Cash Transfers and Dependence in Entrepreneurial Finance

Murillo Campello*
University of Florida & NBER

Guilherme Junqueira†
University of Florida

Spyridon Lagaras‡
University of Illinois

Chi-Yang Tsou§
University of Manchester

September 8, 2026

Abstract

Cash transfer programs increasingly seek to promote business ownership among low-income individuals, yet income-based eligibility rules impose an implicit tax on the success of newly created businesses. Linking Brazil's firm registry to individual payment records from Bolsa Família, the world's largest conditional cash transfer program, we find that businesses founded by beneficiaries are 15.5 percent less likely than observably similar firms to be operating five years after founding, a gap of 7.1 percentage points that is absent at founding and widens with business age. Consistent with income misreporting to retain eligibility, reported income among beneficiary households exhibits excess mass just below the extreme-poverty threshold. To separate the post-entry effect of the means test from selection into business ownership, we estimate a dynamic occupational-choice model in which the benefit is retained when a founder ceases operations and forfeited when the business grows. Eligibility raises a new firm's annual failure risk by seven percentage points against a base rate of ten, and the means test accounts for 18 percent of the survival gap; the remainder reflects who the transfer draws into business ownership. Smoothing the eligibility notch leaves the gap nearly unchanged at substantial fiscal cost, while exempting business income from the means test removes the wedge between operating and closing and eliminates nearly all of that gap. Overall, we conclude that the excess failure of beneficiary businesses reflects the design of the means test rather than the generosity of the transfer.

JEL Codes: D15, G32, I38, J24, L26

Keywords: Cash Transfers; Business Exit; Means Testing; Structural Estimation

*Corresponding author. Warrington College of Business, University of Florida & NBER. Email: campello@ufl.edu.
†Warrington College of Business, University of Florida. Email: gjunqueira@ufl.edu.
‡Gies College of Business, University of Illinois at Urbana-Champaign. Email: lagaras2@illinois.edu.
§Alliance Manchester Business School, University of Manchester. Email: chi-yang.tsou@manchester.ac.uk.

1 Introduction

Cash transfers are among the most widely used instruments of antipoverty policy worldwide, with a large literature establishing significant effects on labor supply and entry into self-employment among low-income individuals (e.g., Bianchi and Bobba (2013)). However, the businesses that recipients start tend to survive and grow less than observably comparable businesses started by non-recipients (e.g., Campello et al. (2026)). Although previous studies provide important insights into the effects of transfers on entrepreneurship, the evidence is limited in important respects. First, existing estimates establish the magnitude of the performance differential without revealing the features of the benefit schedule that contribute to the observed gaps. Consequently, even a credibly identified gap is unable to price the impact of reforms that shift the withdrawal rate, the income base, or the transfer amount. Second, while informative, cross-sectional comparisons between beneficiary and non-beneficiary founders are unable to separate the post-entry effect of the eligibility rules on business performance from performance differences arising from the composition of founders who are drawn into business ownership as a direct result of the transfer.

To overcome these limitations, we analyze longitudinal data on founders and firms to provide novel systematic evidence on the behavior of beneficiaries around the policy thresholds, and we estimate a dynamic model of occupational choice in which the eligibility rules of a means-tested transfer program are embedded. Following cohorts of beneficiary-founded businesses from founding onward allows us to observe survival and growth at the level of the firm, alongside the benefit payments founders receive. We are thus able to estimate the exit response to the benefit schedule separately from differences in the composition of founders. In addition, embedding the schedule in the model allows us to evaluate business survival under alternative designs of the eligibility rules and to test the estimated mechanism against a policy episode in which the income test was suspended.

For our analysis, we draw insights from the Programa Bolsa Família (PBF) in Brazil, which provides a well-suited setting to study the impact of transfer design on the entrepreneurial performance of recipients for multiple reasons. First, PBF is the largest conditional cash transfer program in the world, reaching at least 17 million households, which is equal to about one quarter of Brazilian households. Second, we construct a longitudinal micro-level data set that links the national firm registry, PBF payment records, and the employer-employee administrative registry at the individual level, which allows us to follow sole proprietorships founded between 2014 and 2021 and to identify founders who were program beneficiaries. We complement our evidence with the national household survey, which reports household income and thus allows us to examine reporting behavior around the eligibility thresholds. Prior research on entrepreneurial entry has largely relied on survey evidence, which provides limited–if any–information on the survival and growth of the business itself. Third, the fixed costs of entrepreneurial entry are relatively low in Brazil. The individual microentrepreneur regime allows a

sole proprietor to register at no cost and incur only limited compliance costs, which renders entry feasible for low-income individuals. Fourth, the design of the program allows us to isolate the incentive we study. PBF withdraws the basic benefit at a notch instead of following a phasing-out schedule. A household whose reported income per capita crosses the extreme-poverty line forfeits the benefit in full, a loss equal to one third of total resources for a childless household. Crossing the higher poverty line costs only the smaller benefits tied to the number of children. The size of the benefit reduction at the extreme-poverty threshold makes the incentive to remain eligible large, and the contrast with the poverty line provides a within-program comparison against which we are able to elicit the role of a behavioral response. Finally, the income test was suspended during 2020-2021, which provides an episode that removes the incentive we study and allows us to compare the model's predictions with the observed response.

We begin by documenting the empirical patterns that guide our analysis. Comparing firms founded by beneficiaries with observably similar businesses founded by non-beneficiaries, we find a five-year survival gap of 7.1 percentage points, a 15.5 percent shortfall on a base of 45.7, that is indistinguishable from zero at founding and widens with business age. The timing of the gap matters for policy. If beneficiary firms were simply underperforming at birth, the eligibility rules are irrelevant to the subsequent failure outcomes. A gap that widens with age instead points to conditions that beneficiary founders face while the firms are in operation. Turning to the reported income of beneficiary households, we report bunching at the extreme-poverty threshold and no excess mass at the poverty threshold. The observed pattern establishes that beneficiaries respond to the eligibility rules, since misreporting concentrates where the benefit at stake is the largest and disappeared once the income test was suspended.

Two mechanisms are consistent with our key facts, importantly leading to opposite policy implications. Under the first, the eligibility rules reduce the survival of the businesses beneficiaries operate. Eligibility is assessed on reported income, and a founder who ceases operations retains the benefit while a founder whose business grows becomes ineligible. The resulting wedge distorts the exit decision. We call this mechanism program dependence. Under the second, the gap reflects selection at entry. The insurance value of the transfer reduces the wealth and ability thresholds at which entry is optimal, and recipients found businesses that would have performed poorly under any benefit schedule, a mechanism we label composition. Under program dependence the withdrawal schedule is a policy instrument. Under compositional differences the gap is not a cost the program imposes, and eliminating it would require reducing insurance to low-income households. Isolating the underlying mechanism is therefore impossible by relying on the empirical facts alone. Selection operates on entrepreneurial ability and on actual rather than reported income, neither of which is observed. Program dependence operates on the difference between the value of continuing the entrepreneurial venture and the value of returning to wage employment, and neither value is observed either.

We therefore estimate a model in which both margins are specified, namely entry selection under the insurance value of the transfer and exit under the actual benefit rules. Agents face uninsurable idiosyncratic risk, save in a single asset, and sort between wage employment and business ownership on wealth, entrepreneurial ability, and the returns to each occupation (e.g., Cagetti and De Nardi (2006)). The model incorporates a number of features specific to our institutional setting. First, the schedule is evaluated on reported income, and the scope for underreporting differs between wage employment and self-employment. A formal wage is filed with the government by the employer and can be verified against that filing, whereas self-employment profit is self-declared with no third-party counterpart and can be underreported at a convex cost. Second, the exit hazard falls in the difference between the value of keeping the business open and the value of returning to wage employment, with a single parameter κ governing that response. Setting κ equal to zero removes the exit response and leaves the pure-composition model, in which beneficiary businesses fail more only because beneficiaries enter with lower wealth and lower entrepreneurial ability. The pure-composition model is therefore a testable restriction of the model we estimate.

We estimate nine parameters by simulated method of moments on eighteen moments of the microdata, including the survival rates, the age profile of the adjusted gap, the aggregate entry and exit rates, and the excess mass below the threshold. Five of the eighteen moments are regression coefficients rather than sample averages. We therefore estimate the identical specification on data generated by the model and match the resulting coefficients to their empirical counterparts, as in indirect inference. We find that the exit response is present and precisely estimated. Eligibility for a transfer worth 8 percent of the median formal wage raises a newly-created firm's annual failure probability by about seven percentage points, against an average exit rate of 10 percent. The pure-composition model is rejected by the distance test between the two specifications, and the deterioration in fit falls on the survival dynamics rather than on the static moments.

To decompose the gap into our suggested sources, we run two experiments on the estimated model, holding all parameters at their estimated values and changing one feature at a time. The baseline five-year survival gap is 11.54 percentage points. When we set κ equal to zero and shut down the exit response, the gap falls to 0.92 percentage points, indicating that 92 percent of the gap operates through the exit decision. When instead we suspend the means test and leave κ at the original estimate, the gap falls to 9.46 percentage points, suggesting that 18 percent of the gap originates in the loss of benefits that a successful founder incurs under the means test. The two experiments differ because the value difference governing exit varies between beneficiary and non-beneficiary founders for two reasons. The first is that the means test pays a beneficiary who closes her business and withdraws the benefit from one whose business succeeds, which lowers the value of staying open relative to closing. The second is the difference in the performance gap generated by selection at entry, since beneficiary founders enter with lower ability and lower wealth and their businesses are worth less relative to the wage option.

Setting κ to zero eliminates the exit response to both, while suspending the means test removes only the first. Together our experiments decompose the gap into three parts. Eight percent survives when the exit response is shut down and is therefore composition that requires no behavioral response at all. Eighteen percent originates in the income test. The remaining 74 percent is composition at entry amplified through the exit margin.

Our results carry important implications for the design of transfer programs. First, although nearly all of the gap operates through the exit decision, only a fifth of it originates in the means test itself. This places an upper bound of roughly one fifth on the improvement in survival that any reform of the means test could deliver. Second, the distortion is not a liquidity effect. Transfers in the model relax every recipient's budget constraint, and that support remains in place when we suspend the means test. Consistent with this, performance gaps narrow but do not close in municipalities where the same nominal transfer buys more (e.g., Campello et al. (2026).)

The mechanism requires that founders still stand to lose the benefit after they start a business. Eligibility is assessed against a registry that households update at intervals rather than continuously, and 62 percent of beneficiary founders still receive the transfer one year after founding and 35 percent after five, while only 3.3 percent have household income per capita below the extreme-poverty line while their business operates. Payments therefore continue well after income has risen above the eligibility line. The estimation does not target benefit receipt, and the model reproduces the observed profile.

The model's central prediction can be checked against a policy episode. During 2020–21 Brazil temporarily replaced the means-tested payments with Auxílio Emergencial, a transfer that did not depend on reported income near the eligibility thresholds. The estimated model makes two predictions for this episode. Bunching should disappear, and it does: the excess mass at the extreme-poverty threshold collapses to zero in the 2020–21 survey rounds. And 18 percent of the beneficiary exit gap should close, with a range of 15.9 to 18.2 percent across specifications, while the remainder, being composition, survives the reform. The survival series for the episode is under construction by the companion team, and the comparison will be reported when it arrives.

In sum, a means test has two separate parts: which income counts against eligibility, and how quickly the benefit is withdrawn as that income rises. Two counterfactual experiments on the estimated model, each changing one part and nothing else, trace the survival distortion to the first. Withdrawing the benefit gradually instead of all at once accomplishes almost nothing: cutting the withdrawal rate to a quarter of its value moves the five-year gap from 11.54 to 11.46 percentage points while raising program spending by 125 percent. The exit decision responds to the difference between keeping the benefit after closing and losing it by growing, and a gradual withdrawal leaves that difference in place. Exempting business income from the test removes that loss directly, in the spirit of the earned-income disregards used in US and UK welfare programs, holding the benefit level and the targeting fixed.[1] Section 6.3

[1] In the United States, AFDC and its successor TANF disregard a portion of earnings when benefits are computed, and

reports both experiments in full.

Our paper contributes to several strands of literature. We first add to the empirical literature on how transfers shape entrepreneurship among low-income households and how social insurance affects entrepreneurial entry among individuals facing downside risk. Within the transfer literature, prior studies examine three main margins. On entry, cash transfers and start-up grants can induce self-employment by relaxing financing constraints and increasing individuals' capacity to bear risk (Bianchi and Bobba 2013; Blattman et al. 2014, 2016; de Mel et al. 2014). On productive investment and returns, transfers and capital grants can increase investment and raise the income or profits of existing farms and microenterprises (de Mel et al. 2008; Gertler et al. 2012; Macours et al. 2022; Ghosh and Vats 2022). On occupational mobility, multifaceted programs combining productive-asset transfers with consumption support, training, and coaching can move poor households toward self-employment and higher-return occupations (Banerjee et al. 2015; Bandiera et al. 2017; Balboni et al. 2022).

A related literature shows that social insurance design also shapes entrepreneurial entry: unemployment benefits retained during business formation increase firm creation without reducing entrant quality (Hombert et al. 2020), job-protected leave encourages entrepreneurial experimentation by preserving founders' option to return to wage employment (Gottlieb et al. 2022), and unemployment benefits withdrawn against business income can discourage unincorporated business formation and alter the composition of entrants (Xu 2022). In the context of Bolsa Família, Gerard et al. (2024) trace the program's effects through Brazil's formal labor market. We instead study the post-entry margin, examining how the design of a means-tested transfer program affects the survival of recipient-founded businesses and distinguishing selection at entry from the post-entry incentives created by its eligibility rules.

Our work also adds to the theoretical literature on households facing uninsurable idiosyncratic risk, who save in a single asset under a borrowing limit as in Huggett (1993) and Aiyagari (1994). One branch of this literature puts an occupational choice inside the household problem: Evans and Jovanovic (1989) and Holtz-Eakin et al. (1994) establish that liquidity constraints shape entry into entrepreneurship and survival, Hurst and Lusardi (2004) question how tightly wealth binds, and Quadrini (2000) and Cagetti and De Nardi (2006) develop the quantitative version in which the occupational margin shapes saving and the wealth distribution. Another branch puts a means-tested program inside it: Hubbard et al. (1995) show that asset-tested welfare discourages saving by low-income households, and Low and Pistaferri (2015) estimate the incentive and insurance trade-off of disability insurance on the labor-supply margin. In the first branch, entry and exit respond to wealth and productivity alone; in the second, the means test moves the saving and labor supply of households that never run a business. We give the household of the second branch the occupational choice of the first, and that changes the the-

SSI excludes the first $65 of monthly earnings plus half of the remainder. In the United Kingdom, Universal Credit applies a work allowance before benefits are withdrawn, and newly self-employed claimants are exempt from its minimum income floor during a twelve-month start-up period, the closest existing analogue to the exemption we study. Moffitt (1992) surveys the incentive effects of provisions of this kind.

ory. Because reported income differs across occupations, the same schedule taxes the two occupations differently; because the benefit is kept on the fallback and lost in success, the transfer drives a wedge between the value of operating and the value of returning to work. The means test thereby acts on three margins it reaches in neither branch: whether to found a business, how much capital to put into it, and whether to keep it alive.

The financing side of our model builds on the macroeconomics literature on credit market frictions, in which lenders advance funds only against pledgeable capital because contracts are imperfectly enforced (e.g., Kiyotaki and Moore (1997); Gertler and Kiyotaki (2010); Jermann and Quadrini (2012)); Buera et al. (2015) review its macro-development branch. The applications closest to ours show that constraints of this form hold back wealth-poor entrepreneurs and misallocate capital across producers (e.g., Buera et al. (2011); Buera and Shin (2013); Midrigan and Xu (2014)), and Moll (2014) shows that wealth-proportional borrowing limits are the tractable reduced form of the class. We take the standard setup as is: the entrepreneur pledges a fraction of installed capital, finances the rest from net worth, and the constraint binds. The transfer then enters the entrepreneur's financing directly, as part of the net worth that funds investment.

A further related strand is the structural estimation of dynamic models in corporate finance, including Hennessy and Whited (2005), Hennessy and Whited (2007), Nikolov and Whited (2014), Li et al. (2016), and Warusawitharana and Whited (2016). Our paper departs from the prior literature in one important way. The object we estimate is not a technology or a financing friction but a government program's eligibility rules, and the model confronts not the financial statements of listed firms but the linked administrative records of low-income households and the businesses they found. In this regard, our paper is closest to Low and Pistaferri (2015), who estimate a dynamic model of a means-tested program on household data; their margins are labor supply and program participation, ours are the founding, the scale, and the survival of the recipients' businesses.

Our empirical strategy relates to the literature that uses bunching at kinks and notches to study behavioral responses to nonlinear tax and transfer schedules (Saez 2010; Chetty et al. 2011; Kleven and Waseem 2013), surveyed by Kleven (2016), and to the broader literature on the incentive effects of means-tested programs surveyed by Moffitt (1992). Most closely, Bergstrom et al. (2025) use changes in Bolsa Família's eligibility notch to derive welfare bounds. We use bunching for a different purpose: rather than recover a behavioral elasticity or a sufficient statistic for welfare analysis, we use the location and timing of excess mass as evidence on the reported-income response embedded in our model and as a moment that disciplines its estimation. Bunching appears only at the threshold where crossing entails a large discrete benefit loss and disappears when the income test is suspended, distinguishing a program-induced reporting response from stable features of the income distribution.

Closest to this paper is the companion study by Campello et al. (2026), which provides the empirical foundation for our analysis. The companion causally identifies the effect of cash transfers on

entrepreneurial entry, documents lifecycle performance gaps and post-failure employment outcomes, and addresses selection into cash-transfer entrepreneurship using coarsened exact matching, granular founder- and business-level fixed effects, and sensitivity analyses for selection on unobservables. It further presents program-tenure patterns consistent with a dependence channel. These empirical strategies establish that the performance gaps persist among closely comparable founders and narrow the set of plausible explanations, but they do not structurally decompose the adjusted survival gap into remaining differences in founder composition and post-entry responses to the benefit schedule. Building on this evidence, the present paper develops and estimates a dynamic model that explicitly nests both channels, uses the companion's empirical findings to discipline the model, quantifies their respective contributions to the survival gap, and evaluates counterfactual reforms to the program's eligibility rules.

