjustment cost "hurdle" do we observe a behavioral response. The data clearly supports this, showing a negligible response at low pension levels and a significant, negative MPE only after the payment size surpasses a certain threshold.

5.6 Other margins of adjustment

To validate a welfare analysis which focus solely consumption and labor supply, I investigates other potential margins of adjustment to the unearned income shock: fertility, entrepreneurship, and education. The analysis finds no statistically significant effects on these margins, supporting the modeling choices used in the subsequent welfare analysis.

5.6.1 Fertility

The effect of unearned income on fertility is theoretically ambiguous. An income effect may increase the demand for children (assuming they are a normal good), while a countervailing effect may arise if increased overall consumption raises the marginal value of leisure, potentially crowding out time-intensive child-rearing.

To measure fertility, I use paid maternity leave spells as a proxy, drawing on data from the Brazilian linked employer-employee database available since 2007. In the formal sector, pregnant women are entitled to 3-4 months of paid leave with job protection, providing a reliable measure of births to formal workers (for a detailed description, see (Machado et al., 2024)). As shown in Table 3, the analysis finds no statistically significant effect of pension payments on fertility. However, the estimates are imprecise, making it difficult to rule out moderate size effects.

5.6.2 Entrepreneurship

Unearned income could also plausibly affect entrepreneurship through three channels. It could (i) relax liquidity constraints, providing capital to start a business; (ii) increase the marginal value of leisure, discouraging the demanding work of a new venture; or (iii) reduce risk aversion (assuming DARA[10] utility), thus increasing the appeal of risky business activities.

To measure this outcome, I match the pension database with the Brazilian National Ownership Registry using individual tax IDs. I construct a balanced panel ($N = 120,899$ and $T = 25$) of daughters with pensions starting after 2000, covering the period 2000-2024. The primary outcomes are new business creation (mean 1.70 per 1,000 daughter-years) and joining an existing business as a partner (mean 4.93 per 1,000 daughter-years). The identification strategy mirrors that of the main labor supply analysis.

The results, presented in Table 4 and the event study plots in Figures 33a and 33b, show no significant impact on either measure of entrepreneurship. The point estimates are near zero and precise enough to rule out economically meaningful changes in business creation or ownership, such as a 16% increase in business creation from baseline or a 3.4% increase in business ownership from baseline.

[10] Decreasing absolute risk aversion, for example $u(c) = Ln(c)$ or any other CRRA utility function.

5.6.3 Education

To evaluate the impact of future pension income on educational attainment, I use data from the 2000 and 2010 decennial censuses. The analysis focuses on a specific sample: families with exactly two children, one of whom is a daughter, and with a father or mother in the military forces.

This sample allows for a clean research design. I classify daughters with one brother as the treatment group and daughters with one sister as the control group. This identification strategy takes advantage of a specific rule in the calculation of military survivor pensions, where the pension payment is shared among sisters but not brothers.

I estimate the impact of this additional future income using the following regression model:

$$y_i = \sum_a \delta_a \, \mathbf{1}\{\text{age}_i = a\} \; + \; \beta \;\; \text{Brother}_i + u_i$$

In this equation, the dependent variable y_i is an outcome of interest for the daughter, specifically a dummy variable for attending private school (vs. public school) or college enrollment. The model also controls for a full set of age fixed effects, δ_a.

The hypothesis is that future pension payments can reduce the education premium $L\left(W^c - W^{HS}\right)$. This occurs not by changing the college wage premium $W^c - W^{HS}$, but by changing the equilibrium hours worked over a lifetime dL/dy.

However, as shown in Table 5, the results show no detectable change in educational investments. This null finding is consistent with the theory, given that labor supply responses to income are known to be very small, the resulting change in the education premium would also be very small. A significant educational response would only be likely if the elasticity of college enrollment to the earnings premium was very high.

6 Welfare Analysis

Universal basic income (UBI) can alter the tax base by changing market hours by two channels, an income effect from higher transfers, and a substitution effect from higher taxes. This section presents a simple static framework in which proportional taxes and a lump-sum transfer affect the tax base via labor supply choices and derive sufficient statistics that map reduced-form elasticities into the welfare effect of UBI. Using the link between the curvature of consumption and the ratio of income to substitution effects, I calibrate the welfare effect of basic income for different values of wage elasticity and total redistribution.

6.1 Agents

A continuum of individuals $i \in I$ with mass 1 chooses consumption c_i and labor supply l_i to maximize utility. The wages w_i are heterogeneous and the earnings are taxed at a rate $\tau \in [0, 1)$, and each individual receives a uniform transfer $b \geq 0$. Consumption is:

$$c_i \;=\; (1-\tau)\, w_i\, l_i \;+\; b.$$

Preferences are separable in consumption and labor supply,

$$u_i(c_i, l_i) \;=\; f(c_i) \;+\; h(l_i)\,, \qquad f' > 0,\; f^{"} < 0,\; h' < 0,\; h^{"} > 0.$$

Given (τ, b), and denoting by λ_i the marginal utility of income, the F.O.C. [11] :

$$u_c((1-\tau)\, w_i\; l_i(b,\tau)\; +\; b) = \lambda_i \qquad u_h(\; l_i(b,\tau)\;) = -(1-\tau)\, w_i\, \lambda_i$$

[11] Assuming a interior solution

6.2 Planner

A utilitarian planner finances a uniform basic income b with a proportional tax τ on earnings and chooses (τ, b) subject to a balanced budget to maximize social welfare. Let $W(\tau, b) \equiv \int u_i \, dF_i$ and $Y(\tau, b) \equiv \int w_i \; l_i(b, \tau) \; dF_i$. The government budget is

$$b \; = \; \tau \, Y(\tau, b).$$

An increase in b has a direct effect on utility through higher consumption and an indirect effect through the tax change $\mathrm{d}\tau/\mathrm{d}b$ required to balance the budget and the behavioral responses of l_i that alter Y. Using the envelope theorem (i) The direct effect is the marginal utility of consumption of the basic income recipient, with changes in utility due to re-optimizations of working time being a second order effect (ii) The indirect effect is the marginal utility of consumption of the tax payer multiplied by the slope of increase in the tax rate necessary to finance an increase in the basic income payment. More precisely, the direct and indirect effects are as follows:

$$\frac{\partial W}{\partial b} = \int u_c((1-\tau)\, w_i l_i + b)\, dF_i, \tag{3}$$

$$\frac{\partial W}{\partial \tau}\,\frac{d\tau}{db} = \frac{\displaystyle\int -\, w_i l_i \, u_c((1-\tau)\, w_i l_i + b)\, dF_i}{\displaystyle\int w_i l_i \, dF_i} \times \frac{1 - \tau \frac{\partial Y}{\partial b}}{1 - \frac{\tau}{1-\tau} \frac{1-\tau}{Y} \frac{\partial Y}{\partial (1-\tau)}}, \tag{4}$$

You might recognize that $\frac{\partial Y}{\partial b}$ is the marginal propensity to earn out of unearned income and $\frac{(1-\tau)}{Y} \frac{\partial Y}{\partial (1-\tau)}$ is the taxable income elasticity, which is equal to the uncompensated labor supply elasticity when wages are fixed and other earnings responses are shut down.

6.3 Sufficient statistics

It is useful to express welfare changes in money-metric units by normalizing by the earnings-weighted marginal utility of consumption. Define

$$\tilde{v} \equiv \int u_c(c_i) dF_i, \qquad \hat{v} \equiv \frac{\int u_c(c_i)\, w_i l_i \, dF_i}{\int w_i l_i \, dF_i},$$

The first-order change in money-metric welfare is

$$\frac{d\tilde{W}}{db} = \frac{1}{\hat{v}} \frac{dW}{db} = \frac{\tilde{v} - \hat{v}}{\hat{v}} + \tau \times \frac{(1-\tau)\, \frac{\partial Y}{\partial b} - \varepsilon}{1 - \tau\,(1+\varepsilon)}$$

The first term captures the redistributive gain from transferring one unit to agents with high marginal utility (relative to the earnings-weighted average $\hat{v}$). The second term collects the efficiency cost from behavioral responses: lower hours from more transfers and more taxes, where ε is the taxable income elasticity, and $\frac{\partial Y}{\partial b}$ is the marginal propensity to earn.

6.4 Separability between consumption and leisure

Although standard separable utility functions predict smooth consumption through predictable events like retirement, empirical studies consistently find a sharp drop in expenditure (Bernheim et al., 2001), which appears to reject the strong form of separability. However, this finding has been challenged. Evidence suggests that the magnitude of these drops is often small (Chetty, 2006), and, more importantly, can be partially or fully accounted for by reducing work-related expenses (WRE) such as transportation and food consumed outside of home (Battistin et al., 2009) (Aguiar and Hurst, 2005). This suggests the anomaly is a measurement issue, justifying a weaker form of separability where "true" consumption is simply reinterpreted as expenditure net of these work-related costs.

6.5 Calibration

This section calibrates the marginal welfare effects of a UBI program using the estimated MPE and the sufficient statistics framework, assuming constant elasticities locally.[12]

The marginal welfare change per dollar transferred can be decomposed into a redistribution gain and an efficiency cost (deadweight loss). As derived in Appendix B using a CARA utility approximation, the marginal redistribution gain ($\partial W/\partial b$) is approximately:

$$\frac{\partial W}{\partial b} \approx (1-\tau) \times \gamma \times [CV(w_i l_i)]^2 \tag{5}$$

where τ is the proportional tax rate financing the UBI, $CV(w_i l_i)$ is the coefficient of variation of earnings, and γ is the coefficient of absolute risk aversion. Following (Chetty, 2006), we link risk aversion to labor supply elasticities under separable utility: $\gamma = -\varepsilon^I/\varepsilon^c$, where ε^I is the income elasticity and ε^c is the compensated elasticity.

From Table 2, the average ratio of earnings to total income (pension plus earnings) in our sample is $wl/(y+wl) = 0.463$. Combined with our central MPE estimate of -0.08, this yields an income elasticity $\varepsilon^I = MPE \times (y+wl)/wl = -0.08/0.463 \approx -0.17$. Using the 2010 Brazilian Census (after 2% winsorizing), the coefficient of variation for earnings $CV(w_i l_i)$ is approximately 1.005 (Figure 34). Substituting these values, the marginal redistribution gain becomes a function of the tax rate τ and the compensated elasticity ε^c:

$$\frac{\partial W}{\partial b} \approx (1-\tau) \times \left(\frac{0.17}{\varepsilon^c}\right) \times (1.005)^2 \tag{6}$$

The marginal efficiency cost per dollar transferred, or deadweight loss (DWL), incorporates both the income effect (MPE) and the substitution effect (ε^c). Using the Slutsky

[12] See the discussion in (Kleven, 2021) on the implications of assuming constant elasticities for welfare analysis.

equation to express the uncompensated elasticity $\varepsilon^u = \varepsilon^c + MPE$, the DWL is:

$$DWL = \tau \times \frac{\varepsilon^c - 0.08\tau}{1 - \tau(1 + \varepsilon^c - 0.08)} \tag{7}$$

Figure 7 plots the redistribution gain (Equation 6) and the DWL (Equation 7) against the compensated elasticity ε^c, assuming a tax rate $\tau = 32\%$ (approximating Brazil's government revenue share of GDP). The redistribution gain (red line) is large for small ε^c (implying high risk aversion γ) and decreases as ε^c rises. The DWL (blue line) increases approximately linearly with ε^c. The marginal net welfare gain is positive ($\partial W/\partial b > DWL$) when ε^c is below approximately 0.46.

Figure 8 plots the optimal tax rate τ^* (which equals the optimal UBI size as a fraction of GDP, b^*/Y) that sets the marginal net welfare gain to zero ($\partial W/\partial b = DWL$), as a function of ε^c. Holding the income effect (MPE = -0.08) fixed, a higher compensated elasticity ε^c implies both lower risk aversion[13] (reducing the desirability of redistribution) and higher deadweight loss, thus unambiguously lowering the optimal τ^*.

This calibration exercise is intended to illustrate how the income effect (MPE) and the substitution effect (ε^c) jointly determine the trade-offs inherent in large-scale transfer programs, rather than providing a precise policy recommendation for UBI size.

7 Conclusion

This paper provides new estimates for income effects on labor supply by studying a unique Brazilian pension program that provides large, permanent, and unconditional transfers to military daughters. The central finding is a modest but precisely estimated Marginal Propensity to Earn (MPE) of -0.08 after seven years. This indicates that for every dollar of unearned income received, formal labor earnings fall by only 8 cents. This labor sup-

[13] From $\gamma\varepsilon^c + \varepsilon^I = 0$.

ply adjustment is gradual, operating on the extensive margin through a dynamic process: a temporary spike in job separations at the onset, followed by a persistent decline in the job-finding rate.

The heterogeneity analysis provides critical insights into the economic models that best explain this behavior. The MPE is found to be strongly non-linear, close to zero for small pension amounts and becoming more negative as the transfer size increases. This finding is highly consistent with labor supply models featuring adjustment costs, where individuals only alter their employment status once the income shock is large enough to overcome the fixed costs of doing so. Furthermore, the MPE is strongest for older daughters, despite the lifetime wealth shock being largest for younger recipients. This finding contradicts standard dynamic life-cycle models and instead provides strong support for a static model where liquidity constraints are binding. Younger daughters, while "wealthier" on paper, cannot borrow against their future pension and thus must continue working. The analysis of other margins, finding no significant effects on fertility, entrepreneurship, or education, validates the welfare framework's focus on labor supply and consumption.

These findings have direct and important implications for the debate on large-scale transfer programs like a Universal Basic Income (UBI). By providing a clean estimate of the MPE from an income flow rather than a wealth shock. The small MPE found in this study implies that the efficiency costs, or deadweight loss, from the behavioral response to such a transfer are small. If the compensated labor supply elasticity is small, the welfare calculation suggests that the net effect of a large-scale unconditional transfer program is positive.

References

Aguiar, M. and Hurst, E. (2005). Consumption versus expenditure. *Journal of Political Economy*, 113(5):919–948.

Akee, R. K. Q., Copeland, W. E., Keeler, G., Angold, A., and Costello, E. J. (2010). Parents' incomes and children's outcomes: A quasi-experiment using transfer payments from casino profits. *American Economic Journal: Applied Economics*, 2(1):86–115.

Battistin, E., Brugiavini, A., Rettore, E., and Weber, G. (2009). The retirement consumption puzzle: Evidence from a regression discontinuity approach. *American Economic Review*, 99(5):2209–2226.

Bellon, A., Cookson, J. A., Gilje, E. P., Heimer, R. Z., and Nieuwerburgh, S. V. (2021). Personal wealth, self-employment, and business ownership [house prices, collateral, and self-employment]. *The Review of Financial Studies, Society for Financial Studies*, 34(8):3935–3975.

Bermejo, V., Ferreira, M., Wolfenzon, D., and Zambrana, R. (2025). How do cash windfalls affect entrepreneurship? evidence from the spanish christmas lottery. *Journal of Financial and Quantitative Analysis*, 60(5):2229–2258.

Bernheim, D. B., Skinner, J., and Weinberg, S. (2001). What accounts for the variation in retirement wealth among u.s. households? *American Economic Review*, 91(4):832–857.

Bulman, G., Goodman, S., and Isen, A. (2022). The effect of financial resources on homeownership, marriage, and fertility: Evidence from state lotteries. *NBER Working Papers, National Bureau of Economic Research, Inc.*, 30743.

Böheim, R. and Topf, M. (2020). Unearned income and labor supply: Evidence from survivor pensions in austria. *IZA Discussion Papers*, 13994.

Carvalho, C. C., Corbi, R., and Narita, R. (2018). Unintended consequences of unemployment insurance: Evidence from stricter eligibility criteria in brazil. *Economics Letters, Elsevier*, 162(c):157–161.

Cesarini, D., Lindqvist, E., and Notowidigdo, M. J. (2017). The effect of wealth on individual and household labor supply: Evidence from swedish lotteries. *American Economic Review*, 107(12):3917–3946.

Cesarini, D., Lindqvist, E., Östling, R., and Wallace, B. (2016). Wealth, health, and child development: Evidence from administrative data on swedish lottery players. *The Quarterly Journal of Economics*, 131(2):38–88.

Chetty, R. (2006). A new method of estimating risk aversion. *American Economic Review*, 96:1821–1834.

Colonelli, E., Neto, V. P., and Teso, E. (2024). Politics at work. *Working Paper*.

Daruich, D. and Fernandez, R. (2024). Universal basic income: A dynamic assessment. *American Economic Review*, 114(1):38–88.

de Oliveira, V. T. (2020). *Remuneração e Previdência dos Militares*. https://www.conint.com.br/livro/indice.htm.

Dias, M. and Rocha, R. H. (2021). Racial gaps in the labor market: The role of nonwhite entrepreneurship. *Working Paper*.

Giupponi, G. (2024). When income effects are large: Labor supply responses and the value of welfare transfers. *Working Paper*.

Golosov, M., Graber, M., Mogstad, M., and Novgorodsky, D. (2024). How americans respond to idiosyncratic and exogenous changes in household wealth and unearned income. *Quarterly Journal of Economics*, 139(2).

González, L. (2013). The effect of a universal child benefit on conceptions, abortions, and early maternal labor supply. *American Economic Journal: Economic Policy*, 5(3):160–188.

Hoynes, H. and Rothstein, J. (2019). Universal basic income in the united states and advanced countries. *Annual Review of Economics*, 11(1):929–958.

Huang, H.-H., Han, H.-W., Lo, K.-T., and Yang, T.-T. (2023). Liquidity constraints, cash windfalls, and entrepreneurship: Evidence from administrative data on lottery winners. *Working Paper 2303.17029, arXiv.org.*

Imbens, G. W., Rubin, D. B., and Sacerdote, B. I. (2001). Estimating the effect of unearned income on labor earnings, savings, and consumption: Evidence from a survey of lottery players. *American Economic Review*, 91(4):778–794.

Jones, D. and Marinescu, I. (2022). The labor market impacts of universal and permanent cash transfers: Evidence from the alaska permanent fund. *American Economic Journal: Economic Policy*, 14(2):315–340.

Kleven, H. J. (2021). Sufficient statistics revisited. *Annual Review of Economics Volume*, 13:515–538.

Machado, C., Neto, V. P., and Szerman, C. (2024). Firm and worker responses to extensions in paid maternity leave. *Journal of Human Resources*, Forthcoming.

MaCurdy, T. E. (1981). An empirical model of labor supply in a life-cycle setting. *Journal of Political Economy*, 89(6):1059–1085.

Malkova, O. (2018). Can maternity benefits have long-term effects on childbearing? evidence from soviet russia. *The Review of Economics and Statistics*, 100(4):691–703.

Picchio, M., Suetens, S., and van Ours, J. C. (2018). Labour supply effects of winning a lottery. *Economic Journal*, 128(611):1700–1729.