The rest of the paper is organized as follows. In Section 2, we describe the institutional setting of PBF, the administrative data, and the estimation sample. Section 3 presents the empirical facts. Section 4 develops the model. Section 5 reports the estimation. Section 6 presents the decomposition, the counterfactuals, and the out-of-sample test. Section 7 concludes. Details on the institutional background and data construction are provided in the Internet Appendix. The Internet Appendix also contains a case study of the eligibility notch, additional bunching evidence and robustness checks, complete proofs of the two propositions, and details on the model solution and the estimation procedure.

2 Institutional Background and Data

In this section, we describe the institutions and the data the analysis builds on. Three features of the Brazilian setting do the work in this paper. The transfer is withdrawn at a threshold rather than phased out, the income that determines eligibility is self-reported for the self-employed and third-party reported for wage earners, and eligibility is assessed against a registry that households update periodically rather than continuously. Sections 2.1 to 2.4 take these in turn, together with the business forms available to a low-income founder. Section 2.5 describes the data.

2.1 Bolsa Família

Bolsa Família is the world's largest conditional cash transfer program. It reaches more than 17 million families, about one quarter of Brazilian households, at an average payment of R$170 ($47) per month, or 19% of the minimum wage. The program costs the federal government over R$30 billion ($8 billion) a year, roughly 0.5% of GDP over our sample period.[2]

To qualify, a household registers with the program and reports monthly income per capita below

[2]Between 2014 and 2021 the average exchange rate was R$3.63 to the dollar and the monthly minimum wage was R$900. We give dollar conversions for policy parameters at that rate and leave other values in reais.

a threshold. Before 2023 the program operated two tiers. Taking 2018 as the midpoint of our sample, the *extreme-poverty* threshold was R$89 per capita per month ($25, or 10% of the minimum wage) and the *poverty* threshold was R$178 ($49, or 20%), maintained at exactly twice the extreme-poverty level. Households below the extreme-poverty line received a *basic* benefit, a lump sum independent of family composition, together with *variable* benefits tied to children. Households between the two lines received the variable benefits only. Enrollment is conditional on school attendance and health check-ups, and households must keep that information current, including their income, since errors or failure to update can suspend the benefit.

2.2 The eligibility cliff

Crossing the extreme-poverty threshold removes the entire basic benefit at once. In 2018 the basic benefit was R$89 per month and the variable benefits were R$41 per child aged 0–15 and R$48 per adolescent aged 16–17. A household with two children and per capita income of R$88 collected R$171 in total. The same household at R$90, two reais higher, collected R$82.

The size of that loss is easiest to see in the ratio of post-transfer to pre-transfer income. Deep in extreme poverty the transfer more than doubles a household's resources. At the extreme-poverty threshold the ratio falls from 1.5 to 1.0 for a childless household, which is a loss of one third of total resources triggered by a two-real increase in per capita income. Childless households are 43% of the Brazilian total. Households with children lose less, from 1.5 to 1.15 with one child and to 1.23 with two, but the loss is still discrete.

The poverty threshold is different in kind, not merely in degree. A childless household sees no change in the ratio at all, from 1.0 to 1.0, and households with one or two children see it fall only from 1.08 and 1.11 to 1.0. Withdrawal there is further softened by a two-year protection rule (*regra de proteção*) in force since 2010, under which a household whose income rises above the poverty threshold may keep the benefit for up to two years, provided income stays below half the minimum wage. No such protection exists at the extreme-poverty threshold.

The asymmetry has a sharp implication for where income manipulation should appear. Kleven and Waseem (2013) show that a notch creates a strictly dominated range above the threshold whose width is proportional to the size of the benefit loss. Applying that argument here, the dominated range at the extreme-poverty threshold is approximately R$89 for a childless household, the full basic benefit, while at the poverty threshold a childless household faces no dominated range at all. Manipulation should concentrate at the extreme-poverty line and be absent at the poverty line, and Section 3 shows that it is.

Two later changes bracket our sample. During 2020–21 the temporary Auxílio Emergencial program replaced Bolsa Família with payments of up to R$600 ($165) per household per month, which

suspended the threshold-based structure altogether. In November 2021 Auxílio Brasil introduced a minimum benefit floor of R$400 ($110) per family regardless of composition, made permanent in 2022, which removed the discrete loss at the extreme-poverty threshold. We return to the first of these in Section 6.2, where it serves as an out-of-sample test.

2.3 Who reports income, and how often

Two institutional details determine how the threshold bites, and both enter the model directly.

Reporting is asymmetric across occupations. A formal employee's wage is filed with the government by the employer, through the mandatory annual employment registry (RAIS), and the employee cannot understate it. Self-employment income is not filed by anyone else. A sole proprietor operating under the entry-level tax regime submits a single simplified annual return (DASN-SIMEI), may complete it without an accountant, and faces no third-party statement against which it is checked. The scope for understating income is therefore a feature of the occupation rather than of the individual, and it exists on the entrepreneurial side of the occupational choice and not on the wage side.

Eligibility runs on registered income, not current income. Entitlement is not verified continuously. A household qualifies by registering its income with the program, and that record is revisited at intervals rather than month by month. A household whose income rises therefore keeps receiving the transfer until its record is next reviewed, and receipt persists after the income that justified it has gone.

The gap between the two is large, and it is visible in our data. Among beneficiaries who found a business, 62% are still drawing the transfer a year later, 47% after three years and 35% after five.[3] Measured by income instead, only 3.3% have per-capita household income below the extreme-poverty line while the business operates, and 10.7% are below the poverty line. More than eight in ten of those still being paid are, by income, no longer entitled to be. Withdrawal in this program is administrative and gradual rather than immediate and mechanical, which matters for *when* the distortion we study begins to bind and not only for whether it does.

2.4 Sole proprietorships in Brazil

New business registrations in Brazil grew from two million a year in 2014 to over four million in 2021, and sole proprietorships account for 77% of new establishments. Limited liability companies are second at 14%. The distinction that matters here is that only sole proprietorships may use the individual microentrepreneur regime (MEI), which is the entry point to entrepreneurship for low-income founders.

[3]These are shares among founders observed at each age. The balanced-cohort construction of the same profile reads 76.8, 52.2, and 35.3 percent at ages one, three, and five; the conditioning behind each construction is stated with Table 6.

An MEI business may earn up to R$81,000 ($22,000) a year and employ one worker besides the founder. It pays a fixed R$66.60 ($18.35) a month, less than 1% of the revenue ceiling, and that payment buys access to the public pension system, including retirement, disability and maternity leave at the minimum wage. Compliance is the annual return described above.

A business that outgrows the ceiling moves to the microenterprise regime (ME), which permits revenue up to R$360,000 ($100,000) and nine employees in commerce and services or nineteen in manufacturing. Taxation there is progressive under SIMPLES Nacional, from 4% to 15.5% of gross revenue, and compliance requires an accountant and monthly filings. The step from MEI to ME is the observable margin of business growth in this setting, and we use it as such in Section 3.

2.5 Data and sample

The analysis links four Brazilian administrative sources and one household survey. We describe each briefly, then set out the sample and the variables the model is estimated against.

Sources. The *Cadastro Nacional da Pessoa Jurídica* (CNPJ) is the national firm registry maintained by the Ministry of Finance. It covers every registered business, active and inactive, and records the founding date, the legal form, the business activity, the initial equity capital (*capital social*), and a full set of company events, including whether a business remains active and whether it has moved to a larger legal form. Survival and the MEI-to-ME transition are both read from these events. Bolsa Família payment records from the Ministry of Social Development give the exact amount received by beneficiary and month for 27.3 million beneficiaries between 2013 and 2021, identified by social identification number (NIS). The *Relação Anual de Informações Sociais* (RAIS) is the mandatory annual employment registry and covers every formal employment relationship in Brazil, with monthly salary, job title, payment type, and worker demographics. Household income distributions come from PNAD Contínua, the continuous national household survey run by IBGE, which follows roughly 211,000 households a year and records Bolsa Família participation alongside every income source.

Linking founders to businesses. Neither the firm registry nor the payment records identify a sole proprietorship's founder as an individual. The link is built from the individual tax identifier (CPF), which is almost always embedded in the legal name of a sole proprietorship in the CNPJ registry and which we verify against a major credit bureau. This recovers founders for 16.3 million of the 18.8 million sole proprietorships started between 2014 and 2021, a match rate of 87%. Employment records are then matched to businesses through the CPF and to payment records through the NIS, which yields a complete employment history, a demographic profile, and a Bolsa Família participation history for each founder.

Sample. We keep sole proprietorships, which are the entry point to entrepreneurship for low-income founders and the only form that may use the MEI regime. We keep only the *first* business an individual starts, which keeps experience from an earlier venture out of the estimates. We keep only founders with formal employment in the twelve months before founding, which gives every founder in the sample an observed pre-entry wage and makes the treated and comparison groups alike in labor-market attachment. Bolsa Família participation is tracked from January 2013, businesses are founded between January 2014 and December 2021, and outcomes run to December 2024. The resulting sample is 2,296,485 sole proprietorships, of which 76,219 were founded by Bolsa Família beneficiaries.

Table 1 summarizes the sample. Three contrasts matter for everything that follows. Beneficiary founders are overwhelmingly female (86 against 39 percent), earn a little more than half the pre-entry wage of other founders (R$488 against R$861 monthly), and start with about two-thirds of the initial equity. Their businesses concentrate in service and sales, survive less, grow less, and run into tax trouble more. The fixed-effects strategy of Section 3 absorbs these compositional differences, and the model of Section 4 takes the pre-entry wage and initial equity distributions as inputs, which is how the composition the table displays enters the structural estimation.

[Place Table 1 about here]

Beneficiary-founded businesses. A business is *beneficiary-founded* if its founder was enrolled in Bolsa Família when the business was registered, or had left the program within the preceding twelve months. In 89% of cases the founder is still enrolled at founding. The tag is assigned at founding and never revised. A founder who earns their way out of eligibility while the business operates stays in the beneficiary-founded group. Fixing the tag at founding keeps the comparison with other businesses free of selection on anything the business itself does, and it is the definition the structural estimation targets.

Estimation targets. Five objects carry the identification, and each has a direct counterpart in the data. Business survival by age and by beneficiary status comes from the CNPJ company events. Growth is the MEI-to-ME transition from the same source. The founder's pre-entry wage, which determines eligibility in the fallback state, comes from RAIS, as does the wage a founder earns on returning to employment after a business closes. The initial equity capital recorded at registration disciplines the wealth an entrant brings in, which the model would otherwise be free to choose. Bunching of reported income below the extreme-poverty threshold comes from PNAD Contínua. Section 5 sets out the moments in full.

3 Empirical Facts

In this section, we present the empirical facts that discipline the model. The companion paper (Campello et al. 2026) documents how Bolsa Família shapes entrepreneurial entry and performance; this paper formalizes the mechanisms at play and quantifies their welfare implications, which the companion is silent about. This section summarizes the companion facts the model is asked to explain, and then presents the evidence this paper adds, income concealment tracks the program's design, and stops tracking it when the design changes.

3.1 Facts established in the companion

Three findings from Campello et al. (2026) discipline everything that follows. First, transfers create entrepreneurs. Instrumenting local program participation with the share of households below the manipulation-free poverty threshold, a ten percentage point increase in participation raises transitions into sole proprietorship by 16 percent from non-employment and 8 percent from wage employment, and the beneficiary share of entrepreneurs rises 46 percent. Second, beneficiary-founded businesses close more, and the shortfall grows with age. In the specification with industry, initial equity, and founding-year fixed effects, the survival gap is 0.6 percentage points in the first year and 7.1 by year five, a 15.5 percent shortfall on a base of 45.7. Third, the survivors stay small, beneficiary businesses are 12 percent less likely, in the same specification, to make the MEI-to-ME transition, the legal growth milestone in this setting.

The survival system is not only motivation, it is the estimation target. The beneficiary and non-beneficiary survival curves and the adjusted gap profile enter the moment vector directly, and Table 4 in Section 5 reports the series alongside their model counterparts. The moment vector targets the no-founder-controls variant of the same regression, the estimand the model-side regression replicates exactly; its profile runs from -0.7 percentage points in year one to 6.8 by year five, and the first-year coefficient is indistinguishable from zero in both variants. The age profile matters throughout, the gap is nearly absent at birth and accumulates as the business ages, and Section 4 argues this profile is where a dependence margin and pure founder composition can be told apart.

3.2 Concealment tracks program design

Section 2.2 showed that the two eligibility cutoffs carry very different stakes, the entire basic benefit at the extreme-poverty threshold against modest child-linked benefits, cushioned by the protection rule, at the poverty threshold. Reported income responds exactly where the stakes are large. Figure 1 plots the distribution of pre-transfer income of beneficiary households around each cutoff, with a visible spike just below the extreme-poverty threshold and a corresponding deficit above it, and no comparable

pattern at the poverty threshold.

[Place Figure 1 about here]

Table 2 quantifies the contrast using the excess-mass estimator of Saez (2010) and Chetty et al. (2011), excess mass of 0.23 at the extreme-poverty threshold ($t = 2.94$) against a statistically flat 0.03 at the poverty threshold ($t = 1.08$). Beneficiaries do not misreport indiscriminately. They conceal income only where the program's design makes honesty expensive, and the model's concealment technology in Section 4 is built to reproduce this asymmetry.

[Place Table 2 about here]

3.3 Two placebo checks

The design interpretation is testable, and it passes twice. During 2020–21 the Auxílio Emergencial program replaced means-tested payments with a transfer that did not depend on reported income near the threshold. Re-estimating the same bunching specification on the 2020–21 survey rounds, the excess mass at the extreme-poverty threshold collapses to statistical zero, with point estimates between −0.05 and 0.14 across windows, all insignificant. The same households, facing no income test, stopped bunching. Households outside the program provide the second placebo. Non-beneficiaries face no incentive at either cutoff, and their income distribution shows no excess mass at the thresholds in any period; its only visible irregularities sit at round incomes near R$100, well above the policy cutoff and unrelated to it. Section III of the Internet Appendix reports both sets of estimates in full. Concealment is a response to the transfer schedule rather than a fixed trait of the population, which is the premise of the counterfactuals in Section 6.

4 The Model

In this section we develop the model we use to identify and quantify the dependency channel. The workhorse is deliberately standard. Agents face uninsurable idiosyncratic risk and save in a single asset, as in Huggett (1993) and Aiyagari (1994), and they sort between wage work and business ownership on wealth, ability, and the returns to each occupation, as in Lucas (1978), Quadrini (2000), and Cagetti and De Nardi (2006), with businesses subject to the entry and exit dynamics of Hopenhayn (1992). Keeping the chassis standard means that every result below traces to what we add, not to an exotic preference or technology.

We add three ingredients, each dictated by the institution. First, a means-tested transfer replicates Bolsa Família's two-threshold schedule, a basic benefit below the extreme-poverty line and a

variable benefit up to the poverty line. The schedule tests reported income, and the two occupations report differently. A worker's wage is filed by the employer and cannot be understated. An entrepreneur's profit is self-reported. Second, an entrepreneur can therefore conceal part of that profit, at a cost. Third, exit is a hazard that rises when the business is worth little relative to the worker fallback. The transfer raises the fallback of an eligible founder, which thins the value of staying in business and pushes exit up. This third ingredient is what carries the dependency channel.

4.1 Environment

Time is infinite and discrete, and the economy is populated by a unit continuum of infinitely-lived agents who discount the future at rate β. Preferences are common to everyone. An agent ranks consumption streams by solving:

$$\max_{\{c_t\}_{t\geq 0}} \quad \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \, u(c_t), \qquad u(c) = \frac{c^{1-\sigma}}{1-\sigma}, \tag{1}$$

where σ is the coefficient of relative risk aversion.

Each period an agent occupies one of two roles, worker or entrepreneur, and is described by liquid wealth a, a wage-earning skill e, a permanent entrepreneurial quality q, and, when running a business, a transitory productivity z. The occupational structure and the collateral constraint that goes with it follow Buera et al. (2011) and Buera and Shin (2013). We add a means-tested transfer that the two roles face on different terms, and an exit margin that responds to it.

In both roles the agent solves the same problem in equation (1), choosing consumption together with the saving or investment that goes with it. The two cases differ only in the law of motion of wealth. A worker holds their wealth in the risk-free asset, which accumulates at the gross rate R_f. An entrepreneur's wealth is the net worth invested in the business, and it accumulates out of profit net of depreciation and debt. Both laws of motion are spelled out in the sections that follow.

Prices are taken as given. The agents we model are a small segment of Brazil's labor and credit markets, and their hiring and borrowing are far too small a share of either to move the going rates. The wage per efficiency unit w and the gross risk-free rate R_f are therefore treated as exogenous, in the spirit of a small open economy. Section 4.7 states this closure precisely and is explicit about which market-clearing conditions are imposed and which are not.

The transfer is paid for by a lump-sum tax τ, which leaves the occupational and investment margins undistorted. Throughout, primes denote next-period values.

4.2 The Means-Tested Transfer

The transfer is the source of the mechanism, and we describe it before the agents' problems. Let $\tilde{y}$ denote income *as reported* to the transfer authority. The schedule has two thresholds and is given by:

$$T(\tilde{y}) = \begin{cases} T^{basic} + T^{var}, & \tilde{y} \leq \bar{y}^{EP}, \\ T^{var}, & \bar{y}^{EP} < \tilde{y} \leq \bar{y}^{P}, \\ 0, & \tilde{y} > \bar{y}^{P}. \end{cases} \tag{2}$$

Ordered from the poorest, its three tiers are as follows. *First*, a household in *extreme poverty* ($\tilde{y} \leq \bar{y}^{EP}$) receives the full transfer, the basic plus the variable benefit $T^{basic} + T^{var}$. *Second*, once reported income crosses the extreme-poverty line into ordinary *poverty* ($\bar{y}^{EP} < \tilde{y} \leq \bar{y}^{P}$), the basic benefit stops and only the variable benefit T^{var} remains. *Third*, above the poverty line ($\tilde{y} > \bar{y}^{P}$) the household is ineligible and receives nothing.

The discrete loss of the basic benefit at $\bar{y}^{EP}$ is an eligibility *notch*, a jump in the level of the budget set rather than a change in its slope. Crossing the line by ever so little loses the entire basic benefit at once. The notch draws its force from the income the schedule tests, which is *reported* income, and reported income is not the same object for the two roles. A worker's wage is reported by third parties and cannot be understated. An entrepreneur's profit is self-reported, and the report is therefore a *choice*. The notch makes that choice worth exercising, and, as we show below, it also makes the worker fallback valuable to a founder near the threshold.

Everything that follows turns on the incremental changes the schedule induces, that is, on the *slope* of T near the notch. Just above $\bar{y}^{EP}$ the schedule falls away steeply, $\partial T / \partial \tilde{y} \ll 0$. By reporting slightly less income and falling below $\bar{y}^{EP}$, an agent recovers the entire basic benefit T^{basic}, and the marginal return to keeping reported income down is therefore very high in a narrow band above the line. The margin at $\bar{y}^{P}$ differs in both size and shape. Only the smaller variable benefit T^{var} is at stake there, and its withdrawal is cushioned by the two-year protection rule. The steep region is concentrated in the narrow band just above $\bar{y}^{EP}$.

That slope generates the three responses at the heart of the model, under-reporting that piles reported incomes up just below $\bar{y}^{EP}$, a brake on growth for entrepreneurs near the line, and the dependency channel itself. All three are decisions of the entrepreneur, and we develop them inside the entrepreneur's problem, where they belong.

4.3 Idiosyncratic Processes

Agents differ along three exogenous dimensions. Quality q is a permanent entrepreneurial type, productivity z is a transitory business condition, and e is a wage-earning skill. The first is drawn once and

carried for life, while the latter two evolve stochastically period by period. Together they are the only exogenous sources of heterogeneity in the model. Every other difference across agents, in wealth, in occupation, and in eligibility for the transfer, arises endogenously from decisions taken in response to these three processes. We now specify each in turn.

Quality q. Ability to run a business is given at birth. Each agent draws q once from a log-normal distribution whose logarithm has mean μ_q and standard deviation σ_q, independently of the other states, and carries it for life. The cross-sectional distribution of q is therefore the same at every date, with σ_q governing how dispersed ability is across agents.

Business productivity z. Business conditions vary from year to year, and z captures that variation. It follows an AR(1) in logarithm:

$$\log z' = \rho_z \log z + \sigma_z \, \epsilon', \qquad \epsilon' \sim N(0, 1). \tag{3}$$

Persistence lets a good year make another more likely, while $\rho_z < 1$ pulls z back toward its mean. The process delivers year-to-year dispersion without creating permanent differences between businesses. Permanent differences belong to q, and the separation lets the estimation tell a badly run business from one having a bad year.

Worker skill e. Wage-earning skill varies across workers and over time, and e captures that variation. It follows an AR(1) of the same form, $\log e' = \rho_e \log e + \sigma_e \eta'$. An entrepreneur earns no wage and draws no skill shocks. While running a business, skill moves instead by a deterministic drift,

$$\log e^{+} = \rho_e \log e + \chi_E, \tag{4}$$

where e^{+} is next period's skill for anyone who spent the period running a business, whether they continue or return to wage work. Two forces sit in the expression. Mean reversion $\rho_e < 1$ erodes specialized skill away from wage employment, and the drift χ_E adds the experience of running a business.

4.4 The Worker's Problem

The worker solves a standard income-fluctuation problem, choosing consumption and saving in a risk-free asset, with one addition. The worker can pay the entry cost and start a business. The value function V^W satisfies the Bellman equation:

$$V^W(a, e, q) = \max_{c,\, a' \geq 0} \left\{ u(c) + \beta \, \mathbb{E}_{e'|e} \max \left\{ V^W(a', e', q), \; \bar{V}^E(a', e', q) - F_e \right\} \right\}, \tag{5}$$

subject to:

$$a' = we + R_f a + T(we) - \tau - c, \tag{6}$$

in which $u(c)$ is the worker's CRRA utility and β is the subjective discount factor. The budget constraint is the law of motion of the worker's wealth. The next-period wealth a' is whatever remains of resources after consumption. Those resources consist of labor earnings we, the gross return $R_f a$ on wealth carried into the period, and the transfer $T(we)$, less the lump-sum tax τ. The transfer is evaluated at the wage itself because a worker's income is third-party reported, and a worker therefore takes $T(we)$ as given rather than choosing what to report. The constraint $a' \geq 0$ says that a worker cannot borrow. With no capital to pledge, they can only hold the risk-free asset long.

The inner maximization is the worker's option to remain a worker or become an entrepreneur. Here $\bar{V}^E(a, e, q) \equiv \mathbb{E}_z V^E(a, e, q, z)$ is the expected value of starting a business, with the entrant drawing an initial productivity, and F_e is the entry cost.[4] Define the entry surplus:

$$S(a', e', q) \;\equiv\; \bar{V}^E(a', e', q) - F_e - V^W(a', e', q), \tag{7}$$

where the worker switches occupation and becomes an entrepreneur if and only if $S(a', e', q) \geq 0$. A positive surplus means that the expected value of starting a business, net of the entry cost, is at least as large as the value of remaining a worker.

Because next period's skill e' is not yet realized, the worker forms a probability of entering from current information. Define the entry probability:

$$p(a', e, q) \;\equiv\; \Pr_{e'|e}\left\{ S(a', e', q) \geq 0 \right\} \;=\; \int \mathbb{I}\left\{ S(a', e', q) \geq 0 \right\} \mathrm{d}F(e' \mid e). \tag{8}$$

The entry probability depends on the worker's wealth, quality, and skill. First, wealth a' enters as a choice. A worker who saves more is more likely to find entry worthwhile, which makes wealth a genuine state in the occupational decision. Second, quality q enters as the permanent type, since a more talented agent has more to gain from entering at any wealth and skill. Third, current skill e enters through the distribution of e', and its effect runs in the opposite direction. A higher e' raises the wage a worker earns directly, whereas it raises the value of entrepreneurship only through the fallback the founder would return to on failure. The entry surplus is therefore decreasing in e' over the relevant range, entry is chosen by those whose wage prospects are weakest, and p takes the threshold form $F(\bar{e}(a', q) \mid e)$, where $\bar{e}$ is the skill level at which the worker is indifferent.

The first-order condition with respect to a' yields an intertemporal Euler equation that summarizes the joint consumption-saving and occupational decision.

[4] The entry cost stands in for the one-time outlays of setting up shop, such as renting and outfitting premises, initial inventory, licenses and business registration.

Proposition 1 (Worker's Euler equation)**.** *The worker's optimal saving satisfies:*

$$u'(c) = \beta\Big\{ (1-p)\,\mathbb{E}_{e'|e}\big[\,\mu^W(a',e',q)\,\big|\,stay\,\big] + p\,\mathbb{E}_{e'|e}\big[\,\bar{\mu}^E(a',e',q)\,\big|\,enter\,\big]\Big\}, \tag{9}$$

where p is the entry probability in equation (8), and $\mu^W \equiv \partial V^W/\partial a$ and $\bar{\mu}^E \equiv \mathbb{E}_z\,\partial V^E/\partial a$ are the marginal values of wealth in the two occupations. The condition holds with "$\geq$" when the no-borrowing constraint binds.

Proof. See Section IV.1 of the Internet Appendix. □

The left-hand side is the utility given up by saving one more unit today. The right-hand side is the value of that unit tomorrow, discounted at β. The worker equates the two, as in any consumption-saving problem. The only departure from the canonical case is that the continuation value is not a single object. Next period the worker will be in one of two occupations, and a unit of wealth is worth μ^W in one and $\bar{\mu}^E$ in the other. As a result, the continuation value is the average of the two, weighted by the probability p of entering.[5]

4.5 The Entrepreneur's Problem

The entrepreneur is the object of study. Within the period the entrepreneur makes two decisions, how much capital to operate and how much income to conceal. Exit is not among them. It arrives as a hazard rather than a choice, and the reason is spelled out where the hazard is introduced. We set out the technology and the reporting decision first, then the dynamic problem, and last the exit hazard that carries the dependency channel.

Technology. An entrepreneur of quality q with productivity z operates capital k and hires ℓ efficiency units of labor at the competitive wage w, producing $y = q\,z\,A\,k^{\alpha}\ell^{1-\alpha}$. Labor is hired period by period and is not a dynamic choice, and it can therefore be optimized out. The static problem is:

$$\pi(k,z,q) \;\equiv\; \max_{\ell}\;\big\{\,q\,z\,A\,k^{\alpha}\ell^{1-\alpha} - w\,\ell\,\big\}, \tag{10}$$

with solution:

$$\ell^{\star}(k,z,q) = \left(\frac{(1-\alpha)\,q\,z\,A}{w}\right)^{1/\alpha} k, \qquad \pi(k,z,q) = \Omega\,(q\,z\,A)^{1/\alpha}\,k, \qquad \Omega \equiv \alpha\left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}}. \tag{11}$$

[5] Setting $p = 0$ recovers the textbook case. With no prospect of entry the condition collapses to $u'(c) = \beta\,\mathbb{E}_{e'|e}\left[\mu^W(a',e',q)\right]$, and since $\mu^W = R_f\,u'(c)$ for a worker saving in a risk-free asset, this is the standard Euler equation $u'(c) = \beta R_f\,\mathbb{E}\left[u'(c')\right]$.

Income Concealment. Unlike a worker, whose wage we is reported by third parties and cannot be concealed, an entrepreneur can under-report.[6] We let m denote the fraction of profit concealed. Reported income is then $\tilde{y}_E = (1-m)\pi$, hidden at a cost $\psi(m) = \frac{\phi}{2}m^2\pi$. This cost is *convex* in m. The marginal cost of hiding rises with the fraction concealed, and under-reporting a larger share of profit is therefore disproportionately costly. The entrepreneur chooses concealment to solve:

$$\max_{m} \; T\big((1-m)\pi\big) - \tfrac{\phi}{2}m^2\pi. \tag{12}$$

The first-order condition sets the marginal benefit of a lower report against the marginal cost of hiding it, $-\dfrac{\partial T(\tilde{y}_E)}{\partial \tilde{y}_E} = \phi\, m$, which gives the concealment rule (truncated to $[0,1]$):

$$m^* = -\frac{1}{\phi}\,\frac{\partial T(\tilde{y}_E)}{\partial \tilde{y}_E}. \tag{13}$$

The rule states that concealment is proportional to how steeply the schedule falls at the reported income. Over the flat stretches of T a marginally lower report recovers nothing, $\partial T/\partial \tilde{y}_E = 0$, and $m^* = 0$. Concealment arises only where the schedule declines, and the two thresholds decline very differently. At the extreme-poverty line the entire basic benefit is at stake over a narrow range of income:

$$\left.\frac{\partial T}{\partial \tilde{y}_E}\right|_{\tilde{y}_E \to \bar{y}^{EP}} \ll 0, \tag{14}$$

and an entrepreneur whose profit would place them just above the line recovers the whole of T^{basic} by understating a little. At $\bar{y}^P$ only the smaller T^{var} is at stake and its withdrawal is cushioned, and the slope there is far shallower.

Concealment therefore concentrates almost entirely at the extreme-poverty line. An entrepreneur whose true profit would place them just above the line reports just below instead. The distribution of reported income acquires an excess mass at exactly the point where the benefit is lost. That excess mass is bunching. It is the one place where the mechanism becomes directly visible, and we return to it as evidence.

The Dynamic Problem. We let V^E denote the entrepreneur's value function. It must satisfy the Bellman equation:

$$V^E(a,e,q,z) = \max_{k',m} \; \Big\{u(c) + \beta\big[(1-\pi_d)\,\mathbb{E}_{z'|z}V^E(a',e^+,q,z') + \pi_d\, V^W(a',e^+,q)\big]\Big\}, \tag{15}$$

[6] Wage income is third-party reported. The employer files it directly with the tax authority, and a worker cannot understate we. Self-employment profit, by contrast, is self-reported, and that is the margin m captures.

where π_d is the probability that the business fails and the founder returns to being a worker. It is endogenous, and specifying it is the subject of the next paragraph.

Capital is financed out of the entrepreneur's own net worth and a one-period collateralized loan. Lending is collateralized in the limited-enforcement tradition of Kiyotaki and Moore (1997), Gertler and Kiyotaki (2010), and Jermann and Quadrini (2012), a fraction θ of installed capital can be pledged, and the entrepreneur borrows $b' = \theta k'$, receiving b'/R_f now and repaying b' next period.[7] Capital depreciates at rate δ. Writing $\pi = \pi(k', z, q)$ and $\tilde{y}_E = (1 - m)\pi$, the budget constraint and the law of motion of net worth are:

$$c = a + T(\tilde{y}_E) - \tau - \tfrac{\phi}{2} m^2 \pi + \frac{b'}{R_f} - k', \tag{16}$$

$$a' = \pi + (1 - \delta)k' - b'. \tag{17}$$

Consumption is current net worth, plus the transfer net of tax and concealment cost, less the down-payment $k' - b'/R_f$ the entrepreneur must put up out of their own funds. Next-period net worth is the realized profit plus undepreciated capital, net of the debt repaid.

The continuation value in equation (15) embeds the two ways an entrepreneur can proceed. With probability $1-\pi_d$ the business survives and the founder remains an entrepreneur. With probability π_d the business fails and the founder returns to work, carrying the accumulated skill e^+ and whatever net worth the business leaves behind. Of the two choices in equation (15), concealment enters only the current period and is set by the static rule m^* of equation (13). Capital is the genuinely dynamic decision, and we turn to it next.

The Dependency Channel. That exit probability is where the mechanism operates. We let:

$$s(a, e, q, z) \;\equiv\; V^E(a, e, q, z) - V^W(a, e, q) \tag{18}$$

denote the *operating* surplus, the value of continuing as an entrepreneur over the value of the worker fallback. It is the continuation counterpart of the entry surplus S in equation (7), and at the point of entry the two are linked by $S = \mathbb{E}_z[s] - F_e$. The same comparison governs both margins, and the entry cost is what separates them.

[7] We impose the collateral constraint with equality, and borrowing is therefore always at its limit. The assumption is descriptive for this population before it is convenient. The businesses we study are micro-enterprises, the average beneficiary-founded business starts with about two thousand reais of equity according to Table 1 and essentially none obtains formal credit in its first five years, and at that scale the marginal product of capital sits far above the borrowing rate, which makes how much a founder can pledge, rather than how much they would choose to borrow, the binding determinant of size. Wealth-proportional borrowing limits for entrepreneurs go back to Evans and Jovanovic (1989); Buera et al. (2011) derive quantity limits of this form from limited contract enforcement, and Moll (2014) shows they are the tractable reduced form of that class; our θ plays the same role. Imposing equality then removes b' as a separate choice and leaves capital as the single dynamic decision, which is the computational payoff.

Exit takes a logistic form. Let Λ denote the logistic link:

$$\Lambda(x) \;=\; \frac{\exp(x)}{\exp(x) + \exp(-x)}, \tag{19}$$

which maps an index into a probability, increasing in x and equal to one half at $x = 0$. The exit hazard is:

$$\pi_d(a, e, q, z) = \Lambda\Big(\gamma_0 - \gamma_q \log q - \gamma_z \log z - \kappa\, s\Big), \tag{20}$$

where the first three terms in equation (20) describe the business itself. A higher permanent quality q or a better draw z makes it more viable and less likely to fail, and γ_0 sets the average rate. A hazard built from these alone would say that businesses fail because they are badly run or because they have had a bad year.

The last term says something else, and it is the term this paper is about. It describes nothing about the business. It compares the business with what the founder would have instead, and failure becomes more likely when that alternative is the better of the two. Two businesses identical in q and z can therefore face different exit risks, and the difference lies entirely in what their founders fall back on.

The transfer separates those fallbacks through its conditionality rather than its size. A founder whose business fails returns to wage work, where income is low and the benefit is collected, while an entrepreneur who operates successfully reports income above the line and collects nothing. The fallback is therefore subsidized relative to continuing, which raises V^W without raising V^E by as much, compresses s, and through equation (20) raises exit. The asymmetry is not an assumption. It follows from where the two occupations sit relative to the eligibility line. Means-testing acts as an implicit tax on success, and κ measures how strongly the exit margin responds to it. Setting $\kappa = 0$ nests an exogenous hazard in (q, z).

Exit is a hazard rather than a choice. The two cannot be told apart in the data, since the administrative record shows when a business becomes inactive but not whether the founder chose to close it or was forced out by a lost customer, an illness, or an obligation it could not meet. A threshold rule would attribute all exit to the founder's optimization and would predict a sharp boundary, with every entrepreneur below the line closing and none above it. Observed exit does not look like that, and a hazard that responds to s lets the voluntary and the involuntary components coexist.[8]

Optimality Conditions. The choice of k' is intertemporal. The entrepreneur obtains capital at the price of current consumption, in exchange for a return next period in profit and resale value. The business may not survive to deliver it, in which case the founder is back in wage work. The value of a marginal unit of next-period net worth therefore depends on which state the entrepreneur is in, and we

[8]The specification does not rule out voluntary exit but nests it. As κ grows without bound, equation (20) approaches the indicator $\mathbb{I}\{s < 0\}$, the rule a founder free to close the business at will would follow.

write:

$$\bar{\mu} \equiv (1-\pi_d)\,\mathbb{E}_{z'|z}\,\mu^E(a', e^+, q, z') + \pi_d\,\mu^W(a', e^+, q) \tag{21}$$

for its average. A marginal unit is worth μ^E to a surviving entrepreneur and μ^W to a founder back in wage work, where μ^E and μ^W are the marginal values of wealth in the two occupations.

Proposition 2 (Entrepreneur's capital choice)**.** *Optimal capital equates the cost of a marginal unit, the down payment plus a means-test surcharge, with its expected discounted levered return:*

$$u'(c)\Big(\underbrace{1-\tfrac{\theta}{R_f}}_{\text{down payment}} + \underbrace{\Delta_T}_{\text{surcharge}}\Big) = \beta\,\bar{\mu}\,\big(\underbrace{\pi_k + 1 - \delta - \theta}_{\text{levered return}}\big), \tag{22}$$

where the surcharge is

$$\Delta_T \;\equiv\; \Big[-\frac{\partial T(\tilde{y}_E)}{\partial \tilde{y}_E}\,(1-m^*) + \frac{\phi}{2}(m^*)^2\Big]\,\pi_k \;\geq\; 0, \tag{23}$$

with $\bar{\mu}$ as in equation (21). At an interior concealment optimum the surcharge takes the closed form:

$$\Delta_T \;=\; \phi\, m^*\Big(1-\frac{m^*}{2}\Big)\,\pi_k. \tag{24}$$

Proof. See Section IV.2 of the Internet Appendix. □

One unit of capital costs the down payment $1-\theta/R_f$ that the entrepreneur finances out of consumption. It also raises profit, and with it both the income the schedule claws back and the cost of concealing what is hidden, the two charges the surcharge Δ_T collects. The closed form in equation (24) says where those charges bite. The surcharge is proportional to the concealment the schedule induces, zero over the flat stretches where $m^* = 0$ and largest exactly where the schedule falls steeply, and bunching and staying small are thereby one distortion read on two margins. The slope that piles reported incomes up below $\bar{y}^{EP}$ is the same slope that taxes the next unit of capital. An entrepreneur sitting just below the line faces a choice between growing and keeping the benefit, and for an entrepreneur near the margin the benefit is worth more. It stays small. The eligibility cliff does not only push entrepreneurs out of business. It also holds back those that stay, and it leaves actual capital below what the technology would justify.