Riphahn, R. T. and Wiynck, F. (2017). Fertility effects of child benefits. *Journal of Population Economics*, 30(4):1135–1184.

Robins, P. (1985). A comparison of the labor supply findings from the four negative income tax experiments. *Journal of Human Resources*, 20(4):567–582.

Rocha, R. H. (2021). Labor supply responses to deaths of family members in the presence of survivors. *Working Paper.*

Van Doornik, B., Fazio, D., Schoenherr, D., and Skrastins, J. (2022). Unemployment insurance as a subsidy to risky firms. *The Review of Financial Studies*, 35, Issue 12:5535–5595.

Yonzan, N., Timilsina, L., and Kelly, I. R. (2024). Economic incentives surrounding fertility: Evidence from alaska's permanent fund dividend. *Economics & Human Biology*, 52.

Figure 1: Literature on marginal propensity to earn

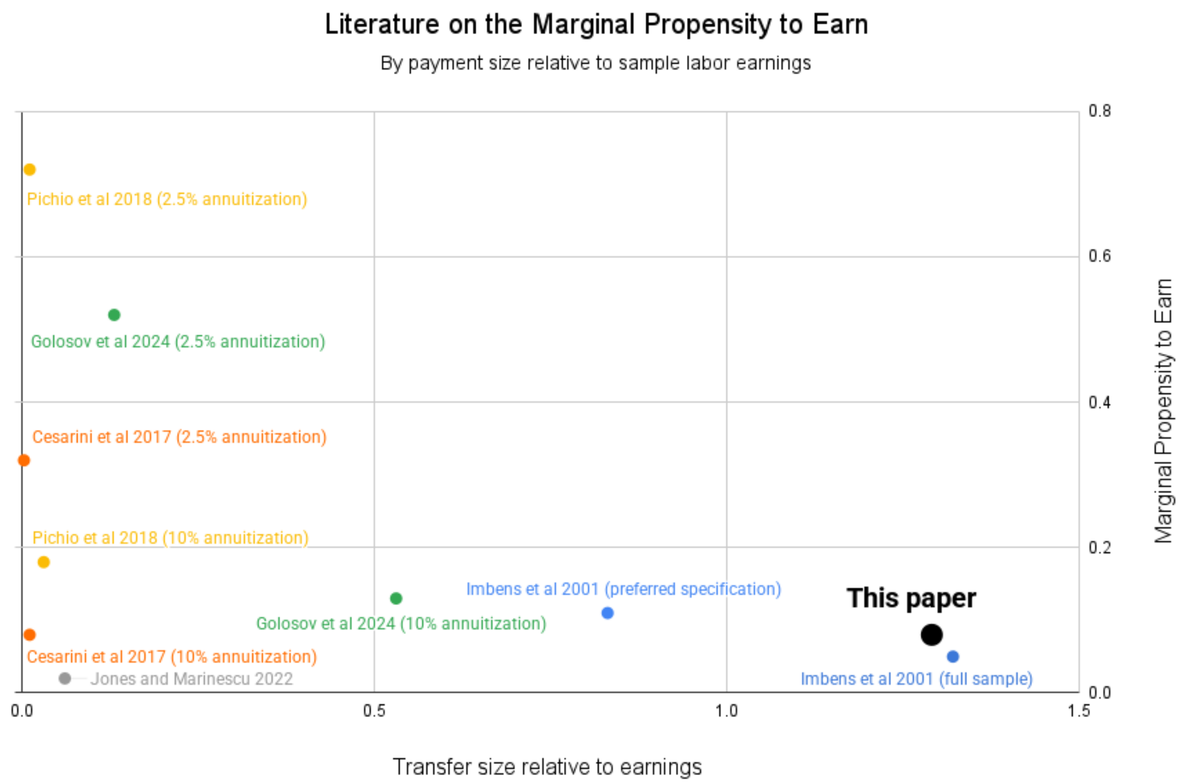

This figure plots estimates of the MPE against the relative size of the income transfer, either flow of income divided by sample mean earnings, or change in wealth multiplied by annuitization rate and divided by sample mean earnings. Our estimate (”This paper”) is characterized by a large transfer size and a small MPE, and is invariant to choice of annuitization. Dots with same color represent the same paper, either a different subsample as in (Imbens et al., 2001) or different annuitization rates. Most studies shown use lottery winnings to estimate the MPE ((Cesarini et al., 2017), (Picchio et al., 2018), (Golosov et al., 2024), (Imbens et al., 2001)). Other estimates are derived from Alaska Permanent Fund dividends (Jones and Marinescu, 2022). For (Imbens et al., 2001), the flow of income was adjusted downward to reflect that the life expectancy in their sample was likely longer than the duration of prize annuity (20 years).

Table 1: Sample Construction

	(1)	(2)	(3)	(4)
	All Beneficiaries	Only Daughters	Start Date $\geq$ 2002	Employed in Jan. 2002
Observations	**258,103**	**159,077**	**114,989**	**30,400**
Payment start year	2008	2009	2015	2015
	(12.7)	(12.1)	(6.4)	(6.3)
Officer when alive	0.57	0.58	0.59	0.67
	(0.60)	(0.61)	(0.61)	(0.47)
Co-beneficiaries	2.36	2.84	2.84	2.80
	(1.58)	(1.66)	(1.64)	(1.58)
Monthly Pension (BRL)	8808	7226	7090	7785
	(7765)	(7091)	(6882)	(7003)
Year of Birth				1961
				(8.3)

This table shows the sample construction process. Column (1) includes all beneficiaries aged 16+. Column (2) restricts the sample to daughters. Column (3) requires the pension to have started in 2002 or later. Column (4) restricts to individuals employed on Jan. 1, 2002. Standard deviations are in parentheses.

Table 2: Summary Statistics of the Final Sample

	Mean	SD	Median	P25	P75
Year of first payment	2015	6.3	2016	2010	2021
Age	47.9	9.4	48	41	54
Monetary Values (BRL)					
Monthly Pension	7785	7002	5522	2890	10174
Monthly Earnings	6025	8953	2953	0	7850
Labor Market Outcomes					
Weekly Hours	24.8	18.5	30	0	40
Formal employment (%)	68	–	–	–	–
Hired (%)	2.20	–	–	–	–
Fired (%)	1.52	–	–	–	–
Quit (%)	0.62	–	–	–	–
College completion (%)	66.4	–	–	–	–

Statistics for the final sample ($N = 30,400$) using monthly data from 2002 to 2016 ($T = 180$), totaling 5,472,000 observations. Monetary values are adjusted to 2025 prices. Variables are unconditional on monthly employment status. For binary variables (indicated by %), only the mean is reported.

Figure 2: Change in total income due to military pensions for daughters.

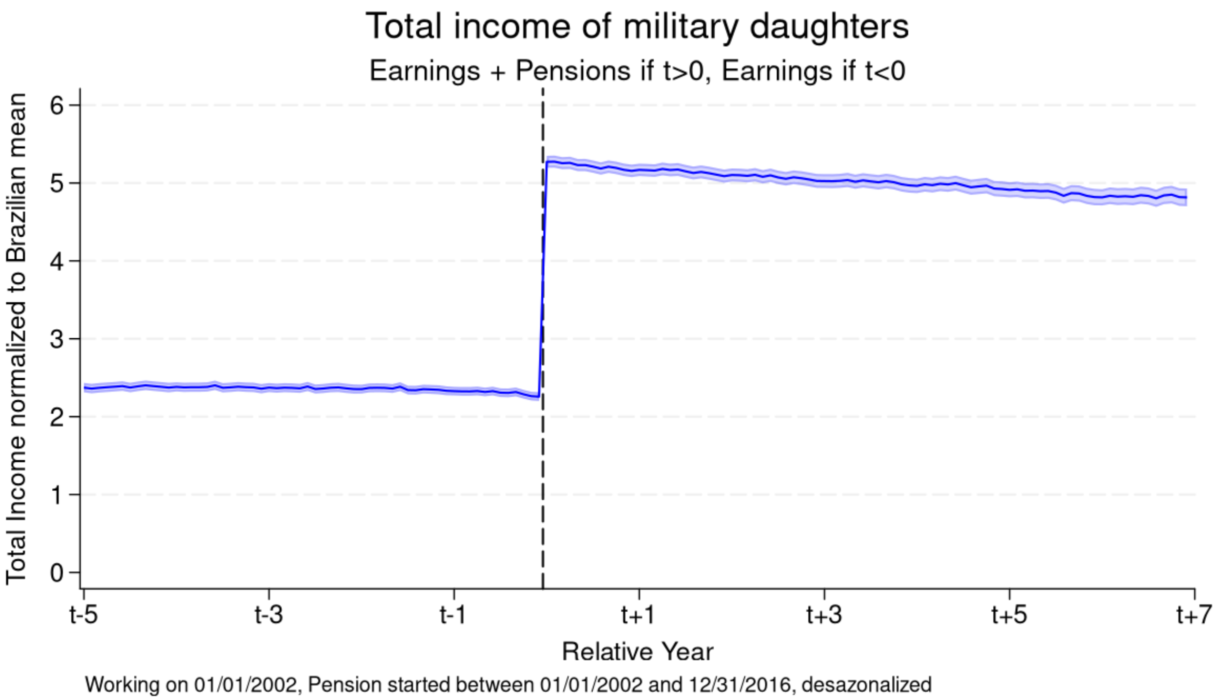

This figure shows the average total income for daughters of military personnel around the year they begin receiving pension benefits ($t = 0$). Total income is defined as labor earnings before the event and labor earnings plus pension payments after the event. Income is normalized by the national mean earnings in 2010. The sample is restricted to daughters employed on January 1, 2002, whose pensions started after January 1, 2002. The shaded area represents a 95% confidence interval. The results show that the pension more than doubles the daughters' total income.

Figure 3: Change in Pension wealth at time of parental death by daughter age

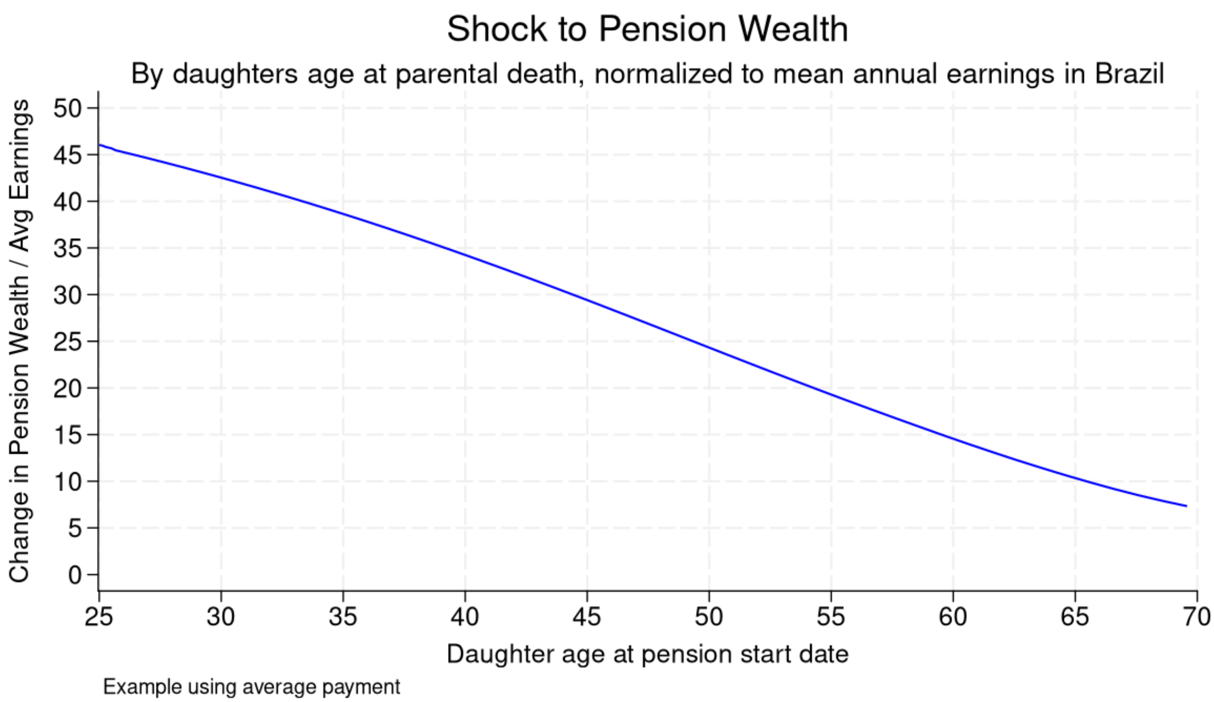

This figure illustrates the simulated lifetime value of the pension wealth shock as a function of the daughter's age at the time of parental death. The shock's formula is : $\Delta W(a) = \frac{Y}{\beta}\left(1-(1-\beta)^{D(a)}\right)$ where $\Delta W(a)$ is the unexpected change in pension wealth at age a, Y is the perpetual annual flow of pension income normalized to national mean earnings, $D(a)$ is the remaining life expectancy for a daughter of age a, β is the annuitization rate. The plot shows that the pension represents a substantially larger wealth shock for younger daughters.

Figure 4: The Impact of Pension Income on Formal Labor Earnings

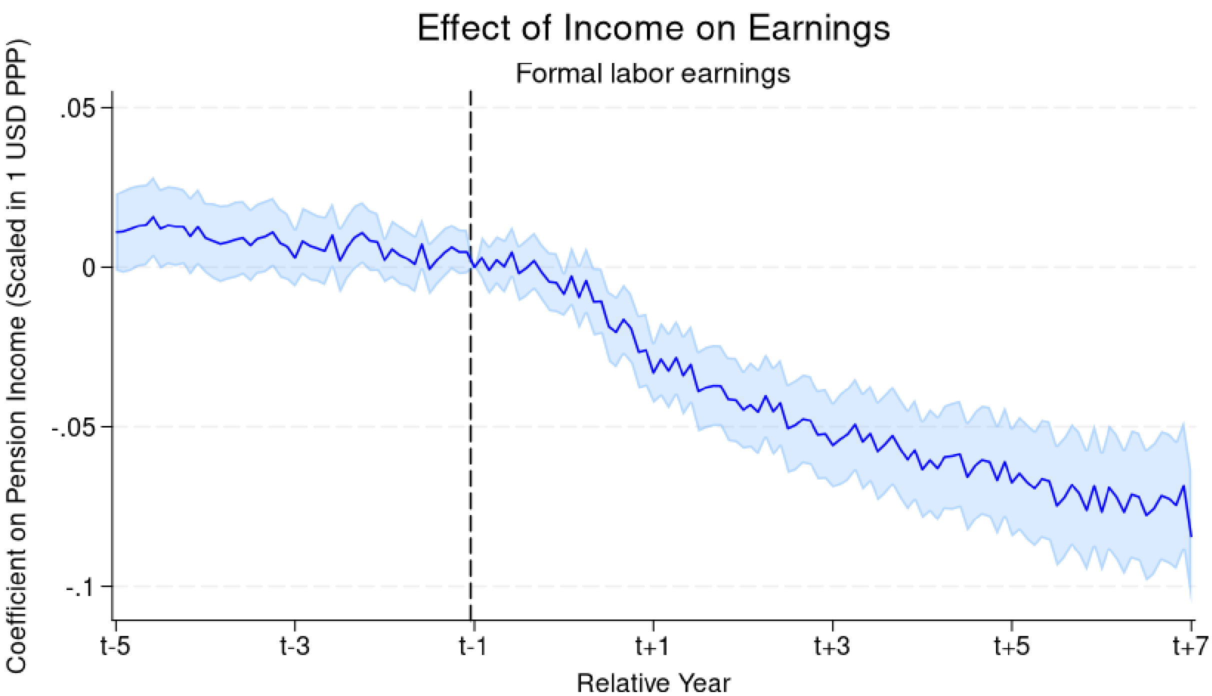

This figure plots the coefficients from an event-study regression of formal labor earnings on pension income. The coefficients β_k represent the dynamic effect of receiving one dollar of pension monthly income on monthly earnings k years relative to the start of the pension. Doze months before the event $(t-1)$ is the omitted reference category. The sample consists of military daughters who were employed on January 1, 2002 and the pension started after January 1, 2002. The shaded area represents a 95% confidence interval. The results show a gradual decline in labor earnings following pension receipt, reaching approximately -0.08 after seven years, implying that daughters reduce their earnings by 8 cents for every dollar of pension income received.

Figure 5: Marginal propensity to earn (MPE) by age at first payment

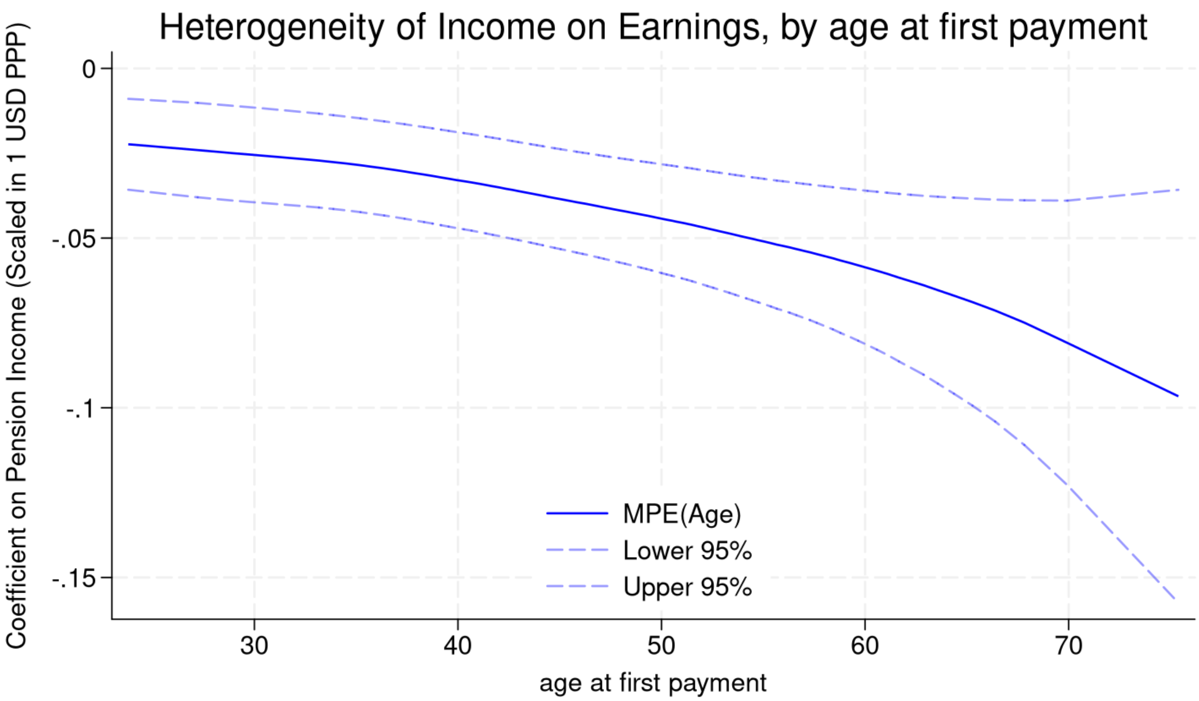

This figure plots the estimated MPE as a flexible function of the daughter's age at the first pension payment. The curve is a parametric approximation estimated by interacting the treatment variable with a cubic spline of age, allowing the effect to vary smoothly. The solid line represents the point estimate, and the dashed lines show the 95% confidence interval. The sample consists of military daughters who were employed on January 1, 2002 and the pension started after January 1, 2002. The MPE is modest and close to zero for the youngest daughters in the sample but becomes substantially more negative as age increases. The figure shows a much stronger labor supply response to the pension income than younger daughters, a puzzle, as the lifetime wealth shock from the pension is largest for the youngest daughters. This puzzle is resolved, however, by considering the role of liquidity constraints. A simple static labor supply model, in which agents cannot borrow against future income, predicts this exact pattern. Younger daughters, while wealthier on paper, cannot access their future pension wealth to finance current consumption and thus continue to work.