The entrepreneur could grow and conceal more to stay eligible, but m^* is already set optimally, and its response therefore contributes nothing to first order. The surcharge holds concealment fixed, and the closed form only substitutes the level at which it is held. On the other side of equation (22), that unit of capital pays off next period as the marginal product π_k plus undepreciated resale $1-\delta$, net of the debt repayment θ, valued at $\bar{\mu}$.

The hazard dampens the investment decision. Although it is determined by the current state and is therefore taken as given when k' is chosen, the hazard still shapes that choice through $\bar{\mu}$. Holding the two marginal values fixed, a higher π_d places more weight on what wealth is worth outside the business and less on what it is worth inside. At a given level of net worth, the inside value is the larger of the two, since an entrepreneur holds consumption down in order to fund capital while a worker with the same wealth does not, and $\bar{\mu}$ therefore falls as the hazard rises. The marginal benefit in equation (22) falls with it and investment weakens. The dependency term reaches investment by this route as well as exit.

The two distortions reach the same place by different routes. Δ_T works directly. An entrepreneur near the threshold gives up part of the transfer by growing, which raises the effective cost of capital and holds investment down. The dependency term works through the horizon. The transfer raises V^W, and for an entrepreneur whose business is going badly the two values already sit close together, close enough that the transfer can carry V^W above V^E. The exit probability π_d rises. A higher π_d puts more weight on μ^W and less on μ^E in $\bar{\mu}$, and a unit of net worth is worth more inside the business than outside it, which leaves the value of carrying capital into next period lower. An entrepreneur who expects to be gone has little reason to build. Investment falls again, for a reason that has nothing to do with where the threshold sits.

Neither channel requires the entrepreneur to choose failure. The first is a response to a schedule, the second a response to a horizon. Both hold capital below what the technology would justify, which is the model's counterpart of the growth shortfall documented in Section 3.

Timing. With both problems in place, the sequence within a period reads as follows.

1. **Incomes accrue.** A worker earns we and receives $T(we)$; an entrepreneur produces $\pi(k', z, q)$, sets m^*, reports $\tilde{y}_E$, and receives $T(\tilde{y}_E)$ net of the concealment cost. Everyone pays the lump-sum tax τ.
2. **Choices.** The worker chooses savings and the entrepreneur chooses capital k'; consumption is the residual. The exit probability π_d is a function of the current state and is therefore already determined when k' is chosen.
3. **Shocks realize.** Worker skill draws e', business productivity draws z', and a founder's skill drifts to e^+.
4. **Occupational transitions.** Each incumbent entrepreneur exits to work with the hazard π_d, and each worker whose entry surplus is non-negative pays F_e, draws z_0, and becomes an entrepreneur. A founder whose business exits enters the next period as a worker, and is tested for eligibility on the wage we^+ rather than on business income.

4.6 Law of Motion for the Distributions

Let μ^W and μ^E denote the measures of workers and entrepreneurs over their respective states, normalized together to the unit population. Next period's entrepreneurs are this period's survivors plus this period's entrants:

$$\mu^{E\prime} = \underbrace{(1-\pi_d)\,\mathcal{T}^E\,\mu^E}_{\text{survivors}} + \underbrace{\mathbb{I}\{S \geq 0\}\,\mathcal{T}_0^E\,\mu^W}_{\text{entrants}}, \tag{25}$$

and next period's workers are this period's stayers plus this period's failures:

$$\mu^{W\prime} = \underbrace{\mathbb{I}\{S < 0\}\,\mathcal{T}^W\,\mu^W}_{\text{stayers}} + \underbrace{\pi_d\,\mathcal{T}^{EW}\,\mu^E}_{\text{failures}}. \tag{26}$$

The interpretation is that a fraction $1-\pi_d$ of incumbents carries on into next period, the rest return to wage work, and the workers whose surplus is non-negative start businesses of their own.[9]

Two features of these flows matter here. An entrant begins producing at once, with the wealth accumulated as a worker, which makes entry wealth an outcome of the model rather than a calibrated object. A founder whose business fails returns to wage work with the skill accumulated while running the business. That second point is where the dependency channel does its work, since the skill a founder brings back sets the wage the failure state pays, and the wage decides whether the transfer is still collected there.

A stationary distribution is a fixed point of equations (25) and (26).

4.7 Market Closure and Stationary Equilibrium

Three objects remain to be closed, the efficiency-unit wage w, the gross risk-free rate R_f, and the government budget. The modeled agents are a *segment of a much larger labor and credit market*, and they are far too small a group to move the prices in it. We therefore take w and R_f as given, at their observed levels, which applies the small-open-economy closure $R_f = R^*$ to both factors rather than to capital alone.

The segment's factor flows are then whatever its policies imply. Entrepreneurs hire $\int \ell(k', z, q)\, d\mu^E$ efficiency units of labor while workers supply $\int e\, d\mu^W$, and the difference is bought from or sold to the wider market at w. Collateralized borrowing $b = \theta k$ and workers' holdings of the risk-free asset are intermediated the same way at R_f. Neither factor market clears within the segment, and neither is re-

[9] Each $\mathcal{T}$ pushes a measure one period forward, sending every state to where the policies and the shocks put it. The superscript names the transition, with $\mathcal{T}^E$ and $\mathcal{T}^W$ keeping an agent in the role it already occupies, $\mathcal{T}_0^E$ taking a worker into business, and $\mathcal{T}^{EW}$ returning a founder whose business failed to wage work. Section V.1 of the Internet Appendix states all four as integrals. We do not set out the measure-theoretic construction of a continuum of agents and of integration over it, and refer instead to Feldman and Gilles (1985) and Judd (1985).

quired to. The segment simply runs a net position with the rest of the economy, and that position adjusts. The resource constraint holds economy-wide, and no goods-market condition is therefore imposed on the segment either.

A fully closed general-equilibrium version, in which w adjusts to equate the two integrals above, is available but unused. Nothing in the mechanism turns on it. Beneficiary entrepreneurs hire labor and raise capital in the broader Brazilian market at the going prices, and internal factor-market clearing is neither expected nor imposed.

The transfer program is financed by the lump-sum tax, which must satisfy:

$$\tau = \int T(we)\, d\mu^W + \int T(\tilde{y}_E)\, d\mu^E. \tag{27}$$

Lump-sum financing is an assumption, and it is worth saying why it costs us little. A distortionary tax would itself change how much people work, save and invest, and those changes would be tangled up with the responses we are trying to measure. The lump-sum tax leaves the exit and investment margins clean, and what we attribute to the transfer's design is therefore not contaminated by how it is raised. The program is also small relative to the public finances it sits in, which makes the deadweight loss of funding it second order for our purposes.

A stationary equilibrium is a collection of value functions $\{V^W, V^E\}$, policies $\{g^W, k', m\}$, an exit hazard π_d, measures $\{\mu^W, \mu^E\}$, and a tax τ with four properties. The value functions solve the Bellman equations and the optimality conditions given the prices (w, R_f) and the hazard. The hazard satisfies equation (20) given the value functions. The measures are invariant under the law of motion. The tax balances equation (27). The wage and the risk-free rate are held at their observed levels throughout. In the estimation we set $\tau = 0$, which finances the program outside the segment, and where the budget is required to balance, τ is pinned by equation (27).

5 Estimation

We estimate most of the model's parameters by simulated method of moments. A first tier is set outside the estimation, from program rules, standard values, or data the model takes as inputs, and Section 5.1 lists it. We then estimate the nine remaining parameters, the entry cost F_e, the quality distribution (μ_q, σ_q), the exit hazard $(\gamma_0, \gamma_q, \gamma_z, \kappa)$, and the productivity process (ρ_z, σ_z), by matching eighteen moments of the administrative data.

5.1 Calibrated parameters

Table 3 lists the first tier. Preferences and technology take standard values, risk aversion $\sigma = 2$, discount factor $\beta = 0.96$, capital share $\alpha = 0.33$, depreciation $\delta = 0.10$, a gross risk-free rate $R_f = 1.04$, and a collateral fraction $\theta = 0.40$. The transfer replicates the program at the scale of the sample period. The median formal wage is the numeraire, the benefit levels $T^{basic} = 0.06$ and $T^{var} = 0.02$ reproduce the 2013–2021 payment ledgers, and the thresholds $\bar{y}^{EP} = 0.70$ and $\bar{y}^{P} = 0.90$ place the bottom quarter of the modeled skill distribution inside eligibility, matching program coverage. The skill process is fitted to the employment registry, an AR(1) on the pre-founding worker panel of 72.1 million RAIS person-years gives $\rho_e = 0.732$ and $\sigma_e = 0.232$, and the experience drift is $\chi_E = 0.05$.

[Place Table 3 about here]

The concealment cost ϕ is calibrated rather than estimated. The criterion is flat in ϕ across two orders of magnitude, concealment moves reported income but almost no resources, and the near-threshold excess mass the model generates is numerically fragile on a discrete state space. We set $\phi = 2.0$, the criterion-minimizing point of the sweep, and use the bunching evidence of Section 3.2 as the qualitative signature of the mechanism rather than as a level the model must hit.

5.2 Estimator

The estimator minimizes

$$Q(\theta) = \left[m(\theta) - m^{data}\right]' \Omega^{-1} \left[m(\theta) - m^{data}\right], \tag{28}$$

where $m(\theta)$ stacks the eighteen moments listed in Table 4 and Ω is an influence-function estimate of their variance–covariance matrix, clustered by municipality for the registry moments and by household for the survey moments. Because the moments are computed on the model's exact stationary distribution rather than on a finite simulated panel, the estimator carries no simulation error and needs no simulation correction. The gap moments deserve one sentence of emphasis. They are the beneficiary coefficients from linear probability survival regressions with industry, initial equity, and founding-year fixed effects interacted, and the model is put through the identical regression on the entrant cells of its exact stationary distribution, with the same fixed effects and the same horizons. Matching the estimator rather than a summary of it keeps the model and the data answering the same question.

[Place Table 4 about here]

5.3 Identification

The empirical difficulty the estimation must overcome is that composition and dependence are observationally equivalent in a survival regression. The beneficiary coefficient is a mixture of the two forces, selection into entry and treatment after entry, both moving it in the same direction at every horizon, and no set of controls can decompose a single coefficient into two forces. The moment vector is built to let the model do what the regression cannot, the survival levels of the two groups discipline composition, and the speed at which the adjusted gap widens with business age is left to carry dependence.

Although every parameter affects every moment to some degree, the moments divide into blocks that carry the identification, and we discuss them in turn.

The two aggregate rates anchor the model's turnover. The entry rate responds first to the entry cost, a higher F_e keeps marginal entrants out, and it responds to little else once the survival system is matched. The exit rate pins the hazard intercept γ_0, which scales failure risk for every business at once. Neither rate can separate why businesses fail, and that burden falls on the survival curves.

The survival levels separate the heterogeneity parameters through their shape. A survival curve falls steeply at young ages and flattens as failure removes the weakest businesses, and how much it flattens depends on what failure selects on. Selection on the permanent type q improves the surviving pool permanently, and the strength of the flattening across horizons therefore moves with γ_q and with the dispersion σ_q that gives selection something to work on. Selection on the transitory state z improves the pool only briefly, mean reversion undoes it at rate ρ_z, and the curvature of the curves at short horizons therefore separates $(\gamma_z, \rho_z, \sigma_z)$ from the permanent block. One block cannot mimic the other, a hazard loaded on z produces curves that converge as horizons lengthen, while a hazard loaded on q produces curves that fan out. The location parameter μ_q is the weakest link in this chain, it shifts profits everywhere at once and is disciplined mainly through the entry margin, and the estimates below flag it accordingly.

The level difference between the beneficiary and the non-beneficiary curves carries the composition channel. Beneficiary founders enter from lower skill and lower wealth, the model's entry margin produces the difference without any dedicated parameter, and matching both curves rather than their gap keeps composition disciplined by levels rather than absorbed into κ.

The dependence parameter reads a feature none of the above can generate, the speed at which the adjusted gap widens with business age. A formal sensitivity calculation in the spirit of Andrews et al. (2017) shows that $\hat{\kappa}$ loads on the year-one and year-two adjusted gaps with near-equal and opposite weights, +0.81 and −0.94, and only lightly on everything else. $\hat{\kappa}$ is identified by the early-horizon slope of the gap profile, not by its level. The economics of the weighting is the composition-versus-treatment distinction. Whatever composition difference entrants carry is carried from the first year, and selection therefore produces a gap that is present at birth and stays flat. A margin operating through exit builds the gap year by year. The slope separates the two, and the slope is what κ is fitted to.

The restricted model makes the same point in reverse. Imposing $\kappa = 0$ collapses the hazard parameters toward collinearity, the γ_q–γ_z correlation reaches 1.00, because without the surplus term the hazard has no independent lever for the early-horizon dynamics and must contort the business block to chase them.

5.4 Estimates

Table 5 reports the estimates with the dependency channel free and shut down. Three features of the free column deserve comment. First, the business block is precisely estimated and economically sensible. Exit falls steeply in the transitory state ($\gamma_z = 2.71$) and more gently in permanent quality ($\gamma_q = 0.50$), matching the registry fact that most young-business failure looks like bad years rather than bad types. The entry cost of 0.77, about nine months of the median wage, is the right order for premises, inventory, and registration in this population. Second, one parameter is individually weak and we flag it. The location μ_q carries a t-statistic of 1.7 inside a near-collinear dispersion complex with σ_q and σ_z; the data pin the dispersion combination, anchored by σ_q ($t = 6.8$), not the individual location. Third, the dependency parameter is estimated at $\hat{\kappa} = 0.284$ with a standard error of 0.052, a t-statistic of 5.44.

[Place Table 5 about here]

The test of the dependency channel is the distance between the two columns. With administrative samples in the millions the J statistic rejects any parsimonious model mechanically, and inference therefore uses the difference across the nested specifications,

$$D = J_{\kappa=0} - J_{\kappa\ free} = 6{,}509.0 - 6{,}257.5 = 251.5 \ \sim \ \chi^2(1), \tag{29}$$

against a critical value of 3.84. The restriction is rejected decisively, and the penalty is paid almost entirely on the survival system, not on the static moments, which if anything slightly favor the restricted model. Removing the margin costs the model exactly the dynamics the margin exists to explain.

The economic magnitude of $\hat{\kappa}$ is easiest to read in hazard points. At the estimates, the gap in annual exit hazards between beneficiary and non-beneficiary founders is about seven percentage points in the founding year and roughly three and a half from the second year on, against an average exit rate of ten percent. Eligibility for a transfer worth eight percent of the median wage nearly doubles a young business's failure risk, not by draining its resources but by improving what its owner returns to. The decay of the wedge with business age follows from the skill drift, the fallback wage recovers as the business survives, and the founder gradually outgrows the eligibility that made closing attractive.

5.5 Fit

Figure 2 plots the model against the data, the four survival curves in levels and the adjusted gap profile, and Table 4 reports the moment-by-moment comparison. The aggregate rates are matched closely, exit to within one percentage point and entry to within 1.2 (model 9.1 percent against a 10.3 percent target). The non-beneficiary survival curve tracks the data to within 3.5 percentage points at every horizon.

The beneficiary curve does not, and we report the failure plainly. The model places beneficiary survival too low at short horizons, about eleven percentage points in the first year, and the adjusted gap opens at about +6 percentage points where the data put the first-year gap near zero. The model front-loads a disadvantage the data build gradually. The overstatement narrows with horizon, to roughly +3 points at five years, and the emergence rate is right, the year-one-to-year-two slope of the gap is about +2.2 points in the data and +2.3 in the model. Since Section 5.3 showed that $\hat{\kappa}$ is identified by precisely this slope rather than by the level, the fit failure and the identification are not the same object, a point we verify directly by re-estimating without the year-one moments.

The front-loading has a diagnosable source. A founder's fallback wage recovers through the experience drift, and the drift carries a typical founder above the eligibility line within a year or two of founding, which concentrates the model's extra exits too early. The data say eligibility persists longer than the model's wage recovery allows, benefit receipt among surviving beneficiary founders is still 62 percent after one year and 35 after five according to Section 2.3, and an enrollment state with administrative persistence is the natural extension. Section 5.7 reports where that extension stands.

[Place Figure 2 about here]

The estimation's discipline does not end with fit. The sample carries a feature few structural estimations enjoy, the policy shut the mechanism off once. The bunching moment is estimated on the 2013–2019 means-tested regime, and switching the income test off inside the model kills the excess mass the concealment margin produces. Brazil ran that switch in 2020–21, and the excess mass at the threshold collapsed to statistical zero in the data (Section 3.2). The mechanism's fingerprint appears where the model says it should and vanishes when the policy that creates it is removed, and Section 6.2 carries the episode's second, registered prediction.

5.6 Untargeted moments

A model fitted to eighteen moments should be checked against what it was not fitted to, and we report two such checks. The first, reported in Table 6, is the benefit receipt profile, the share of surviving beneficiary founders still drawing the transfer at each business age, which the payment ledgers put at 62 percent after one year and 35 after five. Nothing in the moment vector uses receipt after founding, and the model generates the profile entirely from the skill drift and the reporting rules. The second is the

growth margin. The estimation never targets the MEI-to-ME transition, and the model's prediction that beneficiary businesses cross the legal size ladder less often is the investment distortion of Proposition 2 at work, the counterpart of the 12 percent MEI-to-ME shortfall documented in Campello et al. (2026).

5.7 A disclosure on the estimates

The estimates above were sealed in July 2026, and one development postdates them. The estimation ran on a model whose wealth distribution the computational grid silently truncated. With constant returns at the business level, fixed prices under the segment closure, and a heavy-tailed permanent type, the model has no stationary wealth distribution, and the grid ceiling was censoring the divergence. The repair is institutional rather than numerical, the legal revenue ceiling of the ME regime bounds scale, restores stationarity, and makes every reported moment invariant to the wealth ceiling. Re-estimating the repaired model from scratch, with a twenty-start multistart and the full grid ladder, moves the dependency estimate to a precise zero ($\hat{\kappa} = 0.002$, SE 0.005), while the repaired model overshoots the adjusted gaps from composition alone. The two model versions now disagree about a measurable object. The sealed model predicts that the Auxílio episode narrowed the beneficiary exit gap by 18 percent, the repaired model predicts no change, and the registry series that adjudicates between them is under construction. The diagnosed root is receipt persistence, benefit receipt in the data outlives income eligibility by years (Section 2.3) while both model versions test income every period. We have therefore extended the model with an enrollment state. A household that enrolls keeps the benefit regardless of current income and is removed only when an audit, arriving at a rate calibrated to the receipt profile, finds income above the line. Sign-up is immediate, removal is slow, and the wedge the dependency margin feeds on persists as long as receipt does. Production-grid estimates of the extended model are complete. With the audit rate reset to the balanced receipt profile constructed by our data team, the extended model reproduces the level and the year-over-year decay of benefit receipt among surviving founders (Table 6), with a largest miss of 4.0 percentage points across horizons. The dependence parameter is estimated at $\hat{\kappa} = 0.21$, between the sealed 0.28 of Table 5 and the repaired model's zero, and successive deepenings of the search have moved it upward, not down, each improvement in fit arriving with a larger dependence parameter. The moderation is the mechanism at work rather than a retreat, the enrollment state now carries part of the persistence that the sealed model could route only through κ, and what remains is the part of the margin that receipt dynamics cannot imitate. The margin itself is not disposable. In a nested comparison polished from a common point, the restricted model with $\kappa = 0$ comes to rest and stays fixed across repeated deep restarts, while the unrestricted model continues to improve, and the distance between the two has widened with every deepening to 131 criterion points. Formal standard errors and the efficient weighting await the moment covariances now being assembled on the data side. The Auxílio series remains the arbiter among the model versions, and a moderate κ is

the reading most consistent with its suspension-year evidence. We report this openly rather than quietly refitting.

[Place Table 6 about here]

6 Counterfactuals and Validation

The experiments in this section serve two purposes. They translate the estimated parameters into economic magnitudes, what $\hat{\kappa}$ means in survival points rather than in hazard units, and they answer the policy question, which margins of the schedule reach the survival gap and which do not. Table 7 collects the experiments, and the subsections walk through them in turn. Each experiment changes one feature of the transfer schedule at the estimated parameters, holding prices at their observed levels under the segment closure of Section 4.7. The exercises are accordingly local to the estimated program, statements about the margins of Bolsa Família rather than about arbitrarily large reforms.

[Place Table 7 about here]

6.1 Anatomy of the gap

The first experiment freezes the dependence margin and changes nothing else. We set $\kappa = 0$, hold every other parameter at its estimate, and recompute the stationary equilibrium and the cohort survival curves. The five-year beneficiary survival gap falls by 92 percent. On the regression-adjusted definition the gap falls from 10.1 percentage points to 0.5, and the two definitions agree. The freeze spares only the composition term, the part of the gap produced by who selects into business ownership rather than by how the schedule treats them afterward.

The exercise is a decomposition, not a re-estimation, and the two answer different questions. Freezing κ asks what the estimated world looks like with the margin switched off. Whether a model without the margin could fit the data is a separate question, and Section 5.4 answered it: the restriction is rejected with a distance statistic of 251.5 at the sealed estimates, 131 in the enrollment extension of Section 5.7. Together the two exercises bracket the margin. It is necessary to fit the dynamics, and within the fitted model it carries nearly all of the gap.

The freeze also says what the gap is not. It is not a liquidity difference. Transfers in the model relax every recipient's budget, the freeze leaves that support fully in place, and the gap disappears anyway. The freeze removes only the asymmetry of what exit protects, the founder who closes keeps the benefit while the founder who grows loses it.

6.2 The Auxílio episode

During 2020–21 Brazil suspended means testing. Auxílio Emergencial replaced the two-threshold schedule with payments that did not depend on reported income near the eligibility lines, which is as close to switching the mechanism off as policy has come. We committed to the model's prediction for this episode in writing before the corresponding empirical series was produced, and the adjudicating series is built by the companion team with no input from the model; prediction and evidence meet only at the comparison.

The experiment replaces the estimated schedule with a transfer of equal generosity paid without an income test, holding all parameters at their estimates. Its credibility rests on fit, the margins the reform acts on are the exit hazard and the concealment rule, the model matches the aggregate exit rate to within a percentage point (Section 5.5), and it generates the bunching signature the concealment rule exists to produce (Table 4). The prediction has two parts. Removing the income test closes 18 percent of the beneficiary exit gap, with a range of 15.9 to 18.2 percent across grids and specifications, and the remaining four-fifths, being composition, survives the reform. The share, not the level, is the registered object. With no income test the return to under-reporting is zero, $m^* = 0$ in equation (13), and bunching should vanish.

The data have answered the second part. Section 3.2 showed that the excess mass at the extreme-poverty threshold collapses to statistical zero in the 2020–21 survey rounds. The first part is pending, the survival gap for the Auxílio period is under construction from the payment ledgers, and this subsection will report the predicted-against-realized comparison when it arrives. A model that predicts a specific share of a gap in advance can fail. That is the point of registering the number, and whatever the series shows, the comparison is reported here.

6.3 Which margins reach the gap

The Auxílio experiment bounds what any redesign of the schedule can achieve. Four-fifths of the beneficiary exit gap survives even the complete removal of the income test, being composition produced at entry, and no reform of the withdrawal schedule reaches it. The policy question is what reaches the remaining fifth, and the natural-looking instrument does not.

Smoothing the cliff barely moves the gap. Replacing the notch with a linear taper, the basic benefit paid in full up to $\bar{y}^{EP}$ and withdrawn gradually over a band above it, changes the slope of the schedule where concealment responds, and bunching falls accordingly. The exit margin is nearly untouched. Cutting the withdrawal rate to a quarter of its estimated value moves the five-year beneficiary survival gap from 11.54 to 11.46 percentage points, a reduction of 0.08 points, at 125 percent additional program spending, and preserving the benefit floor at an unchanged rate moves the gap not at all. The structure of the hazard says why. The dependence term responds to the surplus s, the distance between

the value of operating and the value of the fallback, and a taper reshapes the schedule's local slope without much changing that distance. Concealment is a response to a slope. Dependence is a response to a gap between states, and a taper treats the first while leaving the second.

Decoupling reaches the margin. In this experiment an operating owner is means-tested on the wage they would earn as a worker, exactly as a worker is, which sets the transfer in the operating state equal to the transfer in the fallback state and drives the wedge between them to zero, holding the benefit level, the targeting, and the program spending where they were. The disciplining comparison is cost. An equally costly across-the-board expansion raises every recipient's resources, leaves the asymmetry intact, and does little to the gap. In the language of comparative statics, the survival gap is highly elastic with respect to the wedge between the operating and fallback states and nearly inelastic with respect to both the slope of the withdrawal and the level of the benefit. The design margin, not the generosity margin, is where the dependence channel lives, in the spirit of the earned-income disregards used in developed-country welfare programs.

6.4 Taking stock

Three statements summarize the quantitative case. The dependence margin is statistically necessary, removing it is rejected with a distance statistic of 251.5 on one degree of freedom at the sealed estimates, and 131 in the enrollment extension of Section 5.7. It is economically large within the model, carrying most of the fitted survival gap, while the larger part of the raw disadvantage is composition and sits beyond the schedule's reach. And it is falsifiable out of sample, one of its two Auxílio predictions is already confirmed in the bunching data and the other is registered and awaiting the series. A model without the margin would not merely fit the dynamics worse. It would conclude that the gap is composition all the way down, price every reform of the schedule at zero, and leave policy with nothing to act on, which is why establishing the margin, and not only measuring the gap, is the paper's contribution.

7 Conclusion

Social programs increasingly hope to seed businesses, and in Brazil the hope is partly realized, Bolsa Família measurably raises entry into business ownership. This paper asked why the businesses it helps create die young, and answered with a mechanism rather than a correlation. A means-tested transfer is kept by a founder who closes and lost by one who grows. The program design itself makes the fallback more attractive than the business it insures, and an exit margin that responds to that comparison is enough to reproduce the age profile of the beneficiary survival gap in the administrative registry.

The margin is not a free parameter doing convenient work. It is precisely estimated, removing it is rejected with a distance statistic of 251.5 on one degree of freedom, and it is identified by the speed at

which the disadvantage emerges rather than by its size, a feature selection alone cannot generate. Within the estimated model the margin carries most of the fitted survival gap. Most of the raw disadvantage, however, reflects composition, the marginal businesses the program pulls into existence, and that layer sits beyond the reach of any withdrawal schedule. The model's sharpest support is out of sample. When Brazil suspended means testing in 2020–21, income bunching at the eligibility threshold vanished, as the mechanism requires, and the model's registered prediction for the same episode, that roughly a fifth of the exit gap disappears, awaits the corresponding series.

The policy lesson inverts the usual instinct. The natural repair for an eligibility cliff is a gentler slope, and the model says the slope is the wrong margin, dependence responds to the distance between operating and not operating, not to the local taper. The repair that reaches the margin is decoupling, letting owners keep benefits while the business operates, in the spirit of earned-income disregards. Transfer programs that want durable businesses from their beneficiaries should attend to what closure protects rather than to what success forfeits.

The model stops short in places we have flagged. Eligibility in the data outlives eligibility in the model, receipt persists years after income has passed the threshold, and Section 5.7 reports the enrollment extension built to close that gap and the re-estimation it is undergoing. The legal size ladder we use as a boundary could itself be modeled as a tax notch, making the growth margin testable at a second threshold. Both extensions sharpen the same question this paper opened, how the design of a safety net shapes what its recipients build.

References

Aiyagari, S. Rao, 1994, Uninsured idiosyncratic risk and aggregate saving, *Quarterly Journal of Economics* 109, 659–684.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro, 2017, Measuring the sensitivity of parameter estimates to estimation moments, *The Quarterly Journal of Economics* 132, 1553–1592.

Balboni, Clare, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil, 2022, Why do people stay poor?, *Quarterly Journal of Economics* 137, 785–844.

Bandiera, Oriana, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman, 2017, Labor markets and poverty in village economies, *Quarterly Journal of Economics* 132, 811–870.

Banerjee, Abhijit, Esther Duflo, Nathanael Goldberg, Dean Karlan, Robert Osei, William Parienté, Jeremy Shapiro, Bram Thuysbaert, and Christopher Udry, 2015, A multifaceted program causes lasting progress for the very poor: Evidence from six countries, *Science* 348, 1260799.

Bergstrom, Katy, William Dodds, and Juan Rios, 2025, Welfare Analysis of Changing Notches: Evidence from Bolsa Família, *American Economic Journal: Economic Policy* 17, 122–161.

Bianchi, Milo, and Matteo Bobba, 2013, Liquidity, risk, and occupational choices, *Review of Economic Studies* 80, 491–511.

Blattman, Christopher, Nathan Fiala, and Sebastian Martinez, 2014, Generating skilled self-employment in developing countries: Experimental evidence from uganda, *The Quarterly Journal of Economics* 129, 697–752.

Blattman, Christopher, Eric P. Green, Julian Jamison, M. Christian Lehmann, and Jeannie Annan, 2016, The returns to microenterprise support among the ultrapoor: A field experiment in postwar uganda, *American Economic Journal: Applied Economics* 8, 35–64.

Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin, 2011, Finance and development: A tale of two sectors, *American Economic Review* 101, 1964–2002.

Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin, 2015, Entrepreneurship and financial frictions: A macrodevelopment perspective, *Annual Review of Economics* 7, 409–436.

Buera, Francisco J., and Yongseok Shin, 2013, Financial frictions and the persistence of history: A quantitative exploration, *Journal of Political Economy* 121, 221–272.

Cagetti, Marco, and Mariacristina De Nardi, 2006, Entrepreneurship, frictions, and wealth, *Journal of Political Economy* 114, 835–870.

Campello, Murillo, Guilherme Junqueira, and Spyridon Lagaras, 2026, Entrepreneurship on a safety net: Evidence from the world's largest cash transfer program, Working paper, University of Florida and University of Illinois.

Chetty, Raj, John N. Friedman, Tøre Olsen, and Luigi Pistaferri, 2011, Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from danish tax records, *Quarterly Journal of Economics* 126, 749–804.

de Mel, Suresh, David McKenzie, and Christopher Woodruff, 2008, Returns to capital in microenterprises: Evidence from a field experiment, *The Quarterly Journal of Economics* 123, 1329–1372.

de Mel, Suresh, David McKenzie, and Christopher Woodruff, 2014, Business training and female enterprise start-up, growth, and dynamics: Experimental evidence from sri lanka, *Journal of Development Economics* 106, 199–210.

Evans, David S., and Boyan Jovanovic, 1989, An estimated model of entrepreneurial choice under liquidity constraints, *Journal of Political Economy* 97, 808–827.

Feldman, Mark, and Christian Gilles, 1985, An expository note on individual risk without aggregate

uncertainty, *Journal of Economic Theory* 35, 26–32.

Gerard, François, Joana Naritomi, and Joana Silva, 2024, Cash transfers and formal labor markets: Evidence from brazil, Working paper.

Gertler, Mark, and Nobuhiro Kiyotaki, 2010, Financial intermediation and credit policy in business cycle analysis, in Benjamin M. Friedman, and Michael Woodford, ed.: *Handbook of Monetary Economics*vol. 3 . pp. 547–599 (Elsevier).

Gertler, Paul J., Sebastian W. Martinez, and Marta Rubio-Codina, 2012, Investing cash transfers to raise long-term living standards, *American Economic Journal: Applied Economics* 4, 164–192.

Ghosh, Pulak, and Nishant Vats, 2022, Safety nets, credit, and investment: Evidence from a guaranteed income program, Working Paper Washington University in St. Louis.

Gottlieb, Joshua D, Richard R Townsend, and Ting Xu, 2022, Does career risk deter potential entrepreneurs?, *The Review of Financial Studies* 35, 3973–4015.

Hennessy, Christopher A., and Toni M. Whited, 2005, Debt dynamics, *Journal of Finance* 60, 1129–1165.

Hennessy, Christopher A., and Toni M. Whited, 2007, How costly is external financing? evidence from a structural estimation, *Journal of Finance* 62, 1705–1745.

Holtz-Eakin, Douglas, David Joulfaian, and Harvey S. Rosen, 1994, Sticking it out: Entrepreneurial survival and liquidity constraints, *Journal of Political Economy* 102, 53–75.

Hombert, Johan, Antoinette Schoar, David Sraer, and David Thesmar, 2020, Can unemployment insurance spur entrepreneurial activity? evidence from france, *Journal of Finance* 75, 1247–1285.

Hopenhayn, Hugo A., 1992, Entry, exit, and firm dynamics in long run equilibrium, *Econometrica* 60, 1127–1150.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, 1995, Precautionary saving and social insurance, *Journal of Political Economy* 103, 360–399.

Huggett, Mark, 1993, The risk-free rate in heterogeneous-agent incomplete-insurance economies, *Journal of Economic Dynamics and Control* 17, 953–969.

Hurst, Erik, and Annamaria Lusardi, 2004, Liquidity constraints, household wealth, and entrepreneurship, *Journal of Political Economy* 112, 319–347.

Jermann, Urban, and Vincenzo Quadrini, 2012, Macroeconomic effects of financial shocks, *American Economic Review* 102, 238–271.

Judd, Kenneth L., 1985, The law of large numbers with a continuum of iid random variables, *Journal of Economic Theory* 35, 19–25.

Kiyotaki, Nobuhiro, and John Moore, 1997, Credit cycles, *Journal of Political Economy* 105, 211–248.

Kleven, Henrik Jacobsen, 2016, Bunching, *Annual Review of Economics* 8, 435–464.

Kleven, Henrik J., and Mazhar Waseem, 2013, Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from pakistan, *Quarterly Journal of Economics* 128, 669–723.

Li, Shaojin, Toni M. Whited, and Yufeng Wu, 2016, Collateral, taxes, and leverage, *Review of Financial Studies* 29, 1453–1500.

Low, Hamish, and Luigi Pistaferri, 2015, Disability insurance and the dynamics of the incentive insurance trade-off, *American Economic Review* 105, 2986–3029.

Lucas, Robert E., Jr., 1978, On the size distribution of business firms, *The Bell Journal of Economics* 9, 508–523.

Macours, Karen, Patrick Premand, and Renos Vakis, 2022, Transfers, diversification and household risk

strategies: Can productive safety nets help households manage climatic variability?, *The Economic Journal* 132, 2438–2470.

Midrigan, Virgiliu, and Daniel Yi Xu, 2014, Finance and misallocation: Evidence from plant-level data, *American Economic Review* 104, 422–458.

Moffitt, Robert, 1992, Incentive effects of the U.S. welfare system: A review, *Journal of Economic Literature* 30, 1–61.

Moll, Benjamin, 2014, Productivity losses from financial frictions: Can self-financing undo capital misallocation?, *American Economic Review* 104, 3186–3221.

Nikolov, Boris, and Toni M. Whited, 2014, Agency conflicts and cash: Estimates from a dynamic model, *Journal of Finance* 69, 1883–1921.

Quadrini, Vincenzo, 2000, Entrepreneurship, saving, and social mobility, *Review of Economic Dynamics* 3, 1–40.

Saez, Emmanuel, 2010, Do taxpayers bunch at kink points?, *American Economic Journal: Economic Policy* 2, 180–212.

Warusawitharana, Missaka, and Toni M. Whited, 2016, Equity market misvaluation, financing, and investment, *Review of Financial Studies* 29, 603–654.

Xu, Wenjian, 2022, Social insurance and entrepreneurship: The effect of unemployment benefits on new-business formation, *Strategic Entrepreneurship Journal* 16, 522–551.

Figure 1: **Income Distributions at the Eligibility Thresholds**

This figure plots the distribution of pre-transfer household per capita income of Bolsa Família beneficiaries around the two eligibility thresholds during 2013–2019. Pre-transfer income is reported income minus expected Bolsa Família payments given family composition and program rules, normalized so that the year-specific threshold equals 100. Panel A centers on the extreme-poverty threshold, below which households qualify for the basic benefit; Panel B centers on the poverty threshold. The vertical line marks the threshold. Data are from PNAD Contínua annual microdata.