Figure 6: Heterogeneity of the MPE by Pension Payment Size

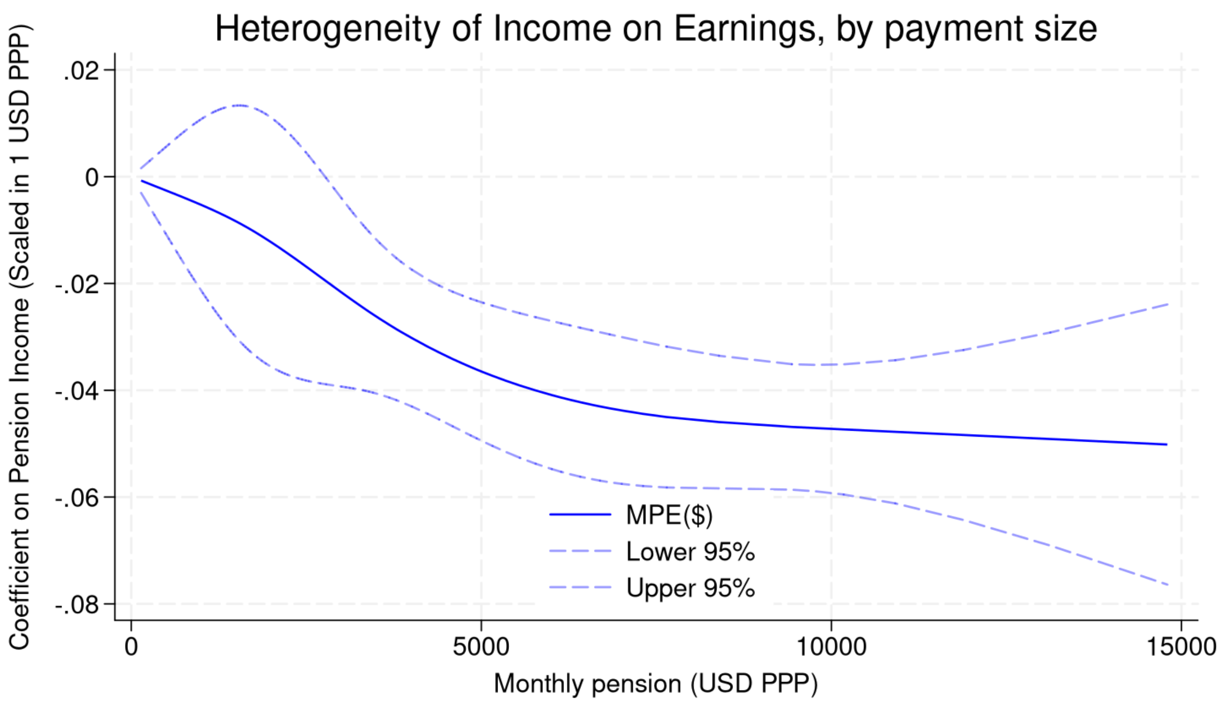

This figure plots the estimated Marginal Propensity to Earn (MPE) as a flexible function of the monthly pension payment size, measured in USD PPP. The curve is estimated by interacting the treatment indicator with a cubic spline of the pension amount. The solid line represents the point estimate, and the dashed lines show the 95% confidence interval. The sample consists of military daughters who were employed on January 1, 2002 and the pension started after January 1, 2002. The MPE is close to zero for small payments, becomes increasingly negative, and then flattens for larger pension amounts. This non-linear relationship is consistent with labor supply models featuring **adjustment costs**. The distribution of payments is concentrated at the lower end of the x-axis: 75% of the sample receives pensions below $5,000 per month, and 90% receive pensions below $10,000 per month.

Figure 7: Changes in Welfare on the margin

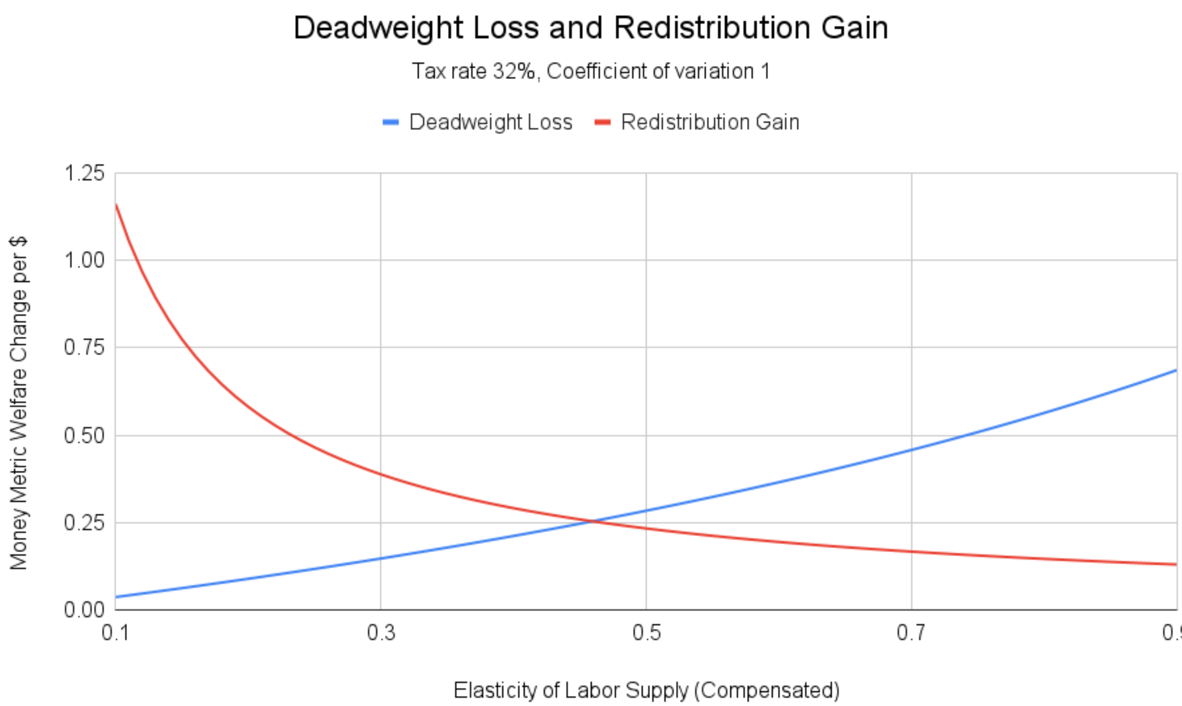

This figure illustrates the welfare implications of the transfer program, plotting the two key components of the welfare change against the compensated elasticity of labor supply. The simulation assumes a tax rate of 32.32%, the total government revenue as function of the GDP for the year of 2024, and a coefficient of variation for income of 1.005, consistent with the Census 2010 after winsorizing 2% of the sample. The blue line represents the Deadweight Loss (efficiency cost) per dollar transferred. This cost is near zero for very low elasticities but rises exponentially. The red line represents the Redistribution Gain per dollar transferred. This gain, which captures the utility increase from moving money from high-income (low marginal utility) to low-income (high marginal utility) individuals, is substantial when the substitution effect is small, due to the link between the ratio of income to substitution effects and the curvature of the utility function.

Figure 8: Optimal Tax and Transfer

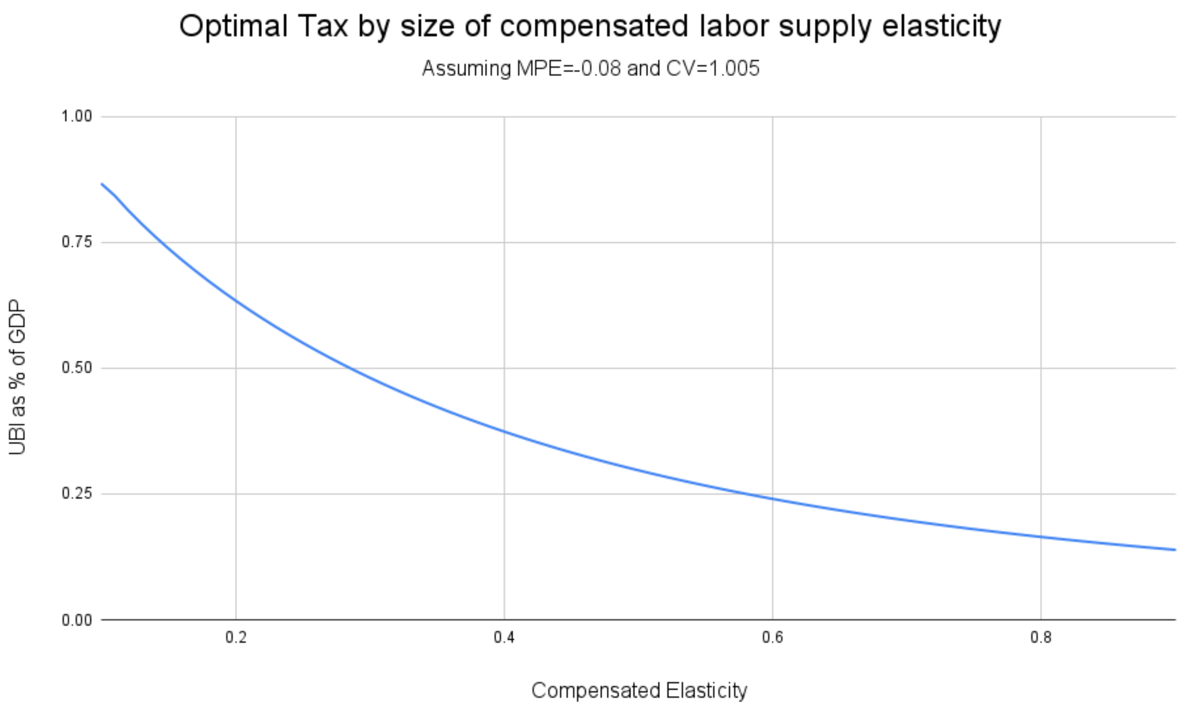

This figure plots the simulated optimal size of a Universal Basic Income (UBI) program, measured as a percentage of GDP, as a flexible function of the compensated labor supply elasticity. The simulation holds the Marginal Propensity to Earn (MPE) fixed at -0.08, the central empirical estimate of this paper. This curve illustrates the core welfare trade-off. A lower compensated elasticity increase the optmal redistribution by two mechanisms (i) Higher curvature of the utility function (ii) Lower deadweight loss.

Appendix A. Record linkage details

Identifiers. The CPF has 11 digits: the first eight are randomly assigned (roughly 100 million combinations), the ninth digit reflects the region of issuance (e.g., applications in the state of São Paulo receive "8"), and the last two are verification digits. For privacy, the open-access public files used here publish only digits 4–9; I refer to this six-digit fragment as *CPF6*. The linkage key is therefore {full name, CPF6}.

Brazilian naming conventions. Brazilian naming conventions help stabilize matches over time. Individuals often have one or two given names (e.g., *Carla*, *Maria Rita*) followed by multiple family names, with the last family name typically being the father's surname (e.g., *Carla Silva Passos*). Upon marriage, women commonly either append the spouse's surname at the end (e.g., *Carla Silva Passos Oliveira*) or drop an intermediate family name (e.g., *Carla Passos Oliveira*), while retaining the father's surname. These structured tail-end changes make it feasible to recognize the same person across files even when a surname is added or removed at marriage.

Matching algorithm. The procedure has two steps. First, I form the candidate groups through a *exact* match in CPF6 and the first given name. Second, within each candidate group, I compute a fuzzy string similarity score for the *full* name and accept a pair as a link if (i) it has the *highest* score in its candidate group and (ii) the score is at least 0.75. Similarity scores are produced by a standard fuzzy matching routine that rewards small edit distances and common prefixes.[14] Pairs with a top score below 0.75 are not linked. The 0.75 threshold was chosen after visually inspecting matched and near-miss pairs to ensure that (a) marriage-consistent changes (one new surname appended at the end or one middle surname dropped) are retained and (b) differences due to abbreviations, spacing or minor typos are accepted, while more substantial rearrangements are not.

Accuracy and validation. Three pieces of evidence speak to the accuracy. (i) Among the accepted links, 65% are *perfect* name matches (score = 1.00), and 86% have scores greater than 0.9; Figure 20 plots the score distribution. (ii) A manual review of 200 randomly selected accepted pairs (stratified by score) confirms that the algorithm behaves as intended. (iii) I validate the approach inside the labor-market data, where the full 11-digit CPF is available: I draw two 10% samples, mask the identifiers to CPF6 in each, run the same blocking and fuzzy scoring as in the cross-dataset linkage, and then verify the identity using the full CPF. This experiment yields a false-positive rate of 0.5%; restricting to imperfect name matches, the false-positive rate is 2.5%; for perfect name matches it is 0.01%.

[14] Raffo (2016), Stata Users' Meeting, fuzzy matching slides.

Appendix B. Gains from redistribution

To approximate the gains from redistribution, first we need to explicitly write the gains from redistribution as function of γ and the level of redistribution τ.

Using a CARA utility function, the marginal utility is $u_c(c_i) = \gamma\, c_i$. Consumption is $c_i = b + (1-\tau)\, w_i\, l_i$. The basic income is $b = \tau Y$ when there is a mass one of agents. If we normalize the monetary values so that $Y = \int_i w_i l_i F(i) = 1$. Then $c_i = \tau + (1-\tau)\, w_i\, l_i$,

Finally, the gain from redistribution is

$$g(\gamma, \tau) = -1 + \frac{\int e^{-\gamma((1-\tau)w_i l_i + \tau} dF_i}{\int w_i l_i e^{(} - \gamma((1-\tau)w_i l_i + \tau) dF_i}$$

First, we can factor out the constant term $e^{-\gamma\tau}$ from the numerator and denominator, which then cancels:

$$g(\gamma, \tau) = \frac{\int e^{-\gamma(1-\tau)w_i l_i} dF_i}{\int w_i l_i e^{-\gamma(1-\tau)w_i l_i} dF_i}$$

Let's define a new constant $k = \gamma(1-\tau)$. Since $\gamma > 0$ and $0 < \tau < 1$, we know $k > 0$. The function simplifies to a ratio of expected values:

$$g(k) = \frac{E[e^{-kY}]}{E[Y e^{-kY}]}$$

We can approximate this function using a first-order Taylor series expansion around $k = 0$:

$$g(k) \approx g(0) + g'(0)\, k$$

The derivative of $g(k)$ is

$$g'(k) = -1 + \frac{\mathbb{E}[Y^2\, e^{-kY}]}{\mathbb{E}[Y e^{-kY}]} \times \frac{\mathbb{E}[e^{-kY}]}{\mathbb{E}[Y e^{-kY}]}$$

then $g'(0) = -1 + \mathbb{E}[Y^2]$ and since $g(0) = 1$, the final expression is

$$g(k) = \big(-1 \,+\, \mathbb{E}[Y^2]\big)\, \gamma\, (1-\tau)$$

However, the distribution of $w_i l_i$ was normalized to have mean one, in a general case the first order approximation of the gains from redistribution are

$$g(\gamma, \tau) = [CV(w_i\, l_i)]^2 \times \gamma \times (1-\tau)$$

Figure 9: Generosity and Scale of Brazilian Social Insurance Programs (Normalized)

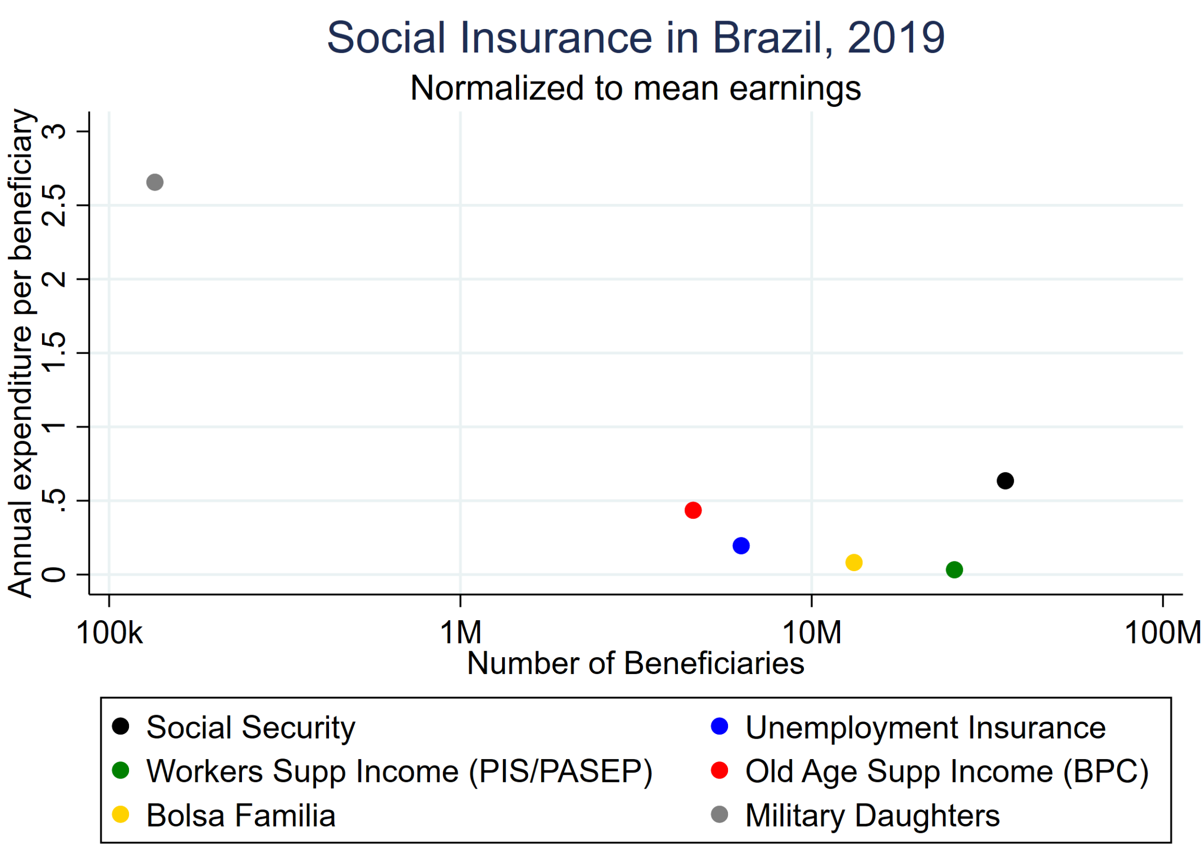

This figure benchmarks the military daughters' pension against other major social insurance programs in Brazil for the year 2019. The x-axis shows the number of beneficiaries on a logarithmic scale, while the y-axis shows the average annual expenditure per beneficiary, normalized to the national mean earnings. The military daughters' program is a clear outlier, providing benefits that are orders of magnitude more generous per person than any other large-scale program.