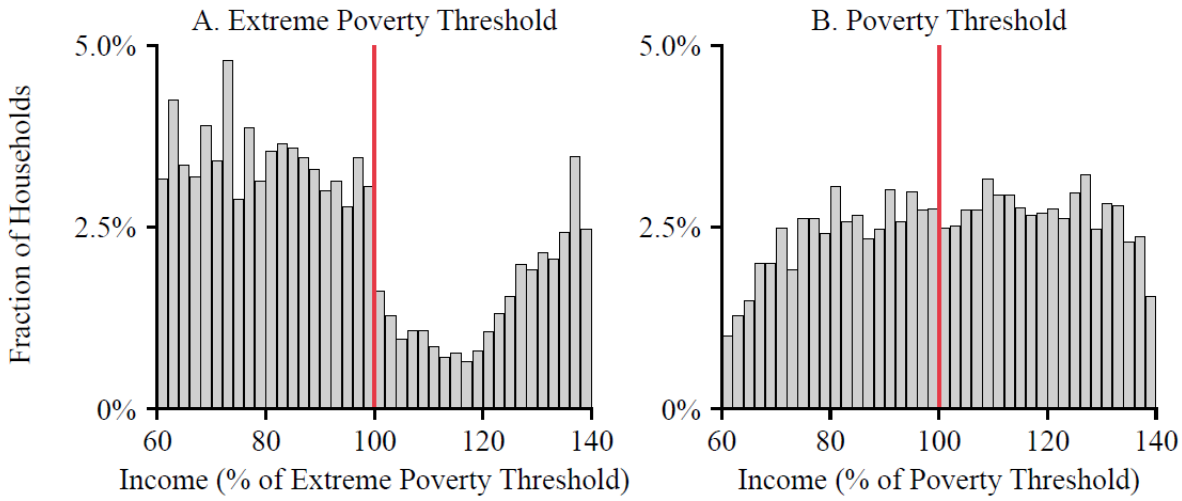


Figure 2: **Model Fit: Survival Levels and the Adjusted Gap**

Panel A plots survival by business age for beneficiary-founded and other businesses, in the model and in the registry data. Panel B plots the adjusted survival gap, the beneficiary coefficient from linear probability survival regressions with industry × initial equity × founding-year fixed effects, in the data and from the identical regression on the model's simulated panel. Model series are computed at the estimated parameters of Table 5.

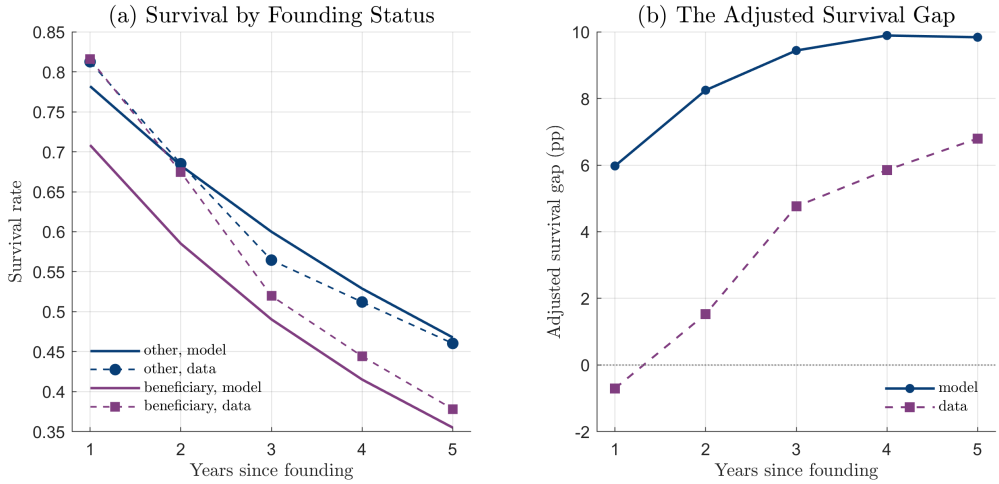


Table 1: **Summary Statistics**

This table compares beneficiary-founded sole proprietorships and their founders with the rest of the sample. The sample includes 2,296,485 sole proprietorships founded between 2014 and 2021 by first-time entrepreneurs with formal employment within twelve months of founding, of which 76,219 are beneficiary-founded. Firm equity is initial equity capital at registration in thousands of constant 2024 reais. Survival, growth (MEI to ME transition), employment creation, tax violation, and active debt collection are measured five years post-founding and reported in percent. Founder characteristics are measured at founding; the pre-entry salary is monthly, in constant 2024 reais. Founder-level differences are significant at the 1 percent level by t-tests. Data are from CNPJ, RAIS, MDS, and PNAD Contínua.

	Beneficiary founders (1)	Other founders (2)	Difference (3)
Panel A: Businesses			
Initial equity (R$ thousands)	2.1	3.2	−1.1
Survival, 1 year (%)	81.6	81.3	+0.3
Survival, 5 years (%)	37.8	46.0	−8.2
MEI → ME, 5 years (%)	4.6	6.8	−2.2
Employment creation, 5 years (%)	1.9	4.3	−2.4
Tax violation, 5 years (%)	23.5	15.1	+8.4
Active debt collection, 5 years (%)	9.7	8.3	+1.4
Panel B: Founders			
Age at founding	35.1	35.3	−0.1
Female (%)	86.5	38.8	+47.7
White (%)	51.1	63.1	−12.0
High school complete (%)	64.9	59.9	+5.1
College complete (%)	2.0	12.6	−10.6
Pre-entry monthly salary (R$)	488	861	−373
Service and sales occupation (%)	50.1	26.4	+23.7
Businesses (thousands)	76.2	2,220.3	

Table 2: **Bunching at Eligibility Thresholds**

This table reports excess mass in the distribution of pre-transfer income of Bolsa Família beneficiary households at the two eligibility thresholds. Pre-transfer income is reported household per capita income minus expected Bolsa Família payments given family composition and program rules. Income is normalized to put the threshold at 100, and each panel restricts to households within the stated exclusion window. Excess mass b is measured relative to a polynomial counterfactual fitted on 2-percentage-point bins. The implied elasticity is $e = b/(1 + b)$ and is not reported when excess mass is statistically insignificant. Confidence intervals are symmetric, $b \pm 1.96\,\mathrm{SE}$, with standard errors from a bootstrap with 1,000 replications clustered at the household level. Crossing the extreme-poverty threshold forfeits the entire basic benefit; crossing the poverty threshold costs only variable benefits under the two-year protection rule. Data are from PNAD Contínua annual microdata, 2013–2019. *, **, and *** denote significance at the 10%, 5%, and 1% level.

Window	Excess mass b	Elasticity e	t-stat	95% CI	Obs.
Panel A: Extreme-poverty threshold (R$89)					
±8%	0.197^{***}	0.165	2.91	[0.064, 0.331]	3,709
±10%	0.232^{***}	0.189	2.94	[0.077, 0.387]	4,507
±12%	0.238^{***}	0.192	2.60	[0.058, 0.417]	5,277
Panel B: Poverty threshold (R$178)					
±8%	0.007	—	0.25	[−0.049, 0.064]	11,973
±10%	0.034	—	1.08	[−0.028, 0.095]	14,744
±12%	0.045	—	1.26	[−0.025, 0.115]	17,436

Table 3: **Calibrated Parameters**

This table lists the parameters set outside the estimation. The median formal wage is the numeraire, and all monetary magnitudes are therefore in units of the median wage. Benefit levels reproduce the 2013–2021 Bolsa Família payment ledgers; the thresholds place the bottom quarter of the modeled skill distribution inside eligibility, matching program coverage. The skill process is fitted to the RAIS pre-founding worker panel (72.1 million person-years, wages winsorized at the 1st and 99th percentiles). The concealment cost is calibrated because the criterion is flat in ϕ; see Section 5.1.

Parameter		Value	Source
Preferences and technology			
Risk aversion	σ	2.0	standard
Discount factor	β	0.96	standard
Capital share	α	0.33	standard
Depreciation	δ	0.10	standard
Risk-free rate (gross)	R_f	1.04	small open economy
Collateral fraction	θ	0.40	
Transfer schedule (units of median wage)			
Basic benefit	T^{basic}	0.06	payment ledgers 2013–2021
Variable benefit	T^{var}	0.02	payment ledgers 2013–2021
Extreme-poverty threshold	$\bar{y}^{EP}$	0.70	program coverage
Poverty threshold	$\bar{y}^{P}$	0.90	program coverage
Skill process and concealment			
Skill persistence	ρ_e	0.732	RAIS pre-founding panel
Skill innovation SD	σ_e	0.232	RAIS pre-founding panel
Experience drift	χ_E	0.05	
Concealment cost	ϕ	2.0	calibrated, criterion flat

Table 4: **Simulated Moments Estimation**

This table reports the eighteen moments the estimation targets, in the data and in the model at the estimated parameters of Table 5. Survival levels are raw survival rates of beneficiary-founded and other businesses by age, in percent. Adjusted gaps are beneficiary coefficients from linear probability survival regressions with industry × initial equity × founding-year fixed effects and no founder controls, in percentage points; the model runs the identical regression on its simulated panel. The survival levels and adjusted gaps are from the registry analysis documented in Campello et al. (2026). The excess mass is measured in the ±10% window at the extreme-poverty threshold and is retained from an earlier vintage of the bunching estimates; it receives essentially no weight in the criterion, the concealment cost is calibrated, and bunching enters the paper as a qualitative signature through the threshold contrast and the Auxílio falsification of Section 3.2 (see Section 5.1). Inference on the dependency channel uses the distance test of Section 5.4.

Moment	Data	Model
Entry rate (annual, %)	10.3	9.1
Exit rate (annual, %)	10.0	10.9
Excess mass at $\bar{y}^{EP}$	0.29	0.13
Beneficiary survival, 1 yr (%)	81.6	70.8
Beneficiary survival, 2 yr (%)	67.5	58.5
Beneficiary survival, 3 yr (%)	52.0	49.0
Beneficiary survival, 4 yr (%)	44.4	41.5
Beneficiary survival, 5 yr (%)	37.8	35.5
Non-beneficiary survival, 1 yr (%)	81.3	78.2
Non-beneficiary survival, 2 yr (%)	68.5	68.3
Non-beneficiary survival, 3 yr (%)	56.5	60.0
Non-beneficiary survival, 4 yr (%)	51.2	52.9
Non-beneficiary survival, 5 yr (%)	46.0	46.8
Adjusted gap, 1 yr (pp)	–0.7	6.0
Adjusted gap, 2 yr (pp)	1.5	8.2
Adjusted gap, 3 yr (pp)	4.8	9.4
Adjusted gap, 4 yr (pp)	5.9	9.9
Adjusted gap, 5 yr (pp)	6.8	9.8

Table 5: **Estimated Parameters**

This table reports simulated method of moments estimates of the nine second-tier parameters, with the dependency channel free (left) and shut down (right). Standard errors in parentheses are computed from the numerical Jacobian at the estimates with the moment variance–covariance matrix Ω; the J statistic is the criterion at the minimum and is distributed χ^2 with degrees of freedom equal to the number of moments minus the number of estimated parameters. With administrative samples in the millions the J statistic rejects mechanically, and inference on the dependency channel uses the distance test $D = J_{\kappa=0} - J_{\kappa\,free} = 251.5 \sim \chi^2(1)$; see Section 5.4.

	κ free		$\kappa = 0$	
Parameter	$\hat{\theta}$	(SE)	$\hat{\theta}$	(SE)
Entry cost F_e	0.773	(0.043)	0.697	(0.010)
Quality location μ_q	0.109	(0.064)	0.140	(0.030)
Quality dispersion σ_q	3.812	(0.560)	3.877	(0.284)
Hazard intercept γ_0	−1.791	(0.319)	−1.672	(0.301)
Hazard quality slope γ_q	0.501	(0.055)	0.567	(0.144)
Hazard productivity slope γ_z	2.710	(0.361)	2.573	(0.655)
Dependency κ	0.284	(0.052)	0	(imposed)
Productivity persistence ρ_z	0.416	(0.009)	0.420	(0.013)
Productivity innovation σ_z	0.899	(0.186)	0.921	(0.055)
J	6,257.5	$[\chi^2(9)]$	6,509.0	$[\chi^2(10)]$

Table 6: **Out-of-Sample Test: Benefit Receipt among Surviving Founders**

This table compares the benefit-receipt profile of surviving beneficiary founders in the data with its model counterpart. The data column is the balanced-cohort share still drawing Bolsa Família at each business age, computed from the payment ledgers; business age four is not available in the balanced build. The model column is the same share computed from the enrollment extension of Section 5.7 at the estimates reported there. The audit rate of the extension is calibrated to the average year-over-year continuation of this profile; the receipt levels do not enter the moment vector of Table 4 and are untargeted.

Still drawing the benefit (%)	Data	Model
At business age 1	76.8	73.1
At business age 2	64.9	60.9
At business age 3	52.2	51.1
At business age 5	35.3	36.3

Table 7: **Counterfactual Experiments**

This table collects the counterfactual experiments of Section 6, computed at the sealed estimates of Table 5; the resolved estimates of Section 5.7 requantify them for the next draft. Each experiment changes one feature of the transfer schedule, holding all parameters fixed, and recomputes the stationary equilibrium and the cohort survival curves. The share closed is the decline in the five-year beneficiary survival gap relative to baseline; on the regression-adjusted definition the dependence freeze moves the gap from 10.1 to 0.5 percentage points (Section 6.1). For the income-test suspension the range across grids and specifications is in brackets, and the suspension pays a transfer of equal generosity without an income test. The taper entry cuts the withdrawal rate to one quarter of its estimated value, moving the gap from 11.54 to 11.46 percentage points, and 0.08 of 11.54 is the 0.7 percent shown. Bunching entries describe the excess mass at $\bar{y}^{EP}$ after the change. Experiments hold the wage and the interest rate at observed levels and are local to the estimated schedule; the decoupling and equal-cost comparisons of Section 6.3 are qualitative in this draft.

Experiment	Share of gap closed (%)	Additional spending	Bunching
Dependence margin frozen ($\kappa = 0$)	92	—	unchanged
Income test suspended (Auxílio)	18 [15.9–18.2]	equal generosity	≈ 0
Linear taper, quarter withdrawal rate	0.7	+125%	reduced
Benefit floor preserved, rate unchanged	0	—	—

Internet Appendix for
‘‘Cash Transfers and Dependence in Entrepreneurial Finance’’

I Data and Sample **IA-2**

I.1 Institutional background of Bolsa Família IA-2

I.2 The MEI and ME regimes IA-2

II The Eligibility Notch in Practice **IA-3**

III Additional Empirical Results and Robustness **IA-3**

III.1 Bunching when the income test was suspended IA-3

III.2 Non-beneficiary households as a placebo IA-4

III.3 Empirical robustness IA-4

IV Proofs **IA-4**

IV.1 Proof of Proposition 1 . . . . . IA-5

IV.2 Proof of Proposition 2 . . . . . IA-7

V Quantitative Appendix **IA-10**

V.1 Transition kernels and the stationary distribution IA-10

V.2 Numerical solution and convergence IA-11

V.3 SMM implementation IA-12

V.4 Estimation robustness IA-13

I Data and Sample

This section collects the institutional and data material behind Section 2 of the main text. It describes the Bolsa Família program and the business regimes in the detail the main text compresses, documents the administrative sources and how they are linked at the individual level, and walks through the construction of the estimation sample filter by filter.

I.1 Institutional background of Bolsa Família

Bolsa Família began in 2004, consolidating four earlier federal transfer programs into a single conditional cash transfer. The program reaches more than 17 million families, over a quarter of Brazilian households, at an average payment of R$170 per month and a fiscal cost of roughly half a percent of GDP. Eligibility runs through the Unified Registry (CadÚnico), operated by municipal governments, where a household declares its composition and income. Enrollment carries conditions, school attendance of 85 percent for children aged 6–15 and 75 percent for adolescents aged 16–17, growth monitoring and health check-ups for young children, and prenatal and postnatal care, monitored by municipal operators and schools. Households must keep the registry current, including their income, and errors or failure to update can suspend the benefit.

The benefit structure has been stable since inception while its levels adjusted for inflation. The extreme-poverty threshold rose from R$50 per capita monthly in 2004 to R$89 in 2018, with the poverty threshold maintained at exactly twice the extreme-poverty level throughout. The basic benefit rose from R$50 to R$89, the variable benefit per child aged 0–15 from R$15 to R$41, and the adolescent benefit reached R$48. Since 2010 a two-year protection rule (*regra de proteção*) has allowed a household whose income rises above the poverty threshold to keep its benefits for up to two years, provided income stays below half the minimum wage. No such protection exists at the extreme-poverty threshold.

Three features matter for our design. Receipt is persistent, beneficiaries in the sample draw transfers for an average of 37 months, and income enters the registry by declaration at registration and updates rather than continuously, which is the administrative persistence documented in Section 2.3. The declared income of the self-employed has no third-party counterpart, which is the reporting asymmetry the model builds on. And the schedule's two cliffs carry very different stakes, the anatomy Section II walks through in reais.

I.2 The MEI and ME regimes

Brazil's simplified business registration system (SIMPLES Nacional) organizes small business taxation in tiers. The individual microentrepreneur regime (MEI) is the entry tier, open to sole proprietorships with annual revenue up to R$81,000 and at most one employee. The proprietor pays a fixed monthly

amount, about R$66.60, less than one percent of the revenue ceiling, which also buys access to the public pension system, including retirement, disability, and maternity benefits at the minimum wage. Compliance is a single simplified annual declaration (DASN-SIMEI), filed without an accountant and checked against no third-party statement. A business that outgrows the ceiling migrates to the microenterprise regime (ME), with revenue up to R$360,000, nine employees in commerce and services or nineteen in manufacturing, progressive taxation under SIMPLES from 4 to 15.5 percent of gross revenue, and monthly filings that in practice require an accountant. Beyond the ME tier, small enterprises (EPP) may earn up to R$4.8 million at rates from 9.5 to 33 percent, and businesses above that ceiling leave SIMPLES for the standard corporate regimes.

Two features matter for our design. Registration as an MEI is free and online, and the estimated entry cost is therefore not a registration fee but the outlay of premises, inventory, and licenses. And the MEI-to-ME migration is recorded in the firm registry as a legal-form event, which is the growth margin of Section 3.1.