Figure 10: Generosity and Scale of Brazilian Social Insurance Programs (USD PPP)

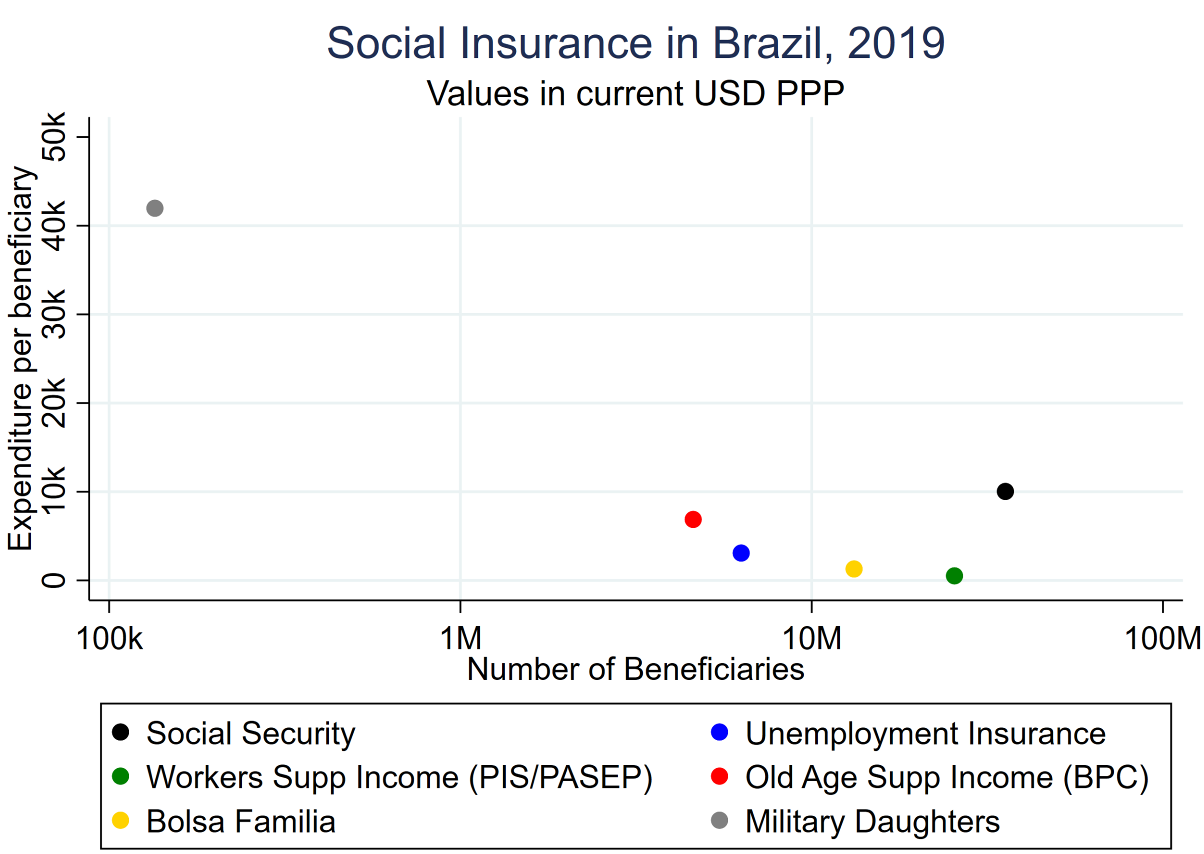

This figure provides an alternative visualization of Brazil's social insurance landscape in 2019, with annual expenditure per beneficiary shown in current US dollars (PPP). As with the normalized version, the military daughters' pension stands out for its exceptionally high per-beneficiary value, highlighting the unique nature of this transfer program.

Figure 11: Distribution of Parental Age at Childbirth

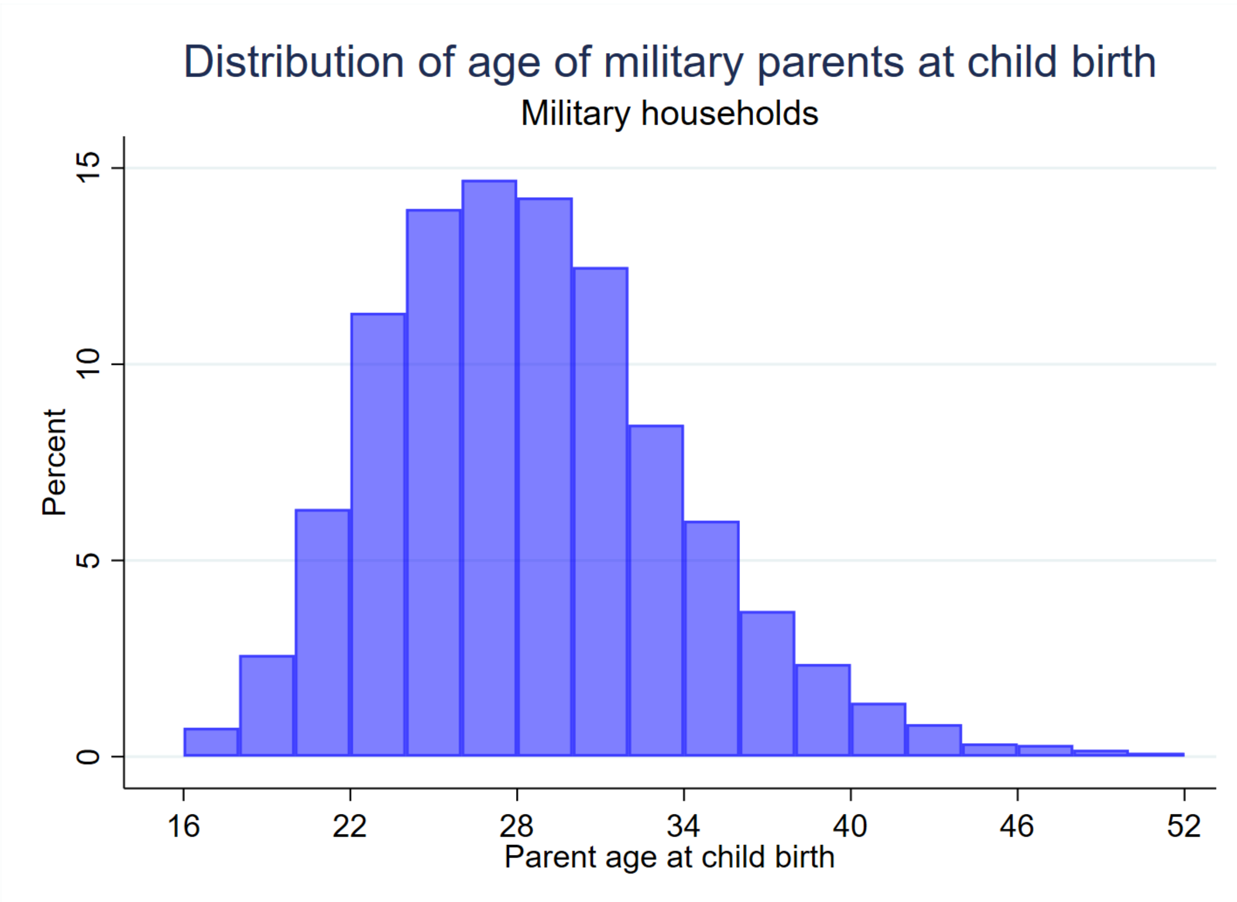

This figure shows a histogram of the military parent's age at the time of their daughter's birth. The distribution is unimodal and right-skewed, with a peak around age 28. This distribution is a key input for simulating the parent's age conditional on the daughter's age.

Figure 12: Cumulative Mortality Risk by Age in Brazil

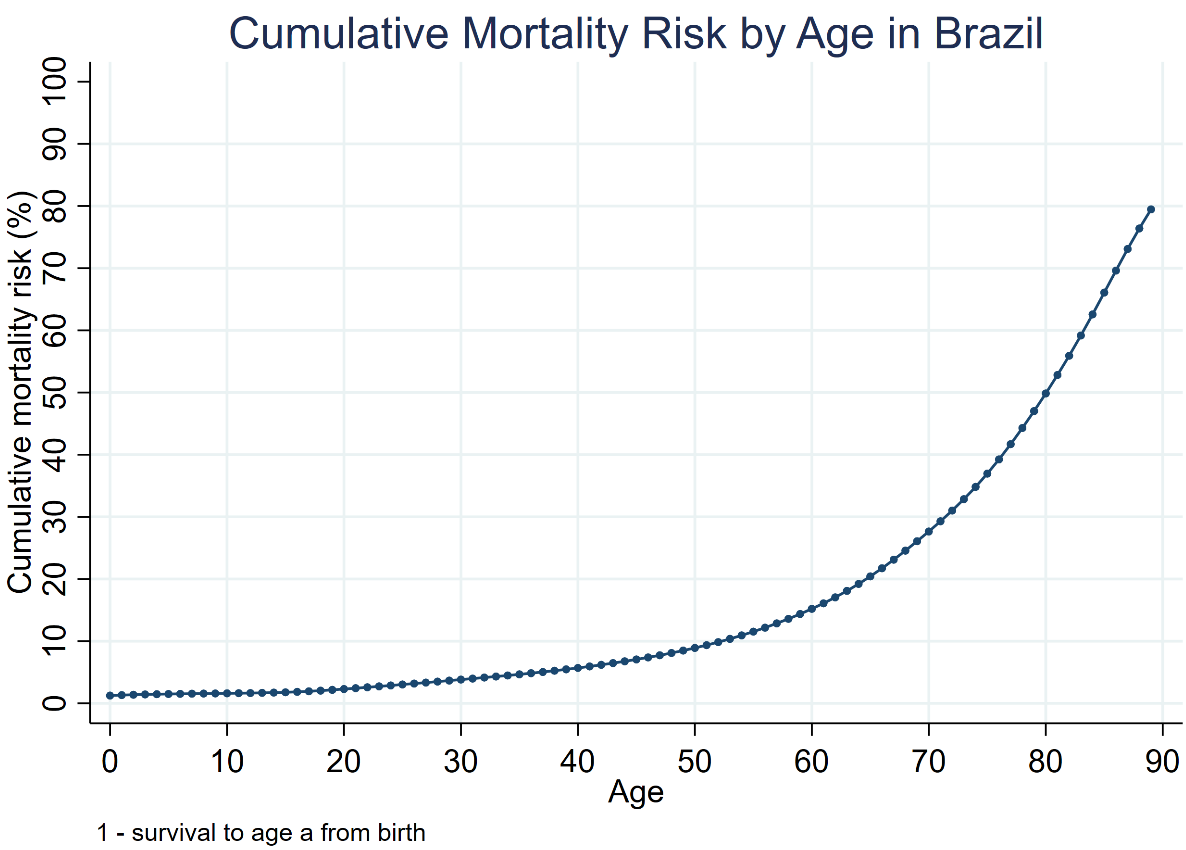

This figure plots the cumulative mortality risk from birth by age for the general Brazilian population, based on national life tables. It represents the probability of dying by a certain age. This mortality schedule is a primary input for simulating the probability of parental.

Figure 13: Annual Mortality Risk by Age in Brazil

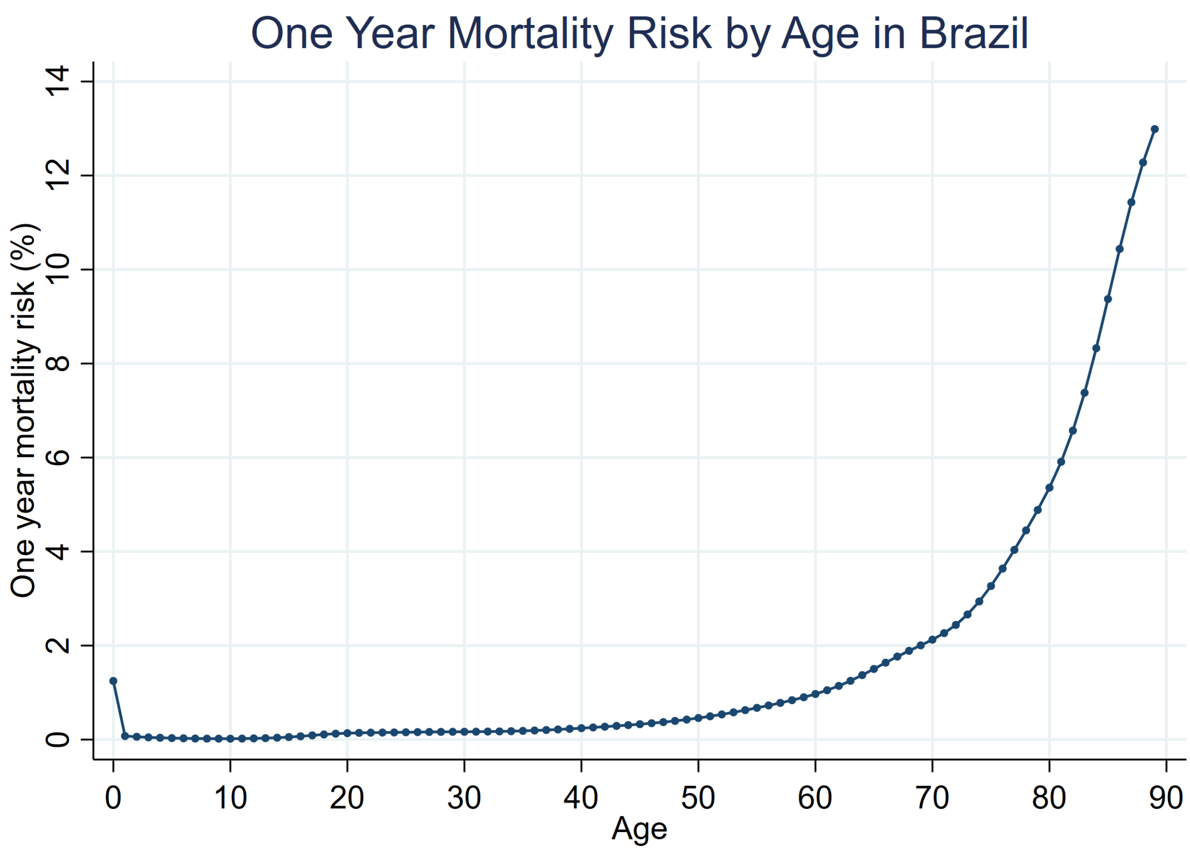

This figure plots the one-year mortality risk, representing the probability of dying in the next year, conditional on having survived to a given age. The data is from Brazilian life tables. This conditional probability is used to calculate the annual risk of parental loss for a daughter of a specific age.

Figure 14: Simulated Annual Probability of Parental Loss by Daughter's Age

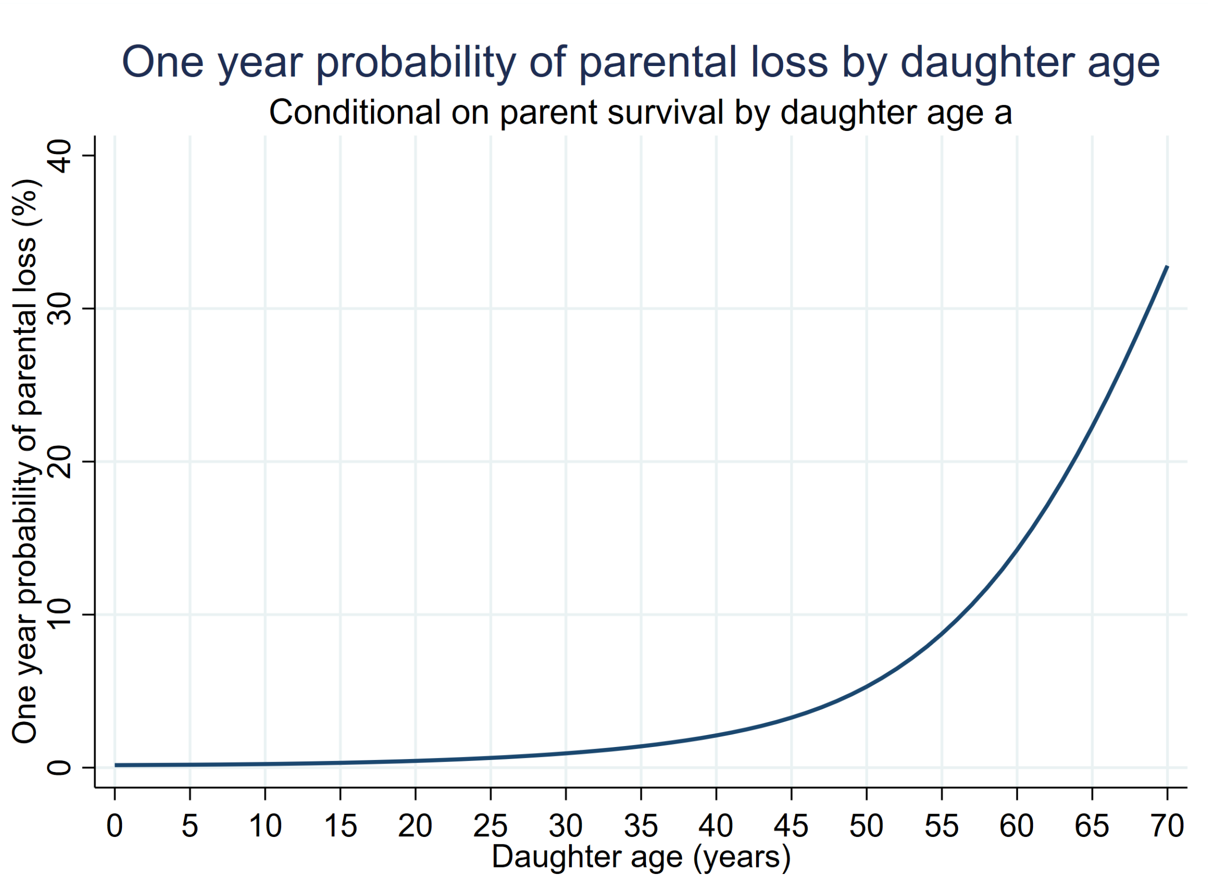

This figure shows the simulated one-year probability that a daughter will experience parental loss, as a function of her own age. This probability is derived by combining the distribution of parental age at childbirth (from the first figure) with the conditional mortality risk schedule (from the third figure). The risk is low during childhood and begins to increase exponentially as the daughter enters her 50s and 60s.

Figure 15: Simulated Cumulative Probability of Parental Loss by Daughter's Age

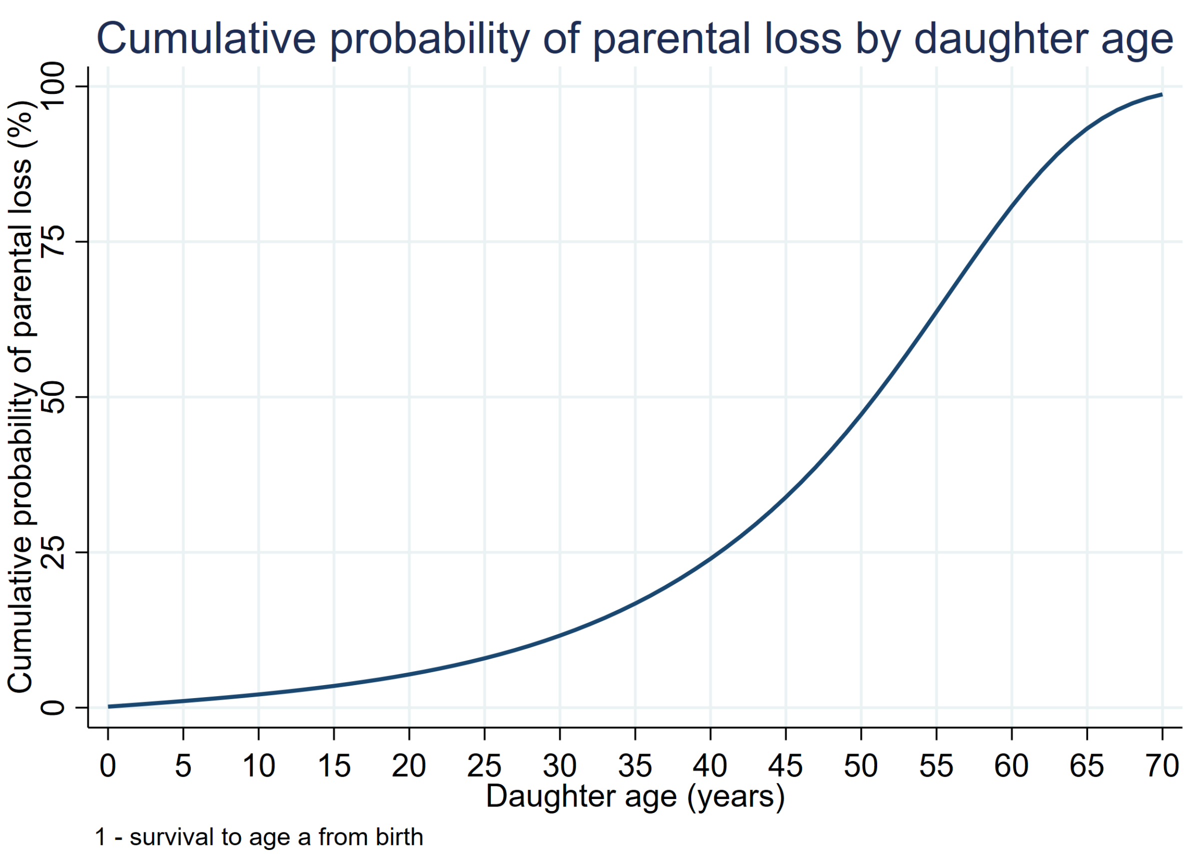

This figure plots the simulated cumulative probability of having experienced parental loss by a certain age. This curve is constructed by aggregating the annual probabilities shown in the previous figure. It shows, for example, that a daughter has approximately a 50% cumulative probability of losing her parent by the time she reaches age 51.

Figure 16: Simulated Parental Life Expectancy Conditional on Daughter's Age

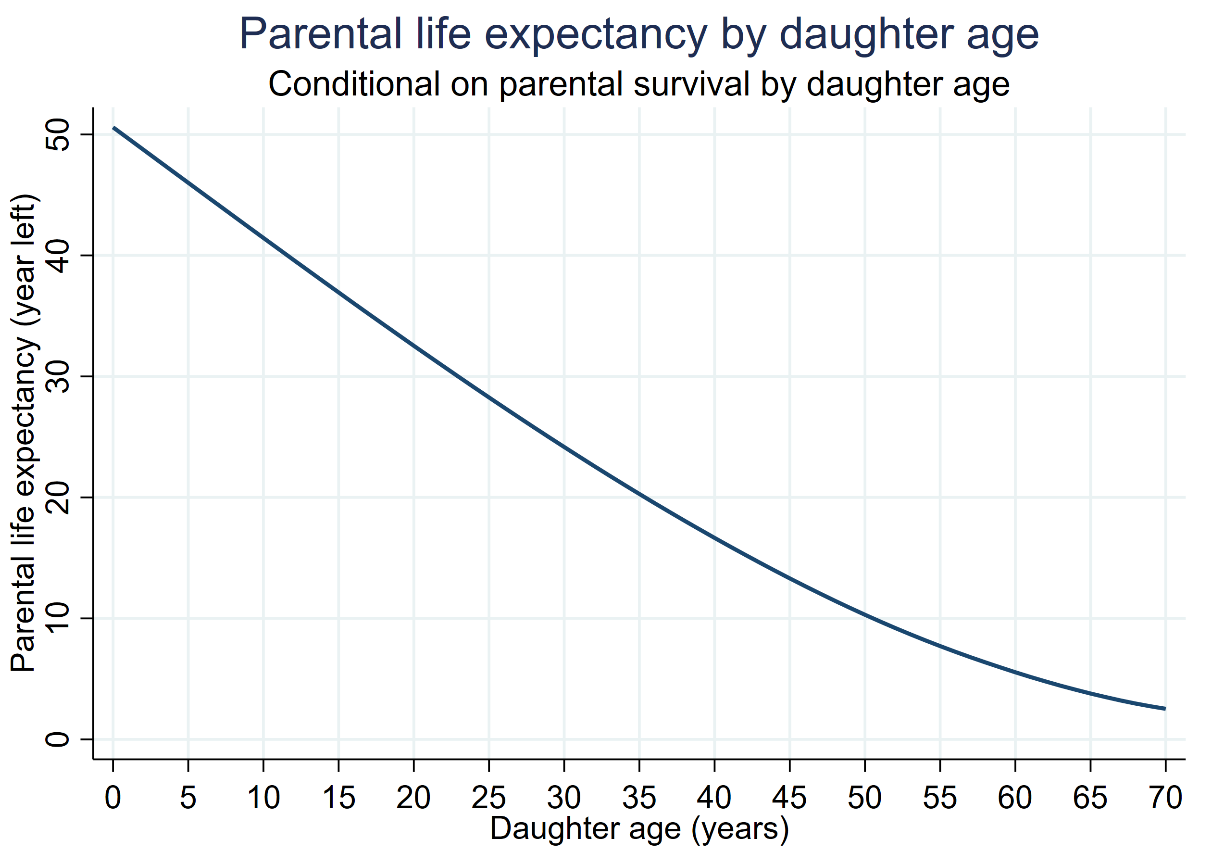

This figure shows the simulated life expectancy of a parent, conditional on their survival up to their daughter's current age. This is calculated using the parental age at daughter birth distribution and the general mortality risk curves shown previously. The declining curve illustrates that as the daughter ages, her surviving parent is also older and thus has a shorter remaining life expectancy.

Figure 17: The Pension Wealth Shock at Parental Death

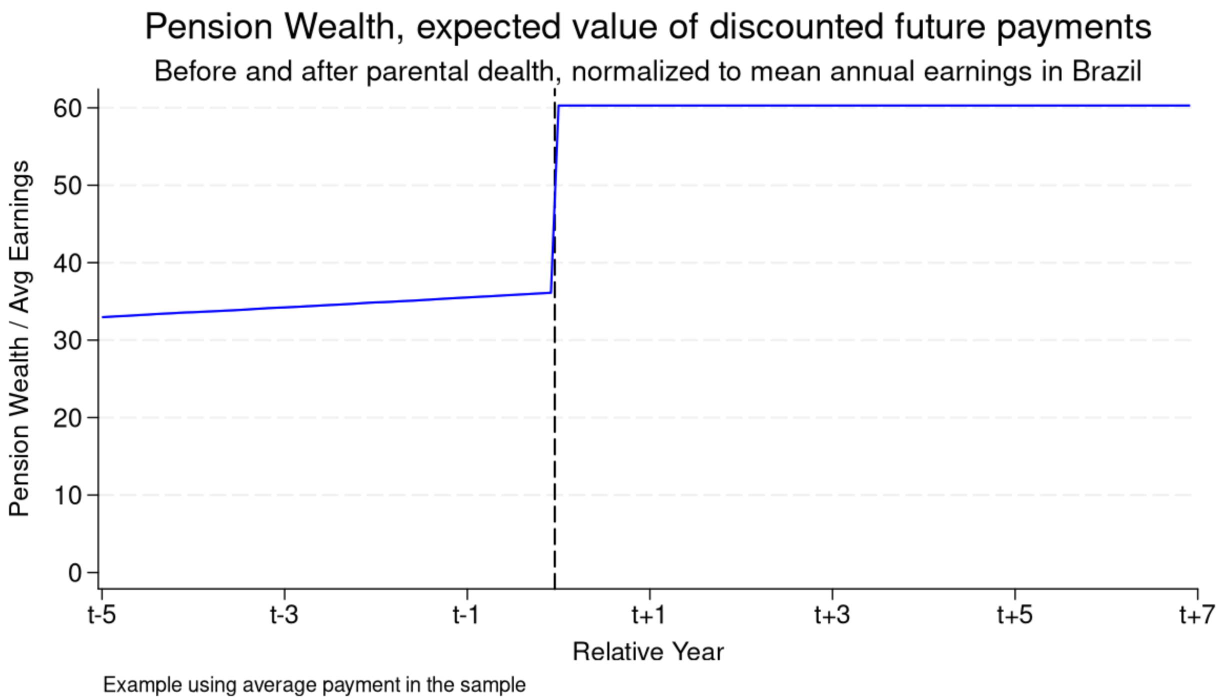

This figure illustrates the expected present discounted value of future pension payments, before and after the death of the military parent. The event occurs at t=0. Expected wealth is lower before the event $t = 0$ because daughters expected to be paid only years ahead, depending on parental life expectancy. Wealth is normalized by national mean earnings. The formula is $W(t) = \frac{Y}{\beta}$ if $t \geq 0$ and $W(t) = \frac{Y}{\beta}(1-\beta)^{D(t)}$ if $t < 0$ where β is the annuitization rate, Y the pension flow, $D(t)$ the parental life expectancy.

Figure 18: Co-residence Rates of Adult Daughters by Parent's Age

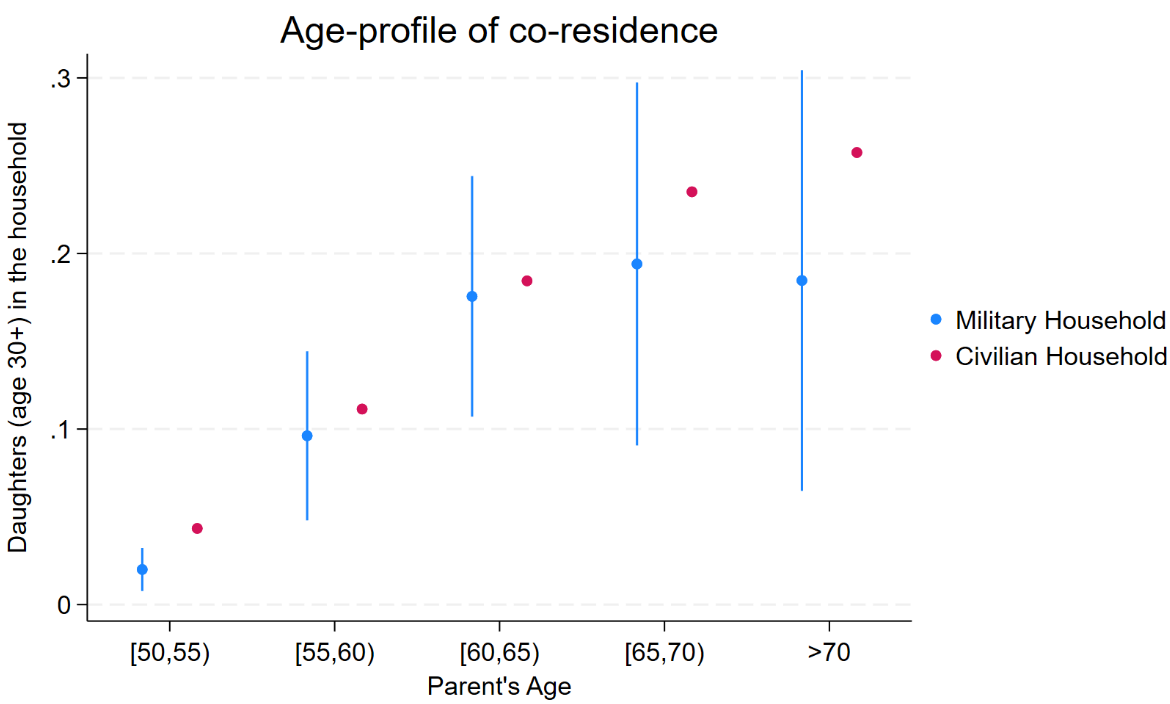

This figure plots the number of adult daughters (age 30+) who co-resides with her parents, binned by the parent's age. It compares rates in military households to those in civilian households. The data suggests that co-residence is relatively uncommon (considering that Fertility Rate was higher than 4 before the 1980s) and that daughters of military personnel are equally likely to live with their parents than their civilian counterparts, supporting the interpretation of the pension as an independent shock to the daughter's resources.

Figure 19: Female Age-Employment Profiles in Brazil, 2000 and 2010

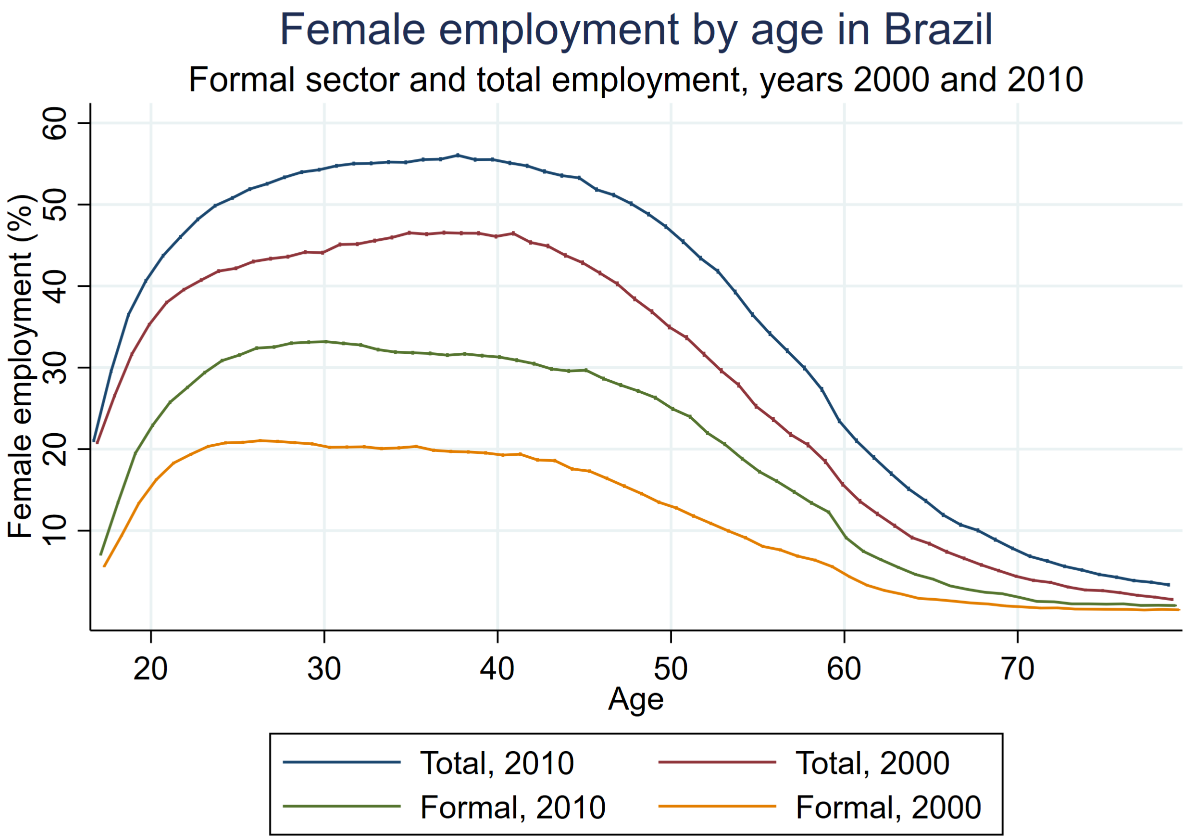

This figure displays the female employment-to-population ratio by age in Brazil for the years 2000 and 2010, separating formal sector employment from total employment. The data shows the classic inverted U-shaped profile, with employment peaking in middle age. A notable trend is the significant increase in both formal and total female labor force participation across all age groups between 2000 and 2010. Considering female formal labor force participation in the year 2000, the sample loss in table 1 from column (3) to column (4) is low, likely due to their higher educational attainment.