II The Eligibility Notch in Practice

The mechanics of the notch are easiest to see in a ledger. Consider a street vendor whose household has two children aged 0–15, facing the 2018 schedule of Section 2.2, a basic benefit of R$89 below the extreme-poverty line and variable benefits of R$41 per child up to the poverty line. Table IA.1 follows the household's total resources as the vendor's reported per capita income rises.

[Place Table IA.1 about here]

The ledger makes three points. Earning more can leave the household with less, a vendor whose income rises from R$89 to R$90 loses R$88 of total resources, and the loss is not recovered until income reaches R$178. Every income in between is strictly dominated by reporting R$89, which is the dominated range of Kleven and Waseem (2013) and the reason reported incomes pile up just below the threshold. And the same arithmetic that rewards under-reporting rewards staying small, an entrepreneur weighing an expansion that would carry reported income across the line faces an effective marginal tax above one hundred percent over the entire band.

III Additional Empirical Results and Robustness

III.1 Bunching when the income test was suspended

During 2020–21 the Auxílio Emergencial program replaced Bolsa Família with emergency payments of up to R$600 per household per month, and the payments did not depend on reported income near

the eligibility thresholds. The episode removes the return to income manipulation while leaving the population and the survey instrument unchanged, and it therefore isolates the design channel from any fixed trait of the households.

Re-estimating the bunching specification of Section 3.2 on the 2020–21 survey rounds, the excess mass at the extreme-poverty threshold collapses. Point estimates range from −0.05 to 0.14 across exclusion windows, and none is statistically significant, against estimates of 0.20 to 0.24, all significant, in the 2013–2019 rounds. The spike visible in the pre-period income distribution disappears entirely. Table IA.2 reports the estimates and Figure IA.1 plots the distributions.

The same households, facing no income test, stopped locating just below the threshold. This is the falsification behind the design interpretation in Section 6.2, manipulation tracks the schedule, not the population.

III.2 Non-beneficiary households as a placebo

Households outside Bolsa Família face no incentive at either threshold, and their income distribution provides a second placebo. Applying the same estimator to non-beneficiary households, in both the 2013–2019 and the 2020–21 periods, yields no excess mass at either threshold, and the formal estimates are negative throughout. The only visible irregularities in the non-beneficiary distribution sit at 111 to 117 percent of the extreme-poverty threshold, corresponding to round incomes of about R$100, well above the policy cutoff and unrelated to it. Table IA.3 reports the estimates and Figure IA.2 plots the distributions.

Bunching appears where and when the program creates discontinuous stakes, and nowhere else, which is the pattern the notch framework of Kleven and Waseem (2013) predicts.

III.3 Empirical robustness

The balanced-cohort receipt profile that disciplines the enrollment extension appears in Table 6, with its conditioning stated there. Alternative error-clustering schemes for the survival regressions and alternative exclusion windows and bin widths for the bunching estimates are being assembled by the data team and join this section in the next revision.

IV Proofs

This section proves the two propositions of Section 4, the worker's Euler equation with the entry option (Proposition 1) and the entrepreneur's capital condition with the means-test wedge (Proposition 2). The maintained assumptions are stated once, (A1) through (A3) with the first proof and (A4) and (A5) with

the second, and each proof flags the one step where the economics does the work, the entry margin drops out of the worker's Euler equation because the two occupations are worth the same at the switching point, and the concealment response drops out of the capital condition because concealment is already optimal.

IV.1 Proof of Proposition 1

Throughout the proof, (a, e, q) is the worker's current state, and all derivatives with respect to wealth are denoted by μ with the appropriate superscript, $\mu^W \equiv \partial V^W/\partial a$, $\mu^E \equiv \partial V^E/\partial a$, and $\bar{\mu}^E \equiv \mathbb{E}_z\, \mu^E$. We maintain three assumptions, standard in this class of models.

(A1) The conditional distribution $F(e' \mid e)$ admits a density $f(e' \mid e)$; in particular it places no atom on the set where the entry surplus is exactly zero.

(A2) $V^W(\cdot, e, q)$ and $V^E(\cdot, e, q, z)$ are differentiable in wealth on the interior of its domain, with derivatives bounded on compact sets. Differentiability of the value functions in this class follows from the envelope arguments of Benveniste and Scheinkman (1979) and, without concavity, of Milgrom and Segal (2002).

(A3) u is CRRA and satisfies the Inada condition $u'(c) \to \infty$ as $c \to 0$.

Step 1: restatement. Define cash on hand

$$\omega \;\equiv\; we + R_f a + T(we) - \tau, \tag{IV.1}$$

and the budget constraint then reads $c = \omega - a'$, with the worker's problem a one-dimensional choice of savings:

$$V^W(a, e, q) \;=\; \max_{a' \in [0,\, \omega]} \Big\{ u(\omega - a') \;+\; \beta\, \mathbb{E}_{e'|e}\, W(a', e', q) \Big\}, \tag{IV.2}$$

where

$$W(a', e', q) \;\equiv\; \max\left\{ V^W(a', e', q),\; \bar{V}^E(a', e', q) - F_e \right\} \tag{IV.3}$$

is the continuation value after next period's occupational choice, and $S(a', e', q) = \bar{V}^E(a', e', q) - F_e - V^W(a', e', q)$ is the entry surplus of equation (7).

Step 2: the continuation value is differentiable in savings, and its derivative ignores the entry margin. Fix (e, q). As shown in the text, the surplus is strictly decreasing in e' over the relevant range, and entry is therefore a threshold event, $S(a', e', q) \geq 0$ if and only if $e' \leq \bar{e}(a', q)$, where the threshold is defined by $S(a', \bar{e}(a', q), q) = 0$ when an interior crossing exists. Split the expectation at the threshold:

$$\mathbb{E}_{e'|e}\, W(a', e', q) = \int_{-\infty}^{\bar{e}(a', q)} \left[\bar{V}^E(a', e', q) - F_e \right] f(e' \mid e)\, de' + \int_{\bar{e}(a', q)}^{\infty} V^W(a', e', q)\, f(e' \mid e)\, de'. \tag{IV.4}$$

Both integrands are differentiable in a' by (A2), dominated on compacts, and the limits of integration are differentiable in a' wherever $\bar{e}$ is. The Leibniz rule gives

$$\frac{\partial}{\partial a'}\,\mathbb{E}_{e'|e}W = \int_{-\infty}^{\bar{e}} \bar{\mu}^E(a',e',q)\,f(e'\mid e)\,de' + \int_{\bar{e}}^{\infty} \mu^W(a',e',q)\,f(e'\mid e)\,de' + \frac{\partial \bar{e}}{\partial a'}\,f(\bar{e}\mid e)\Big[\big(\bar{V}^E(a',\bar{e},q)-F_e\big)-V^W(a',\bar{e},q)\Big]. \tag{IV.5}$$

The bracketed term is $S(a',\bar{e}(a',q),q)$, which is zero by the definition of the threshold. The boundary term therefore vanishes, the worker's future occupational switch has no first-order effect on the value of saving, because at the switching point the two occupations are worth exactly the same. The remaining derivative is

$$\frac{\partial}{\partial a'}\,\mathbb{E}_{e'|e}W = \mathbb{E}_{e'|e}\Big[\mathbb{I}\{\text{enter}\}\,\bar{\mu}^E(a',e',q) + \mathbb{I}\{\text{stay}\}\,\mu^W(a',e',q)\Big]. \tag{IV.6}$$

Step 3: conditional-expectation form. By the definition of the entry probability, $p(a',e,q) = F(\bar{e}(a',q)\mid e)$, and by the definition of conditional expectation,

$$\mathbb{E}\big[\mathbb{I}\{\text{enter}\}\,\bar{\mu}^E\big] = p\cdot\mathbb{E}\big[\bar{\mu}^E\,\big|\,\text{enter}\big], \qquad \mathbb{E}\big[\mathbb{I}\{\text{stay}\}\,\mu^W\big] = (1-p)\cdot\mathbb{E}\big[\mu^W\,\big|\,\text{stay}\big]. \tag{IV.7}$$

Step 4: first-order condition. Let $\varphi(a') \equiv u(\omega-a') + \beta\,\mathbb{E}_{e'|e}W(a',e',q)$ denote the objective in (IV.2). By (A2) and Step 2, φ is differentiable on $(0,\omega)$ with

$$\varphi'(a') = -\,u'(\omega-a') + \beta\,\frac{\partial}{\partial a'}\,\mathbb{E}_{e'|e}W. \tag{IV.8}$$

The corner $a'=\omega$ is never optimal, it implies $c=0$ and $\varphi'(\omega^-)=-\infty$ by (A3). If the optimum is interior, $\varphi'(a')=0$; substituting (IV.6) and (IV.7) and writing $c=\omega-a'$ delivers equation (9):

$$u'(c) = \beta\Big\{(1-p)\,\mathbb{E}_{e'|e}\big[\mu^W(a',e',q)\,\big|\,\text{stay}\big] + p\,\mathbb{E}_{e'|e}\big[\bar{\mu}^E(a',e',q)\,\big|\,\text{enter}\big]\Big\}. \tag{IV.9}$$

If instead the no-borrowing constraint binds, $a'=0$, optimality requires $\varphi'(0)\le 0$, which is the same expression with ''$\ge$'' in place of equality. This exhausts the two cases and establishes the proposition.

Step 5: the envelope condition used in the text. Current wealth enters φ only through cash on hand, with $\partial\omega/\partial a = R_f$. Let $g(a)$ denote the optimal savings policy. Differentiating $V^W(a,e,q) = u(\omega(a)-g(a)) + \beta\,\mathbb{E}\,W(g(a),e',q)$ with respect to a,

$$\mu^W(a,e,q) = u'(c)\,R_f + g'(a)\Big[-u'(c) + \beta\,\tfrac{\partial}{\partial a'}\mathbb{E}\,W\Big]. \tag{IV.10}$$

At an interior optimum the bracket is zero by Step 4; at a binding corner the policy is locally constant,

$g'(a) = 0$. In both cases

$$\mu^W(a, e, q) = R_f\, u'(c), \tag{IV.11}$$

which is the envelope condition invoked in the footnote to Proposition 1, with $p = 0$ it reduces the Euler equation to the textbook $u'(c) = \beta R_f\, \mathbb{E}\, u'(c')$. ■

Remark. The threshold form in Step 2 uses the monotonicity of the surplus in e'. If the surplus is not monotone, the same conclusion follows from (A1) and (A2) directly, W is the maximum of two differentiable functions and is therefore locally Lipschitz in a' and differentiable except on the crossing set $\{S = 0\}$, which has F-measure zero under (A1); dominated convergence then justifies differentiating under the expectation, and (IV.6) holds as stated.

IV.2 Proof of Proposition 2

The proof uses assumptions (A2) and (A3) of Section IV.1 and two additional ones.

(A4) The transfer schedule T is continuously differentiable in reported income. This is the smoothed schedule of the quantitative implementation; at the notch itself the schedule is a steep smooth step, and T' denotes its slope.

(A5) The exit hazard π_d is a function of the current state (a, e, q, z) only, as specified in the timing of Section 4, and is therefore taken as given when (k', m) are chosen.

Step 1: restatement. Substitute the collateral constraint at equality, $b' = \theta k'$, into the budget constraint and the law of motion. Writing $\pi = \pi(k', z, q)$ for current profit and $\tilde{y}_E = (1 - m)\pi$ for reported income, consumption and next-period net worth become functions of the two choices:

$$c(k', m) = a + T\big((1 - m)\pi\big) - \tau - \tfrac{\phi}{2} m^2 \pi + \left(\tfrac{\theta}{R_f} - 1\right) k', \tag{IV.12}$$

$$a'(k') = \pi + (1 - \delta - \theta)\, k'. \tag{IV.13}$$

The entrepreneur's problem is

$$\max_{k' \geq 0,\, m \in [0,1]} \Phi(k', m) \equiv u\big(c(k', m)\big) + \beta\Big[(1 - \pi_d)\, \mathbb{E}_{z'|z} V^E(a'(k'), e^+, q, z') + \pi_d\, V^W(a'(k'), e^+, q)\Big]. \tag{IV.14}$$

Three facts simplify what follows. By (A5), π_d does not depend on (k', m). The skill drift e^+ is deterministic and does not depend on (k', m). And next-period net worth (IV.13) does not depend on m, concealment moves reported income, not resources carried forward.

Step 2: derivatives of the budget objects. Write $\pi_k \equiv \partial\pi/\partial k'$. Differentiating (IV.12) and

(IV.13),

$$\frac{\partial c}{\partial k'} = T'\big(\tilde{y}_E\big)\,(1-m)\,\pi_k - \tfrac{\phi}{2}m^2\,\pi_k + \tfrac{\theta}{R_f} - 1, \tag{IV.15}$$

$$\frac{\partial c}{\partial m} = -\,\pi\left[T'\big(\tilde{y}_E\big) + \phi\, m\right], \tag{IV.16}$$

$$\frac{\partial a'}{\partial k'} = \pi_k + 1 - \delta - \theta, \qquad \frac{\partial a'}{\partial m} \;=\; 0. \tag{IV.17}$$

Equation (IV.15) collects the three ways one more unit of capital touches the current budget, it raises reported income and so moves the transfer by $T'\cdot(1-m)\pi_k$, it scales the concealment cost by $\frac{\phi}{2}m^2\pi_k$, and it costs the down-payment $1-\theta/R_f$. Equation (IV.16) uses $\partial\tilde{y}_E/\partial m = -\pi$ and $\partial\big(\frac{\phi}{2}m^2\pi\big)/\partial m = \phi m\pi$.

Step 3: derivative of the continuation value. By (A2) the value functions are differentiable in wealth and their derivatives are dominated on compacts, and the derivative therefore passes under the (finite-state) expectation over z'. Using (A5), the chain rule through a' gives

$$\frac{\partial}{\partial k'}\,\beta\Big[(1-\pi_d)\,\mathbb{E}_{z'|z}V^E + \pi_d V^W\Big] = \beta\,\bar{\mu}\,\frac{\partial a'}{\partial k'}, \qquad \bar{\mu} \equiv (1-\pi_d)\,\mathbb{E}_{z'|z}\,\mu^E(a', e^+, q, z') + \pi_d\,\mu^W(a', e^+, q), \tag{IV.18}$$

which is the average marginal value of net worth in equation (21). Because $\partial a'/\partial m = 0$, the continuation value does not depend on m at all.

Step 4: the concealment rule. Differentiate (IV.14) in m. By Step 3 only the current-period term contributes:

$$\frac{\partial \Phi}{\partial m} = u'(c)\,\frac{\partial c}{\partial m} = -\,u'(c)\,\pi\left[T'\big(\tilde{y}_E\big) + \phi\, m\right]. \tag{IV.19}$$

Since $u'(c) > 0$ and $\pi > 0$ for an operating business, an interior optimum $m \in (0,1)$ requires the bracket to vanish,

$$m^* = -\,\frac{1}{\phi}\,T'\big(\tilde{y}_E\big), \tag{IV.20}$$

which is the concealment rule (13). At the corner $m = 0$, optimality requires $\partial\Phi/\partial m \le 0$, that is $T' \ge 0$, no concealment where the schedule does not fall; at $m = 1$ it requires $T' \le -\phi$. The truncation of the rule to $[0,1]$ stated in the text covers all three cases.

Step 5: the capital condition. The choice of k' affects Φ both directly and through the optimal concealment $m^*(k')$, which varies with k' because reported income does. The total derivative is

$$\frac{d\Phi}{dk'} = \frac{\partial \Phi}{\partial k'} + \left.\frac{\partial \Phi}{\partial m}\right|_{m=m^*} \cdot \frac{dm^*}{dk'}. \tag{IV.21}$$

At an interior m^* the second factor's first term is zero by (IV.19), and at a strict corner m^* is locally constant in k', and $dm^*/dk' = 0$. In either case the second term vanishes, the response of concealment to capital contributes nothing to first order, which is the envelope observation made after the proposition,

Δ_T holds m fixed. The condition for k' therefore uses the partial derivative only:

$$\frac{\partial \Phi}{\partial k'} = u'(c)\,\frac{\partial c}{\partial k'} + \beta\,\bar{\mu}\,\big(\pi_k + 1 - \delta - \theta\big) \;=\; 0 \tag{IV.22}$$

at an interior optimum. Substituting (IV.15) and collecting the transfer and concealment terms into

$$\Delta_T \;\equiv\; \Big[-\,T'\big(\tilde{y}_E\big)\,(1-m) + \tfrac{\phi}{2}m^2\Big]\,\pi_k, \tag{IV.23}$$

which is nonnegative because the schedule is non-increasing, $T' \le 0$, and $m \le 1$, the first-order condition reads

$$u'(c)\left(1 - \tfrac{\theta}{R_f} + \Delta_T\right) = \beta\,\bar{\mu}\,\big(\pi_k + 1 - \delta - \theta\big), \tag{IV.24}$$

which is equation (22) with the surcharge (23). If the optimum is at the lower corner $k' = 0$, the same expression holds with ''$\ge$'' replacing equality; in the model this corner is essentially never active, an entrepreneur for whom operating is not worthwhile exits through the hazard rather than by choosing zero capital. The empirically relevant corner is the upper one. With the legal revenue ceiling of Section 4, capital is chosen from $[0, \bar{k}(z,q)]$, and at a binding ceiling the condition holds with ''$\le$'', the marginal benefit of capital still exceeds its marginal cost and the statutory limit, not the first-order condition, determines size. The share of operating businesses at the ceiling is reported with the estimation results. The corner $c = 0$ is excluded by (A3).

Step 6: the closed form. At an interior optimum the two pieces of the surcharge combine into one expression. Substituting the concealment rule $T'(\tilde{y}_E) = -\phi\, m^*$ into the definition of Δ_T,

$$\Delta_T \;=\; \Big[\phi\, m^*(1-m^*) + \tfrac{\phi}{2}(m^*)^2\Big]\pi_k \;=\; \phi\, m^*\Big(1 - \tfrac{m^*}{2}\Big)\pi_k, \tag{IV.25}$$

which is equation (24), strictly positive whenever concealment is active, $m^* \in (0,1]$. Wherever the schedule makes concealment worthwhile, the same slope raises the effective price of capital in (22), which is the growth distortion discussed in the text. Together with Steps 4 and 5 this establishes the proposition. ∎

Remark. Under the constant-returns technology of Section 4, $\pi(k', z, q) = \Omega(qzA)^{1/\alpha}k'$, the marginal product $\pi_k = \Omega(qzA)^{1/\alpha}$ does not depend on k', and all curvature in the problem comes from preferences and the value functions. Nothing in the proof uses this, the argument goes through for any differentiable $\pi(\cdot, z, q)$.