Figure 20: Quality of Name Matching Between Datasets

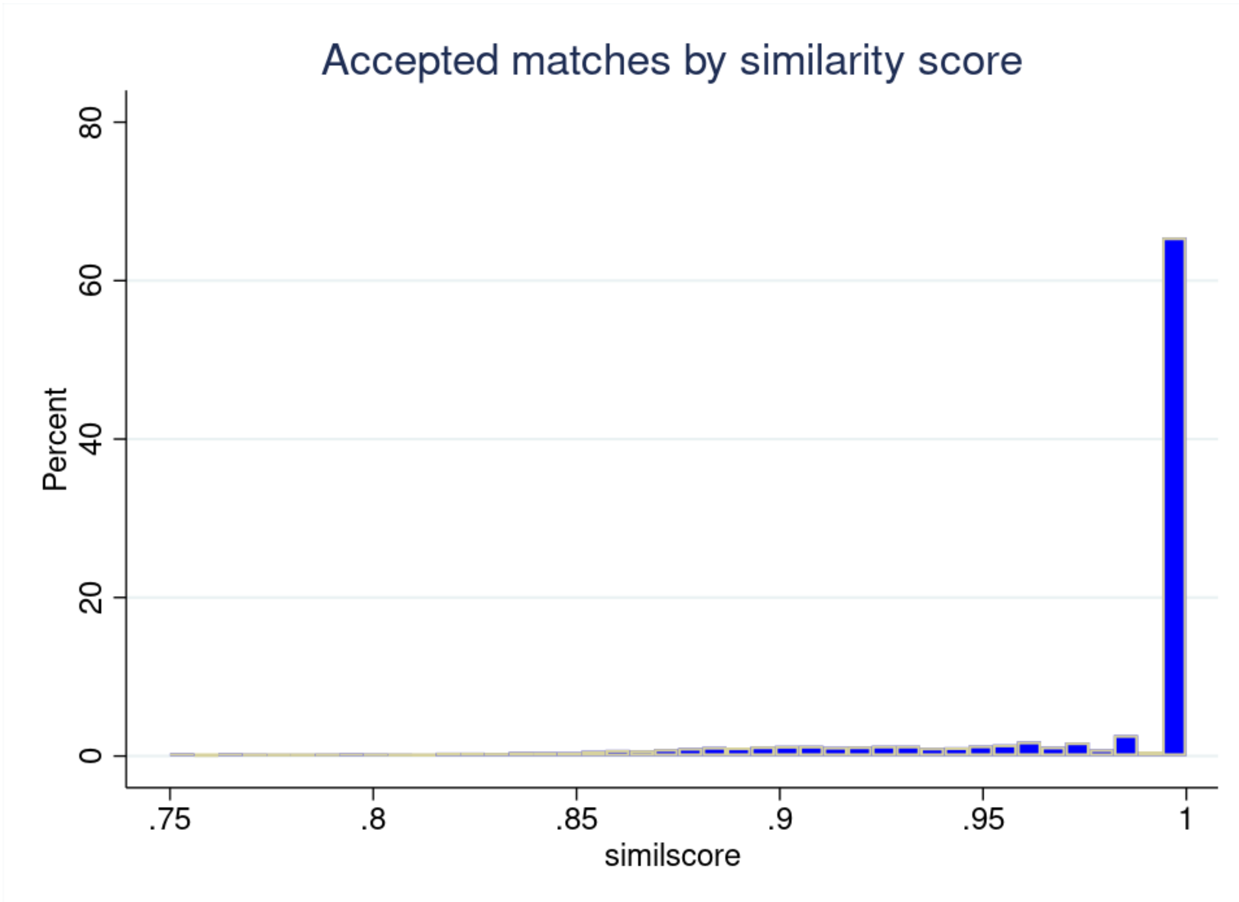

This figure displays a histogram of the string similarity scores (similscore) for the name-matching algorithm used to link the pension administrative data with the formal labor market records. All individuals compared share digits 4 to 9 of their tax code (which spans 1 million combinations). The overwhelming concentration of scores at or very near 1 indicates a very high-quality match, providing confidence that the same individuals are being tracked across datasets. Similarity scores near 1 indicates a difference in names likely of a single letter due to misspelling.

Figure 21: Administrative Delay Between Death and First Pension Payment

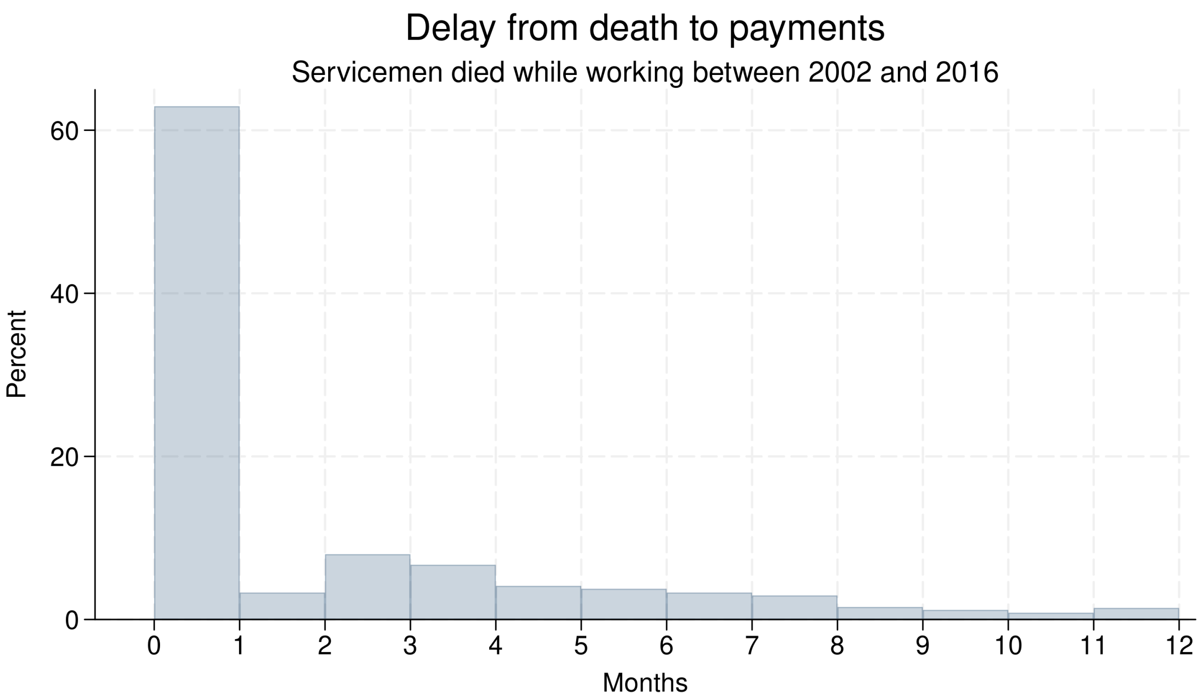

This figure shows a histogram of the number of months between the servicemen's death and the daughter's first pension payment. The sample is restricted to servicemen who died while working between 2002 and 2016. Although 60% receive their pension in the same month, it can take up to a year between death and payment.

Figure 22: Distribution of Beneficiaries per Servicemen

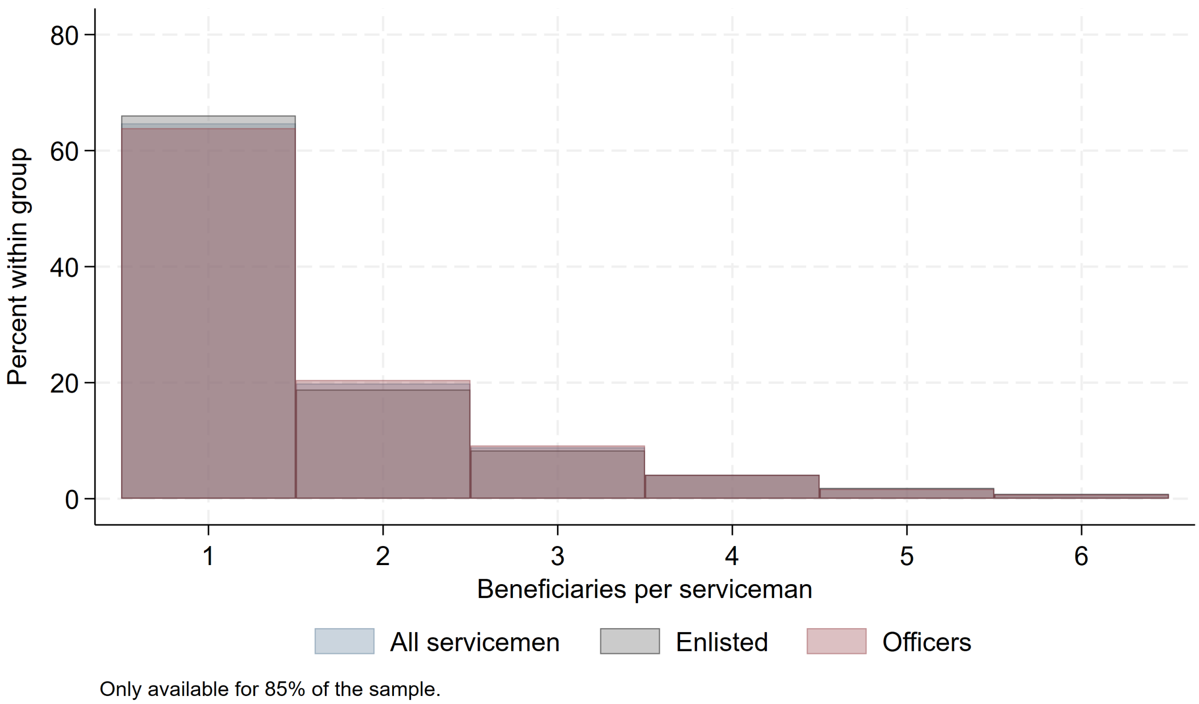

This figure shows a histogram of the number of beneficiaries associated with a single servicemen's pension, for the 85% of the sample where this information is available. It is separated for officers and enlisted personnel.

Figure 23: Distribution of Daughter's Age at First Pension Payment

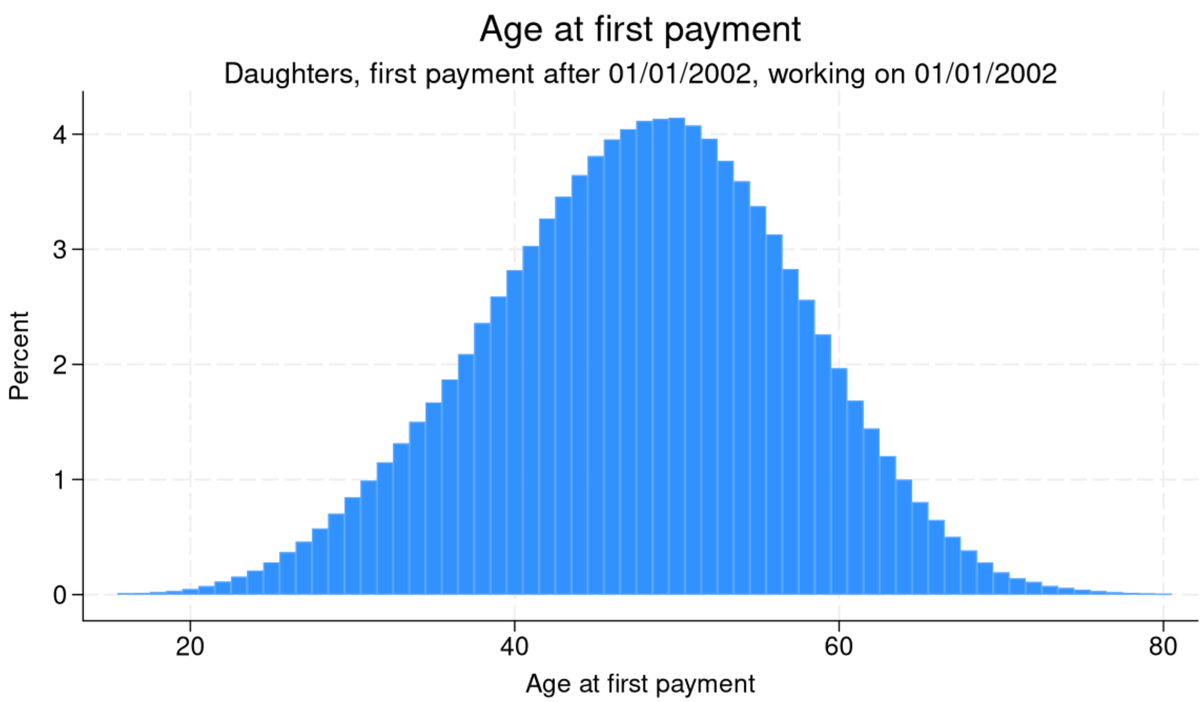

This figure presents a histogram showing the distribution of age for daughters in the analysis sample, specifically at the time they receive their first pension payment. The sample includes daughters whose first pension payment occurred after January 1, 2002, and who were actively working on January 1, 2002. The distribution is approximately normal, centered around a mean age of approximately 50, with ages ranging from the early 20s to around 80. This reflects the broad age spectrum of beneficiaries at the onset of their pension.

Figure 24: Distribution of Pension Payments Relative to Mean Earnings

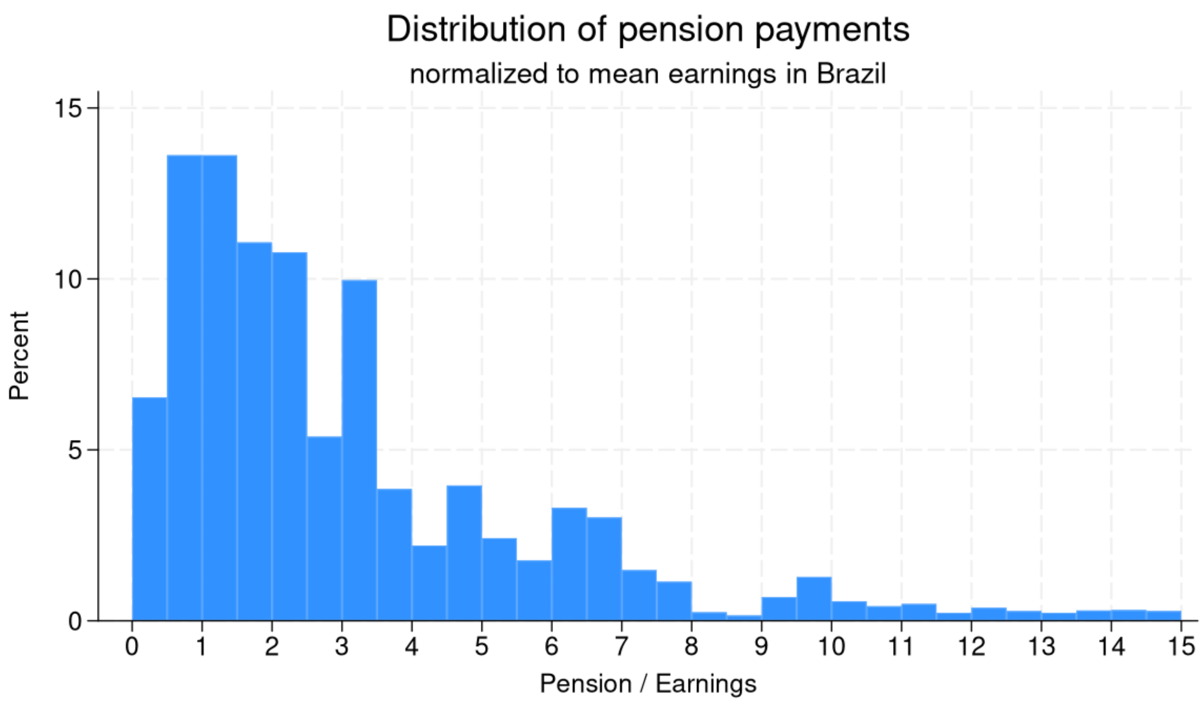

This figure shows a histogram of monthly pension payments for daughters in the final analysis sample. The x-axis plots the pension value normalized by the average monthly labor earnings in Brazil, providing a measure of the pension's generosity relative to the national wage. The distribution is highly right-skewed, with a median payment above 2 times the mean earnings. While most beneficiaries receive pensions up to 4 times the average wage, a long tail indicates that a small fraction of individuals receives exceptionally large payments.

Figure 25: Labor Supply Response vs. Estimates from Lottery Literature

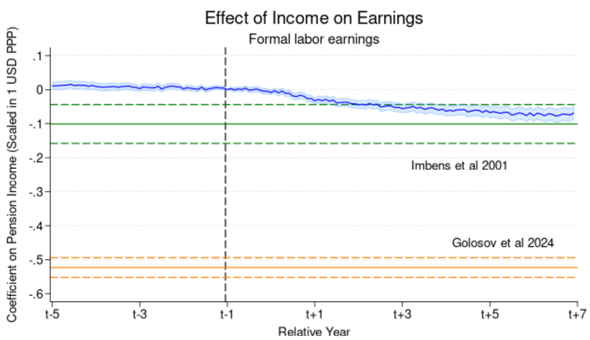

(a) Panel A: Comparison with Preferred Specifications

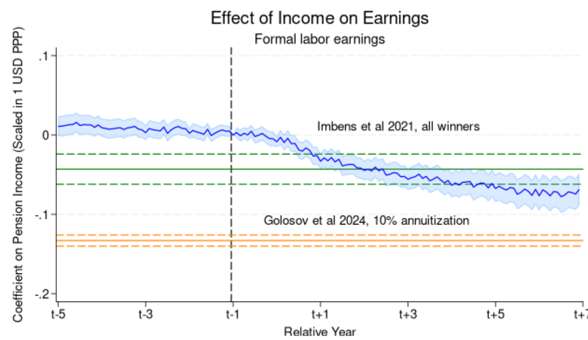

(b) Panel B: Comparison with Alternative Specifications

The blue line plots the coefficients from an event-study regression of formal labor earnings on pension income. The horizontal lines represent point estimates and 95% confidence intervals from prominent studies. Panel A compares our dynamic estimates to the preferred specifications from Imbens et al. (2001) and Golosov et al. (2024). My results are smaller than (Imbens et al., 2001) and much smaller than (Golosov et al., 2024). Panel B show how results can be reconciled, by using a larger annuitization rate in (Golosov et al., 2024) and by considering the sample of all winners in (Imbens et al., 2001). Not shown in the plot, results from (Cesarini et al., 2017) can also be reconciled by adopting a larger annuitization rate of 10% per year.

Figure 26: Effect of Pension Receipt on Formal Employment

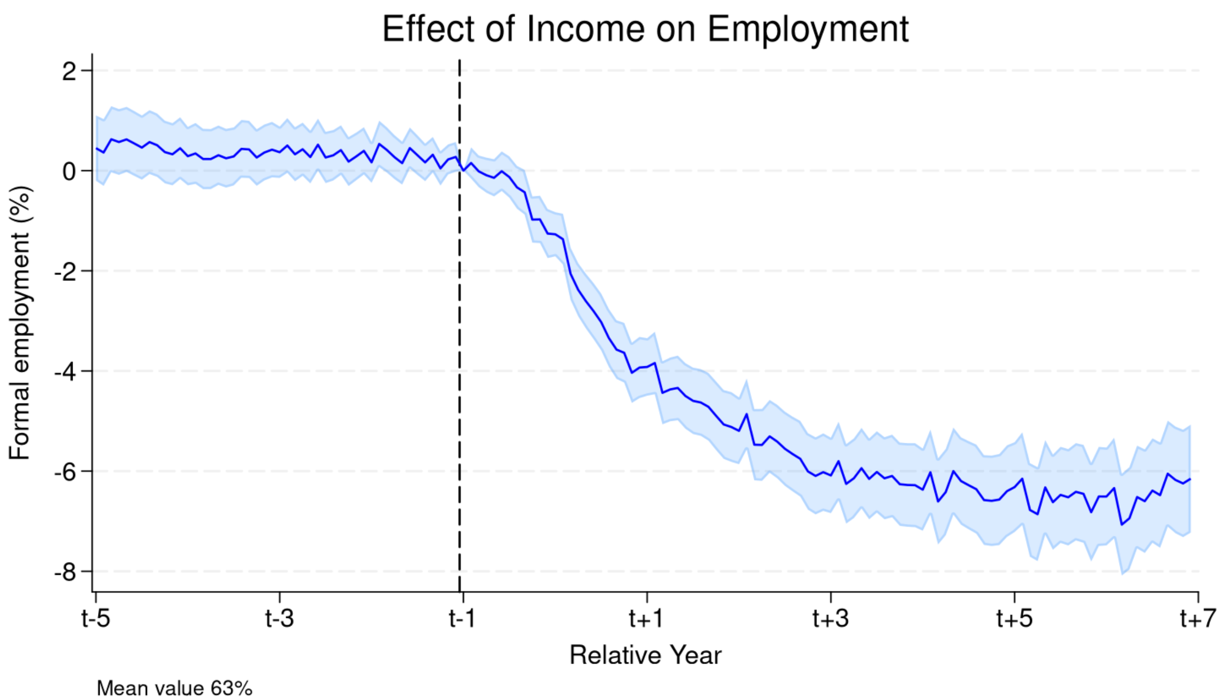

This figure shows the effect of pension receipt on the probability of formal employment. The y-axis represents the percentage point change in the likelihood of being employed. There is some anticipation effect on the 12 months before the start of the payment, and since there is some delay between death and payment, death might have occurred already for those cases. Following the start of the pension, the employment rate declines steadily, stabilizing at approximately 6 percentage points below the baseline after four years. This demonstrates a significant labor supply response on the extensive margin. The mean formal employment rate in the sample is 63%.

Figure 27: Effect of Pension Receipt on Weekly Hours Worked

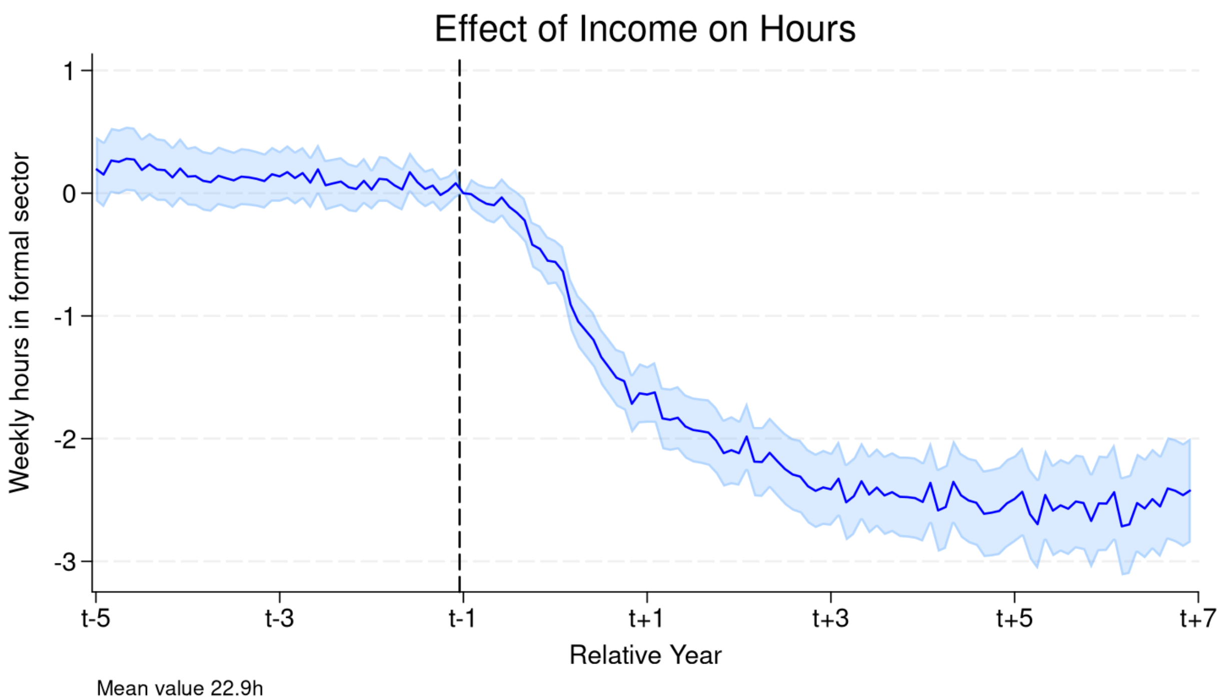

This figure shows the effect of pension receipt on weekly hours worked in the formal sector. The y-axis represents the change in hours relative to the year before the event. Hours are recorded as zero for the non-employed. There is some anticipation effect on the 12 months before the start of the payment, and since there is some delay between death and payment, death might have occurred already for those cases. Pension receipt leads to a sharp and persistent decline in hours worked, with a total reduction of approximately 2.5 hours per week after five years. This indicates no labor supply adjustment on the intensive margin. The mean weekly hours in the sample is 22.9.

Figure 28: Effect of Pension Receipt on Formal Labor Earnings

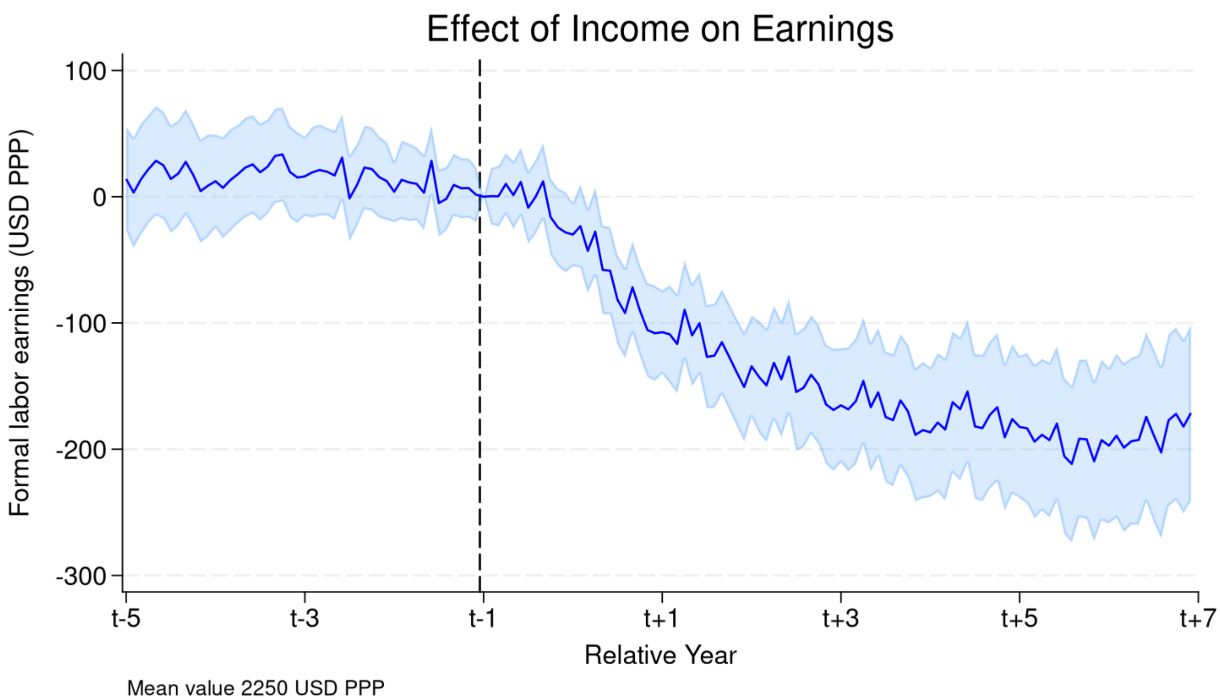

This figure plots the coefficients from an event-study regression showing the effect of pension receipt on the level of monthly formal labor earnings, measured in USD PPP. The sample consists of daughters who were employed on January 1, 2002. Earnings are recorded as zero for the non-employed. There is some anticipation effect on the 12 months before the start of the payment, and since there is some delay between death and payment, death might have occurred already for those cases. Following the start of the pension, earnings gradually decline by approximately $180 per month after seven years, relative to a sample mean of $2,250. The shaded area represents a 95% confidence interval.

Figure 29: Effect of Pension Receipt on Hiring Rate

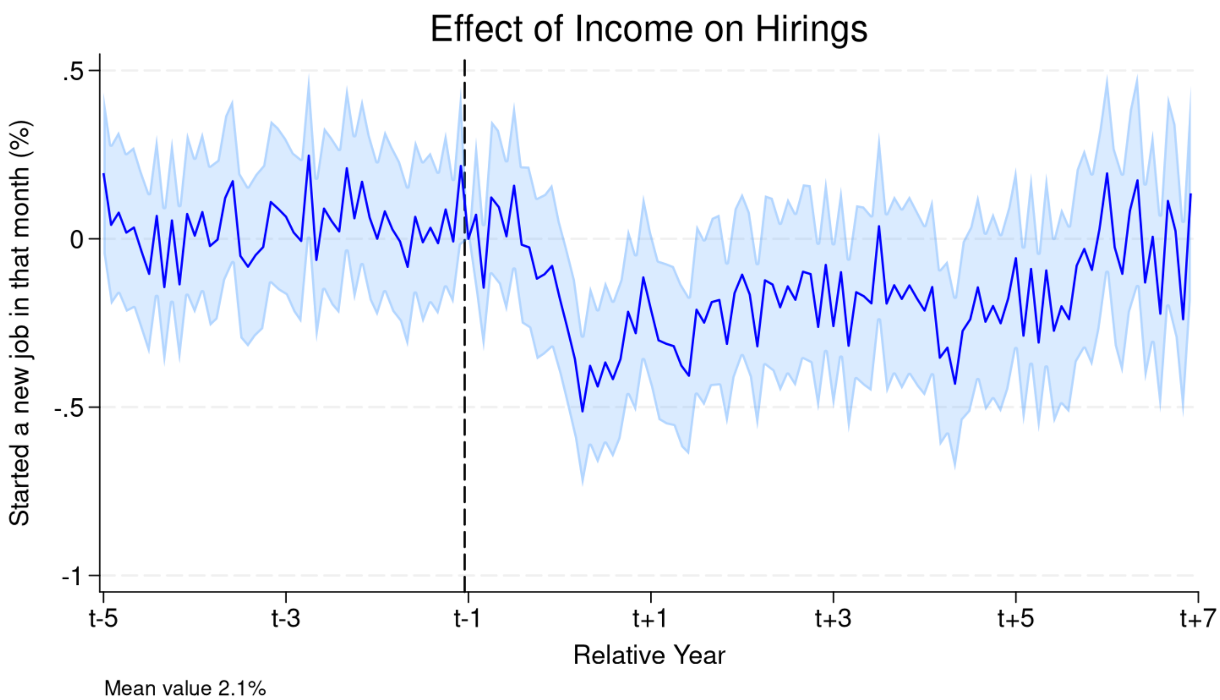

This figure plots the effect of pension receipt on the probability of starting a new job in a given month. The y-axis shows the percentage point change in the hiring rate. The mean hiring rate in the sample is 2.1% per month. After the pension begins, there is a persistent negative effect, with the hiring rate dropping by approximately 0.2-0.4 percentage points, 10% to 20% of the baseline. Although the estimates are noisy, this suggests that pension income reduces job-seeking behavior and transitions into new employment.

Figure 30: Effect of Pension Receipt on Involuntary Job Separations

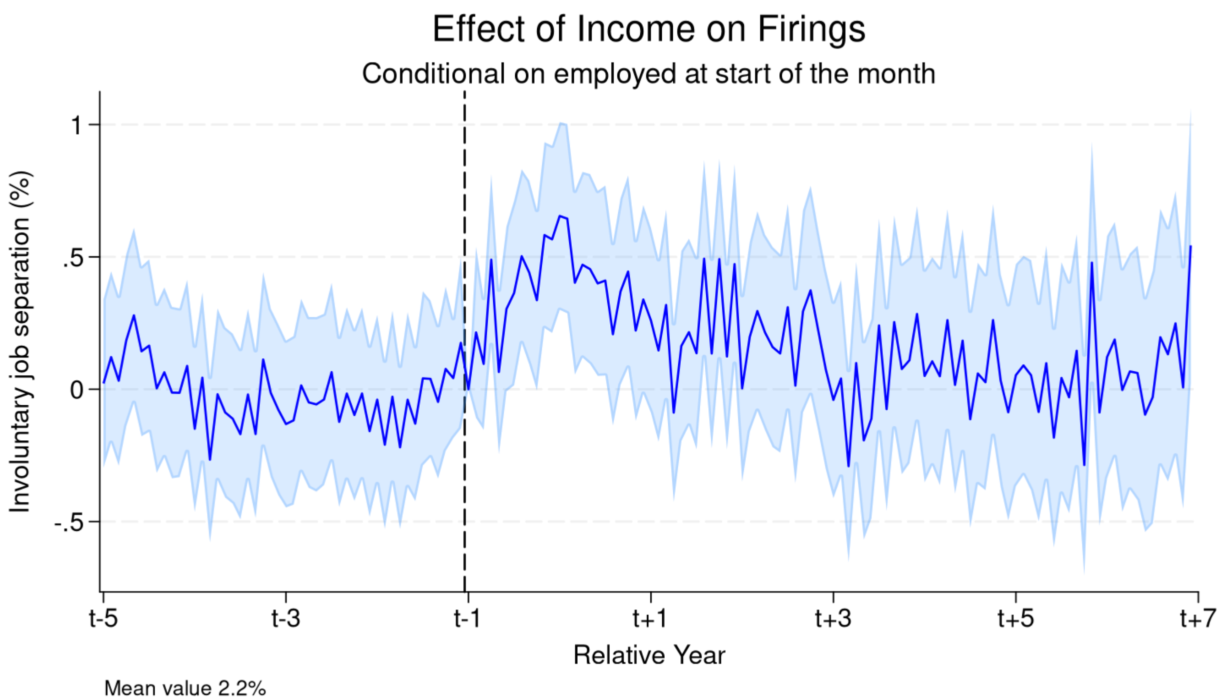

This figure shows the effect of pension receipt on the probability of involuntary job separation (firings), conditional on being employed at the start of the month. The y-axis represents the percentage point change in the firing rate. The spike in firings is large, about 30% of the baseline, but short lived, and return to baseline between 1 and 3 years after the pension started.

Figure 31: Heterogeneity of the MPE by Education Level

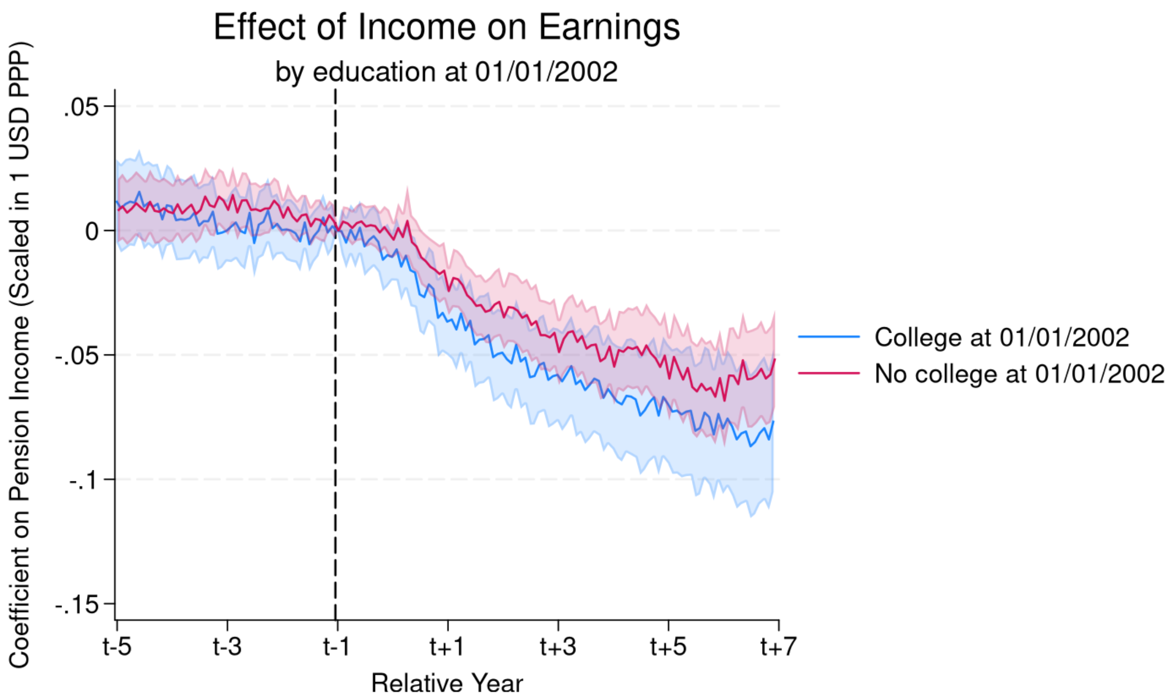

This figure plots the Marginal Propensity to Earn (MPE) separately for daughters with and without a college degree at baseline (January 1, 2002). The pre-event trends for both groups are flat and statistically indistinguishable from zero. After pension receipt, both groups reduce their labor earnings. While the point estimates for the non-college group are slightly more negative initially, the confidence intervals for the two series largely overlap, suggesting there are no statistically significant differences in the labor supply response across education levels.

Figure 32: Heterogeneity of the MPE by Baseline Wage

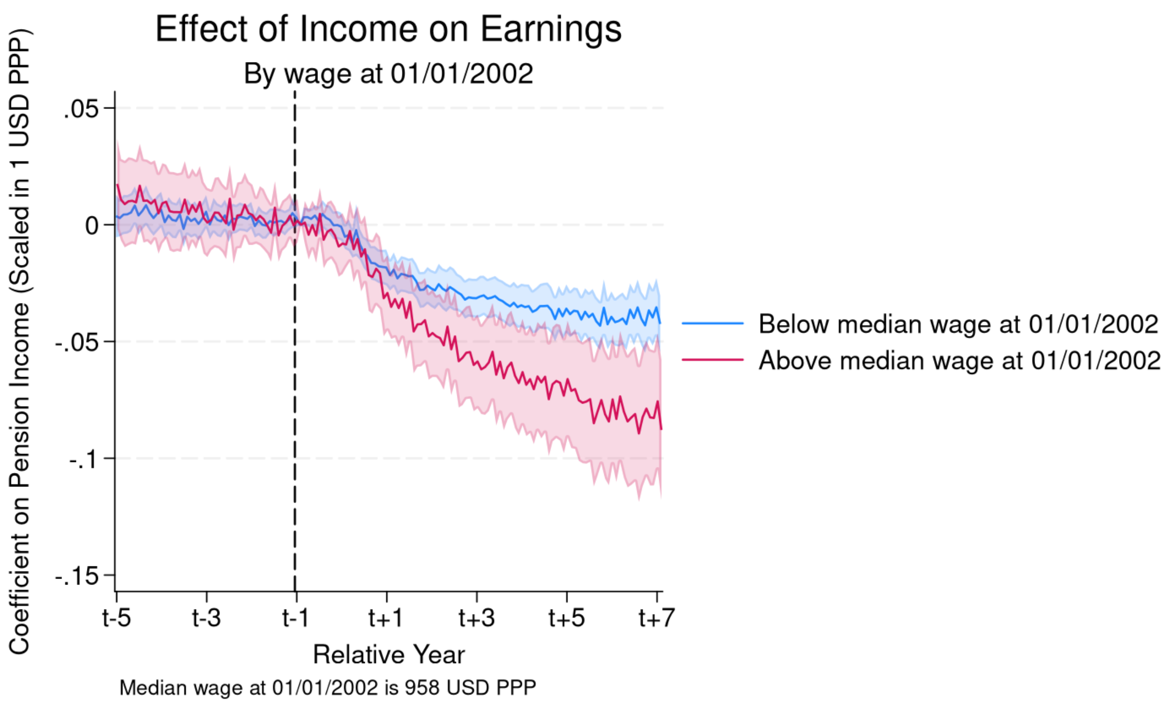

This figure plots the Marginal Propensity to Earn (MPE) separately for two groups: daughters earning above and below the median wage at baseline (January 1, 2002). The median monthly wage was $958 USD PPP. Both groups exhibit a negative labor supply response, but the effect is substantially stronger and more persistent for those with above-median baseline earnings. This suggests that higher-income individuals have a larger labor supply elasticity with respect to unearned income.

Table 3: Effect of pension payment on female birth rate

	($Age \leq 35$)	($35 < Age \leq 45$)	($Age \leq 45$)
Treatment	4.96	3.81	3.31
	(10.80)	(3.54)	(3.23)
Constant	18.50**	11.49***	13.2***
	(9.37)	(2.07)	(2.04)
Individual FE	Yes	Yes	Yes
$Year \times Month$ FE	Yes	Yes	Yes
Age FE	Yes	Yes	Yes
Observations	87,366	433,079	520,451
R-squared	0.018	0.014	0.014

The outcome is the birth rate, births per 1000 women-year. Age is defined as the age at the start date of pension. No significant changes in fertility for different ages at the beginning of the pension. The average birth rate in the sample is similar to the general population, 15.2 births per 1000 women-year.

Table 4: Effect of Pension Income on Entrepreneurship

	(1)	(2)
Dependent Variable:	Business Creation	Business Ownership
Treated	0.079	-0.123
	(0.092)	(0.178)
Constant	1.670***	6.687***
	(0.040)	(0.078)
Individual FE	Yes	Yes
Year FE	Yes	Yes
Observations	3,022,475	3,022,475
R-squared	0.05	0.06

The outcome is the number of business created or joined normalized by 1000 person-year. The panel was constructed by matching the pension administrative dataset with the business ownership administrative dataset, restricting to individuals with pension start date after 2000, and businesses created or joined after 2000.

Figure 33: Impact of Income on Business Activity Over Time

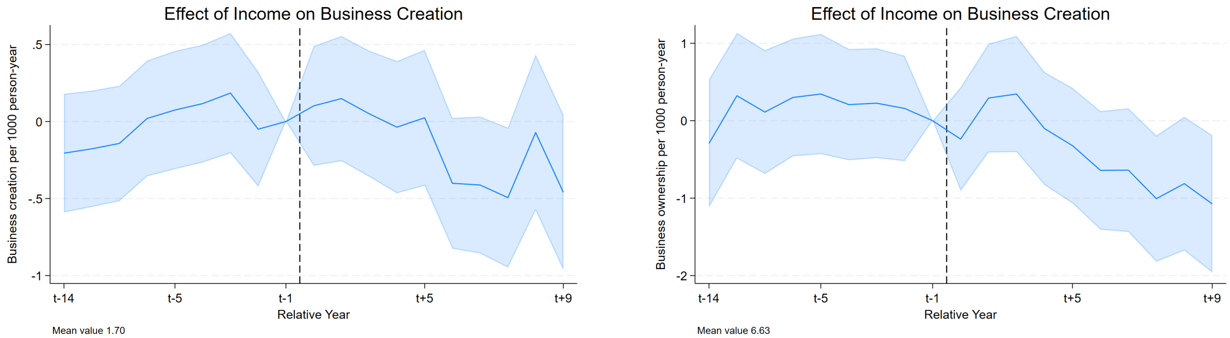

(a) Effect on Business Creation (b) Effect on Business Ownership

This figure shows the effect of pension receipt on business creation and ownership. The left panel shows the effect on new business creation per 1000 person-year (mean value 1.70), and the right panel shows the effect on business ownership per 1000 person-year (mean value 6.63). In the post-treatment period, the point estimates for both business creation and ownership remain statistically indistinguishable from zero.

Table 5: Effect of Treatment on Educational Outcomes

	Private school	College
Treatment	-0.029	-0.0003
	(0.035)	(0.104)
Constant	0.488***	0.527***
	(0.028)	(0.085)
Age FE	Yes	Yes
Observations	3,026	338
R-squared	0.100	0.228

No significant changes in educational investments, private school enrollment for children of all ages, or college enrollment for children 18+, when comparing daughters with a brother versus daughters with a sister (who need to share the future pensions) in two children households.

Figure 34: Distribution of earnings in Brazil, 2010

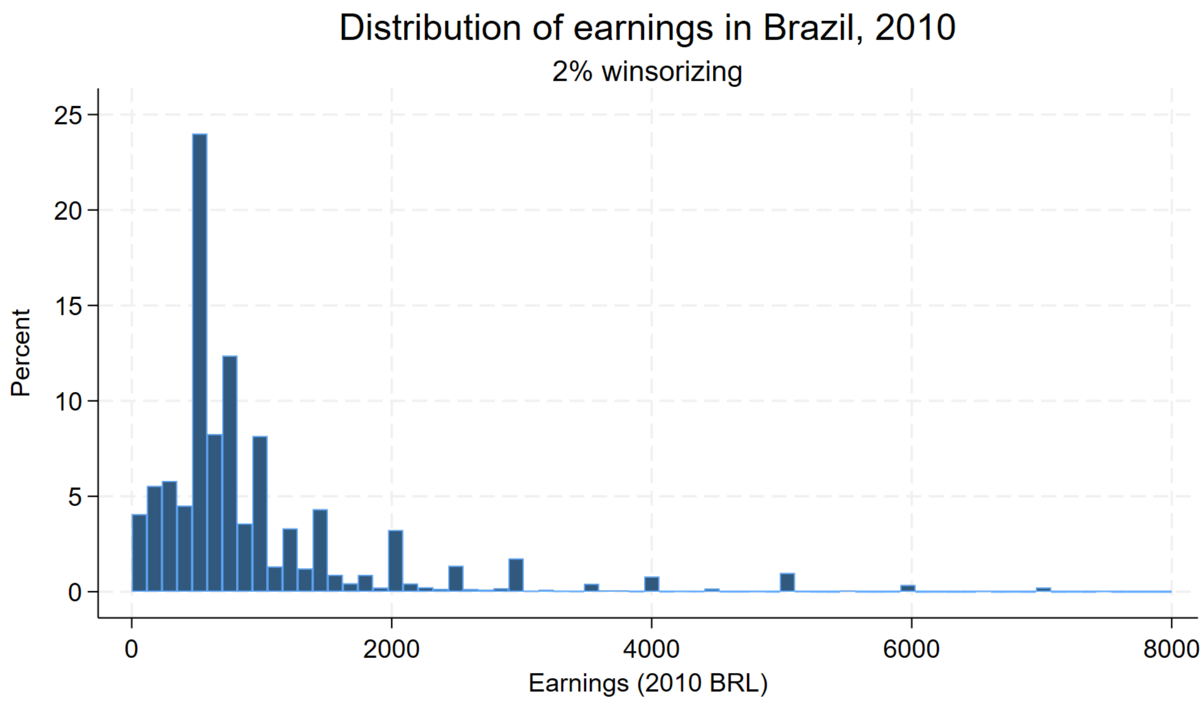

This figure displays a histogram of the earnings distribution in Brazil for 2010, which serves as a key input for the welfare analysis in Section 6. The distribution, which has been 2% winsorized to handle extreme outliers, is highly right-skewed. Values in 2010 BRL. The coefficient of variation for earnings in the winsorized data is 1.005.

Figure 35: Net welfare gain by τ, when $\varepsilon^c = 0.25$

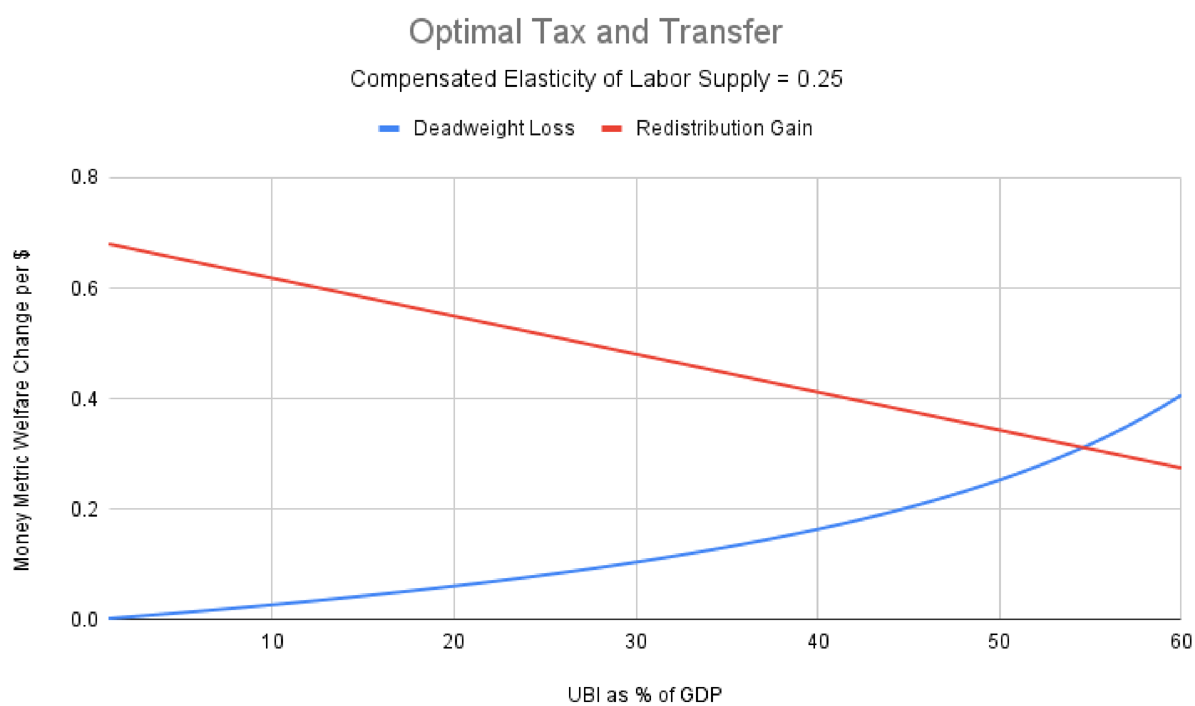

This figure plots the marginal social benefit and marginal social cost of a Universal Basic Income (UBI) program. The x-axis shows the size of the UBI as a percentage of GDP, and the y-axis shows the "Money Metric Welfare Change per \$. The analysis assumes a compensated elasticity of labor supply of 0.25. The optimal UBI size is where these two lines intersect, at approximately 54% of GDP. Before this point, the marginal benefit of redistribution exceeds the marginal cost. After this point, the marginal cost of tax distortion is greater than the marginal benefit, suggesting that further expansion of the UBI would reduce net social welfare.

Figure 36: Net welfare gain by τ, when $\varepsilon^c = 0.50$

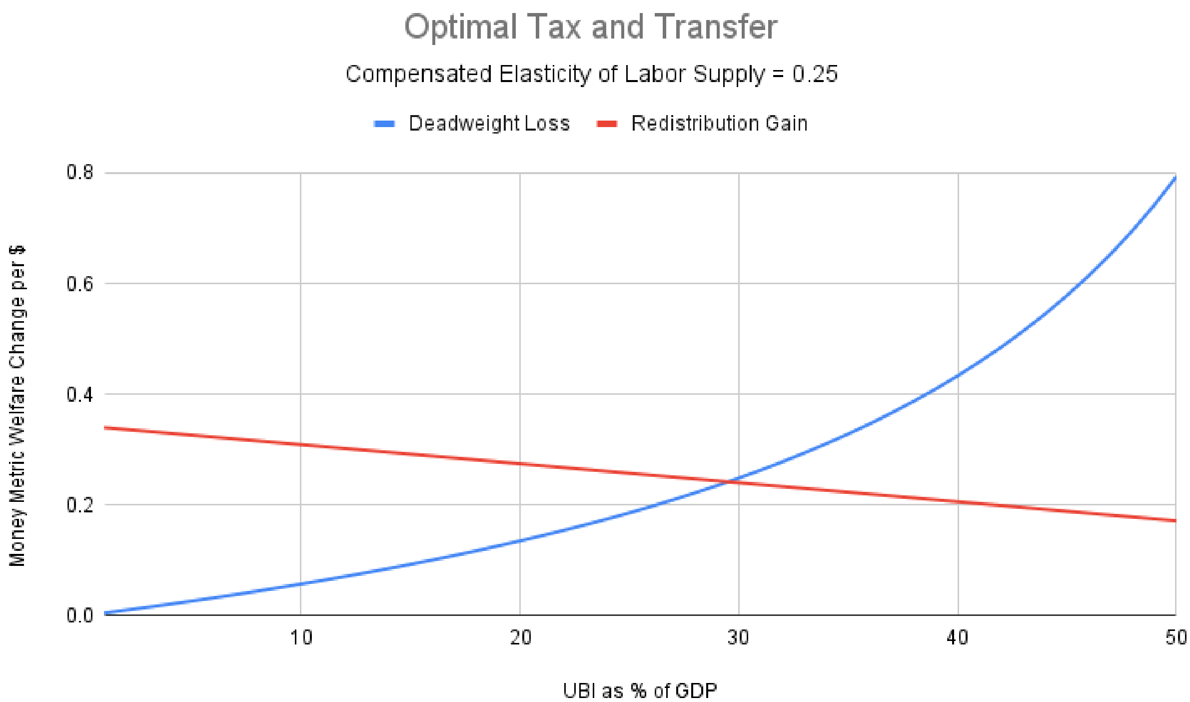

This figure plots the marginal social benefit and marginal social cost of a Universal Basic Income (UBI) program. The x-axis shows the size of the UBI as a percentage of GDP, and the y-axis shows the "Money Metric Welfare Change per $. The analysis assumes a compensated elasticity of labor supply of 0.50. The optimal UBI size is where these two lines intersect, at approximately 29% of GDP. Before this point, the marginal benefit of redistribution exceeds the marginal cost. After this point, the marginal cost of tax distortion is greater than the marginal benefit, suggesting that further expansion of the UBI would reduce net social welfare.

Figure 37: Net welfare gain by τ, when $\varepsilon^c = 0.75$

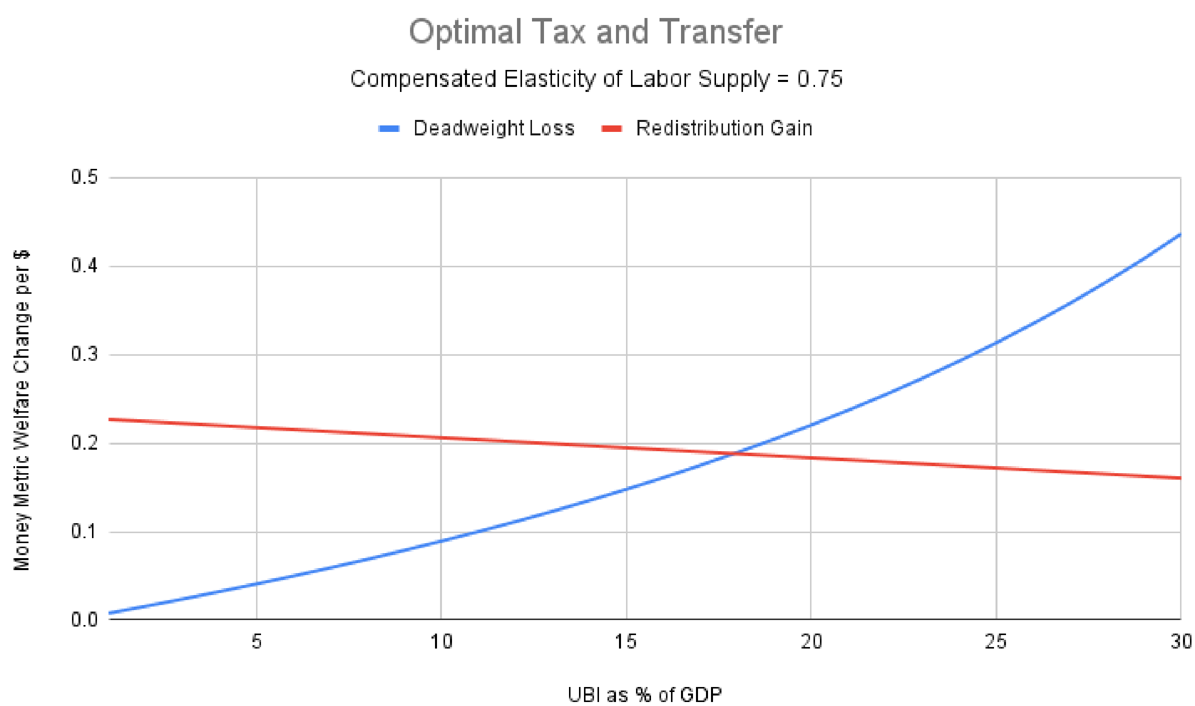

This figure plots the marginal social benefit and marginal social cost of a Universal Basic Income (UBI) program. The x-axis shows the size of the UBI as a percentage of GDP, and the y-axis shows the "Money Metric Welfare Change per \$. The analysis assumes a compensated elasticity of labor supply of 0.75. The optimal UBI size is where these two lines intersect, at approximately 18% of GDP. Before this point, the marginal benefit of redistribution exceeds the marginal cost. After this point, the marginal cost of tax distortion is greater than the marginal benefit, suggesting that further expansion of the UBI would reduce net social welfare.