V Quantitative Appendix

This section documents the computation behind Sections 4 and 5. It states the law of motion as explicit transition kernels and describes how the stationary distribution is computed, sets out the discretization, the value function iteration, and the convergence evidence across grid refinements, including the wealth-ceiling diagnosis and the revenue-cap repair, and details the estimator, how each moment is computed on the exact stationary distribution, how the weighting matrix is built, and how the search, the standard errors, and the distance test are implemented. A final subsection collects robustness of the estimation itself.

V.1 Transition kernels and the stationary distribution

A worker's state is $s_W = (a, e, q)$ and an entrepreneur's is $s_E = (a, e, q, z)$. Four objects from the solution of the individual problems enter the law of motion, the worker's savings policy $g^W(s_W)$, the entry surplus S of equation (7) evaluated at next period's skill, the entrepreneur's next-period net worth $a'_E(s_E) = \pi + (1 - \delta - \theta)\, k'(s_E)$ implied by the capital policy, and the exit hazard $\pi_d(s_E)$. The skill of anyone who ran a business moves by the deterministic drift $e^+(e)$, and entrants draw their initial productivity from a fixed distribution G_0, uniform on the discretized support in the implementation.

The two laws of motion of Section 4, written on rectangles $A \times E \times Q \times Z$ of the state space, are

$$\begin{aligned}\mu^{E\prime}(A\times E\times Q\times Z) = &\int \big(1-\pi_d(s)\big)\,\mathbb{I}\{a'_E(s)\in A\}\,\mathbb{I}\{e^+(e)\in E\}\,\mathbb{I}\{q\in Q\}\;P_z(z,Z)\;d\mu^E(s)\\ &+\int\!\!\int \mathbb{I}\{S(g^W(s),e',q)\geq 0\}\,\mathbb{I}\{g^W(s)\in A\}\,\mathbb{I}\{e'\in E\}\,\mathbb{I}\{q\in Q\}\;G_0(Z)\;dF(e'\mid e)\,d\mu^W(s),\end{aligned} \tag{V.1}$$

$$\begin{aligned}\mu^{W\prime}(A\times E\times Q) = &\int\!\!\int \mathbb{I}\{S(g^W(s),e',q)<0\}\,\mathbb{I}\{g^W(s)\in A\}\,\mathbb{I}\{e'\in E\}\,\mathbb{I}\{q\in Q\}\;dF(e'\mid e)\,d\mu^W(s)\\ &+\int \pi_d(s)\,\mathbb{I}\{a'_E(s)\in A\}\,\mathbb{I}\{e^+(e)\in E\}\,\mathbb{I}\{q\in Q\}\;d\mu^E(s),\end{aligned} \tag{V.2}$$

where P_z is the transition kernel of the productivity process. The first term of (V.1) is the survivor operator $\mathcal{T}^E$ and the second the entrant operator $\mathcal{T}_0^E$; the terms of (V.2) are the stayer operator $\mathcal{T}^W$ and the failure operator $\mathcal{T}^{EW}$.

On the discrete state space, policies take continuous values between wealth nodes, and mass is assigned to the two neighboring nodes by the standard lottery, in the manner of Young (2010), the weight on each neighbor proportional to proximity, which preserves the mean of the assigned wealth exactly. The stationary pair (μ^W, μ^E) is computed by iterating (V.1)–(V.2) forward from a uniform initialization, normalizing total mass to one, until the sup-norm change falls below 10^{-6}.

V.2 Numerical solution and convergence

Discretization. Wealth is a log-spaced grid on $[a_{\min}, a_{\max}]$ with $a_{\min} = 0.001$ in median-wage units, about R$28. The skill and productivity processes are discretized by the method of Tauchen (1986) with a three standard deviation span. The permanent type q takes the quantile midpoints of its log-normal distribution, node i at the $(i - \frac{1}{2})/n_q$ quantile with equal weights, which respects the mass of the heavy tail where moment-matching quadrature does not. The estimation ladder uses four grids:

Grid	n_a	n_z	n_q	n_e
BASE	50	7	5	5
B	100	15	9	7
C	150	21	11	9
D (production)	200	27	13	11

Value function iteration. Savings and capital are chosen continuously, by a coarse scan followed by golden-section refinement, with an incumbent probe at the previous optimum to guard against local ripples; continuation values are evaluated by linear interpolation in wealth, and each golden-section evaluation costs a single one-dimensional interpolation. Iterations are accelerated by Howard policy-improvement sweeps, thirty per update, and terminate when the sup-norm change in values falls below 10^{-4}. Solution accuracy for discretized value function iteration of this type is analyzed by Kirkby (2017) and Santos and Peralta-Alva (2005), and the sensitivity of model conclusions to solution method by Hatchondo et al. (2010); the ladder below is our practical implementation of their advice, report the object of interest across successive refinements rather than at a single resolution.

Convergence ladder. The headline decomposition is recomputed on grids of increasing resolution, holding the permanent type dimension at its production value ($n_q = 13$) and refining the others. The dependency share of the five-year survival gap moves from 100.8 percent at BASE to 96.5, 92.4, and 90.3 percent at the B, C, and D refinements, a monotone approach with successive increments shrinking toward the production value, which is the pattern expected of a converging discretization.

The wealth ceiling and the revenue cap. An unbounded wealth grid is not innocuous in this model. With constant returns at the business level, fixed prices under the segment closure, and a heavy-tailed permanent type, the wealth accumulation of the most able entrepreneurs need not admit a stationary distribution, and on a bounded grid the missing stationarity surfaces as mass piling on the top wealth node, 85 to 90 percent of entrepreneur mass at every ceiling between 20 and 500 wage units in our diagnostics, with moments that drift as the ceiling moves. The repair is institutional rather than numerical. The businesses in the sample are bound by the legal revenue ceiling of their regime, R$360,000 for the ME tier,

about 13 times the median wage, and imposing it bounds capital at $\bar{k}(z,q) = y_{cap}\,\alpha/\pi(1,z,q)$ under constant returns. With the cap in place the top-node mass drops to zero and every reported moment is invariant to the wealth ceiling. The MEI ceiling, about 2.9 times the median wage, is deliberately not used, it is a regime boundary that a growing business crosses by migrating to ME, the migration is an outcome our empirical growth margin measures, and hard-coding it would forbid the very transition the model is asked to explain.

V.3 SMM implementation

Moment computation. All eighteen moments are computed on the exact stationary distribution rather than on a simulated panel, so the criterion carries no simulation error. The survival moments follow founding cohorts by distribution-forward iteration, the entrant measure implied by the stationary worker distribution and the entry policy is aged one year at a time under the hazard and the policies, and survival at horizon h is the surviving mass share. A founder is tagged as a beneficiary if the pre-entry worker wage we falls below the extreme-poverty threshold, mirroring the beneficiary-at-founding definition of the data, the tag is fixed at entry, and grid nodes straddling the threshold are split pro rata so that cohort composition is continuous under grid refinement.

The adjusted-gap moments replicate the data estimand. Entrant cells are formed by initial-equity band, five bands, and beneficiary status, each cell's cohort is aged forward, and the beneficiary coefficient is recovered from the linear probability projection of cell survival on the beneficiary indicator and equity-band effects, the model counterpart of the industry × equity × founding-year specification with no founder controls. Survival is linear in the initial measure, and cells sharing a regressor row therefore merge exactly and the computation costs at most ten forward passes.

Weighting. Ω is an influence-function estimate of the moment variance–covariance matrix constructed from the underlying microdata, clustered by municipality for the registry blocks, with the covariances between survival levels and adjusted gaps included, and clustered by household for the survey block; covariances across data sources are set to zero. Its condition number is about 3.3×10^3, and the weighting matrix is $W = \Omega^{-1}$. A diagonal-Ω alternative is reported in Section V.4.

Optimization. The criterion is minimized by Nelder–Mead with a simplex-spread tolerance of 10^{-2} and parameter bounds enforced by penalty. The search runs in two stages. A multistart stage draws starting points around a reference vector, each coordinate scaled by a factor between one half and two, or moved additively within ±0.5 when the reference coordinate is near zero, and runs the full search from each of at least twelve seeds on the BASE grid. A polish stage then warm-starts the best BASE points up the ladder, each leg of the chain sized to its wall-clock budget from the measured per-iteration cost,

with every leg's output saved under a distinct name and the next leg warm-started from it. Estimates are accepted when the criterion improvement between successive production-grid legs falls below one percent.

Standard errors and sensitivity. The Jacobian G of the moment vector in the parameters is computed by central differences at the estimates. Standard errors are the square roots of the diagonal of $(G'\Omega^{-1}G)^{-1}$. The sensitivity matrix of Andrews et al. (2017), $\Lambda = -(G'WG)^{-1}G'W$, maps local moment perturbations into parameter movements, and the κ row of Λ underlies the identification discussion of Section 5.3.

The distance test. The restricted model drops κ from the search vector rather than pinning it at zero by equal bounds, which would degenerate the simplex along that dimension. Both specifications use the same eighteen moments and the same W, and $D = J_{\kappa=0} - J_{\kappa\, free}$ is referred to $\chi^2(1)$.

V.4 Estimation robustness

Dropping the first-year moments. Because $\hat{\kappa}$ is identified by the early-horizon slope of the gap profile while the model's known fit failure sits in the first-year levels, we re-estimate excluding the three year-one moments.

Alternative weighting. Re-estimating with the diagonal of Ω in place of the full matrix checks that the off-diagonal structure, in particular the gap-level covariances, is not driving the estimates.

The concealment cost. The criterion is flat in ϕ across two orders of magnitude, a range of roughly 6,247 to 6,369 in criterion units, which is why ϕ is calibrated at the criterion-minimizing value rather than estimated, and why no conclusion in the paper leans on its level.

The quality distribution. The moments never pin the parametric family of q. Bounded families do not remove the pre-cap divergence, the estimator pushes the bound as it pushed σ_q, and after the revenue cap the family is second order, capped types all earn the capped scale regardless of the tail shape beyond the cap.

Taper mechanics. The cliff-to-slope counterfactual uses a piecewise-linear phase-out that pays the basic benefit in full up to $\bar{y}^{EP}$ and withdraws it linearly over a band above, so that only the withdrawal rate varies. Widening the smoothing parameter of the estimated schedule instead would also cut benefits below the threshold, at a smoothing width of 0.40 the poorest households collect only 85 percent of

the basic benefit, conflating a schedule reform with a benefit cut, which is why the taper is a separate instrument in the code.

References

Aiyagari, S. Rao, 1994, Uninsured idiosyncratic risk and aggregate saving, *Quarterly Journal of Economics* 109, 659–684.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro, 2017, Measuring the sensitivity of parameter estimates to estimation moments, *The Quarterly Journal of Economics* 132, 1553–1592.

Balboni, Clare, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil, 2022, Why do people stay poor?, *Quarterly Journal of Economics* 137, 785–844.

Bandiera, Oriana, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman, 2017, Labor markets and poverty in village economies, *Quarterly Journal of Economics* 132, 811–870.

Benveniste, Lawrence M., and José A. Scheinkman, 1979, On the differentiability of the value function in dynamic models of economics, *Econometrica* 47, 727–732.

Bianchi, Milo, and Matteo Bobba, 2013, Liquidity, risk, and occupational choices, *Review of Economic Studies* 80, 491–511.

Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin, 2011, Finance and development: A tale of two sectors, *American Economic Review* 101, 1964–2002.

Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin, 2015, Entrepreneurship and financial frictions: A macrodevelopment perspective, *Annual Review of Economics* 7, 409–436.

Buera, Francisco J., and Yongseok Shin, 2013, Financial frictions and the persistence of history: A quantitative exploration, *Journal of Political Economy* 121, 221–272.

Cagetti, Marco, and Mariacristina De Nardi, 2006, Entrepreneurship, frictions, and wealth, *Journal of Political Economy* 114, 835–870.

Campello, Murillo, Guilherme Junqueira, and Spyridon Lagaras, 2026, Entrepreneurship on a safety net: Evidence from the world's largest cash transfer program, Working paper, University of Florida and University of Illinois.

Chetty, Raj, John N. Friedman, Tøre Olsen, and Luigi Pistaferri, 2011, Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from danish tax records, *Quarterly Journal of Economics* 126, 749–804.

de Mel, Suresh, David McKenzie, and Christopher Woodruff, 2008, Returns to capital in microenterprises: Evidence from a field experiment, *The Quarterly Journal of Economics* 123, 1329–1372.

Evans, David S., and Boyan Jovanovic, 1989, An estimated model of entrepreneurial choice under liquidity constraints, *Journal of Political Economy* 97, 808–827.

Feldman, Mark, and Christian Gilles, 1985, An expository note on individual risk without aggregate uncertainty, *Journal of Economic Theory* 35, 26–32.

Gerard, François, Joana Naritomi, and Joana Silva, 2024, Cash transfers and formal labor markets: Evidence from brazil, Working paper.

Gertler, Mark, and Nobuhiro Kiyotaki, 2010, Financial intermediation and credit policy in business

cycle analysis, in Benjamin M. Friedman, and Michael Woodford, ed.: *Handbook of Monetary Economics*vol. 3 . pp. 547–599 (Elsevier).

Gottlieb, Joshua D, Richard R Townsend, and Ting Xu, 2022, Does career risk deter potential entrepreneurs?, *The Review of Financial Studies* 35, 3973–4015.

Hatchondo, Juan Carlos, Leonardo Martinez, and Horacio Sapriza, 2010, Quantitative properties of sovereign default models: Solution methods matter, *Review of Economic Dynamics* 13, 919–933.

Hennessy, Christopher A., and Toni M. Whited, 2005, Debt dynamics, *Journal of Finance* 60, 1129–1165.

Hennessy, Christopher A., and Toni M. Whited, 2007, How costly is external financing? evidence from a structural estimation, *Journal of Finance* 62, 1705–1745.

Holtz-Eakin, Douglas, David Joulfaian, and Harvey S. Rosen, 1994, Sticking it out: Entrepreneurial survival and liquidity constraints, *Journal of Political Economy* 102, 53–75.

Hombert, Johan, Antoinette Schoar, David Sraer, and David Thesmar, 2020, Can unemployment insurance spur entrepreneurial activity? evidence from france, *Journal of Finance* 75, 1247–1285.

Hopenhayn, Hugo A., 1992, Entry, exit, and firm dynamics in long run equilibrium, *Econometrica* 60, 1127–1150.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, 1995, Precautionary saving and social insurance, *Journal of Political Economy* 103, 360–399.

Huggett, Mark, 1993, The risk-free rate in heterogeneous-agent incomplete-insurance economies, *Journal of Economic Dynamics and Control* 17, 953–969.

Hurst, Erik, and Annamaria Lusardi, 2004, Liquidity constraints, household wealth, and entrepreneurship, *Journal of Political Economy* 112, 319–347.

Jermann, Urban, and Vincenzo Quadrini, 2012, Macroeconomic effects of financial shocks, *American Economic Review* 102, 238–271.

Judd, Kenneth L., 1985, The law of large numbers with a continuum of iid random variables, *Journal of Economic Theory* 35, 19–25.

Kirkby, Robert, 2017, Convergence of discretized value function iteration, *Computational Economics* 49, 117–153.

Kiyotaki, Nobuhiro, and John Moore, 1997, Credit cycles, *Journal of Political Economy* 105, 211–248.

Kleven, Henrik Jacobsen, 2016, Bunching, *Annual Review of Economics* 8, 435–464.

Kleven, Henrik J., and Mazhar Waseem, 2013, Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from pakistan, *Quarterly Journal of Economics* 128, 669–723.

Li, Shaojin, Toni M. Whited, and Yufeng Wu, 2016, Collateral, taxes, and leverage, *Review of Financial Studies* 29, 1453–1500.

Low, Hamish, and Luigi Pistaferri, 2015, Disability insurance and the dynamics of the incentive insurance trade-off, *American Economic Review* 105, 2986–3029.

Lucas, Robert E., Jr., 1978, On the size distribution of business firms, *The Bell Journal of Economics* 9, 508–523.

Midrigan, Virgiliu, and Daniel Yi Xu, 2014, Finance and misallocation: Evidence from plant-level data, *American Economic Review* 104, 422–458.

Milgrom, Paul, and Ilya Segal, 2002, Envelope theorems for arbitrary choice sets, *Econometrica* 70, 583–601.

Moffitt, Robert, 1992, Incentive effects of the U.S. welfare system: A review, *Journal of Economic Literature* 30, 1–61.

Moll, Benjamin, 2014, Productivity losses from financial frictions: Can self-financing undo capital misallocation?, *American Economic Review* 104, 3186–3221.

Nikolov, Boris, and Toni M. Whited, 2014, Agency conflicts and cash: Estimates from a dynamic model, *Journal of Finance* 69, 1883–1921.

Quadrini, Vincenzo, 2000, Entrepreneurship, saving, and social mobility, *Review of Economic Dynamics* 3, 1–40.

Saez, Emmanuel, 2010, Do taxpayers bunch at kink points?, *American Economic Journal: Economic Policy* 2, 180–212.

Santos, Manuel S., and Adrian Peralta-Alva, 2005, Accuracy of simulations for stochastic dynamic models, *Econometrica* 73, 1939–1976.

Tamkoç, M. Nazım, and Gustavo Ventura, 2025, Rules and regulations, managerial time and economic development, Policy Research Working Paper 10761 World Bank.

Tauchen, George, 1986, Finite state markov-chain approximations to univariate and vector autoregressions, *Economics Letters* 20, 177–181.

Warusawitharana, Missaka, and Toni M. Whited, 2016, Equity market misvaluation, financing, and investment, *Review of Financial Studies* 29, 603–654.

Young, Eric R., 2010, Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations, *Journal of Economic Dynamics and Control* 34, 36–41.

Figure IA.1: **Income Distributions of Beneficiaries, 2020–2021**

During 2013–2019, bunching concentrates just below the extreme-poverty threshold (Figure 1). This figure plots the same distribution for beneficiary households during 2020–21, when Auxílio Emergencial suspended the income test, and the excess mass is gone. Income is normalized to put the year-specific threshold at 100, and the vertical line marks the threshold. Data are from PNAD Contínua annual microdata.

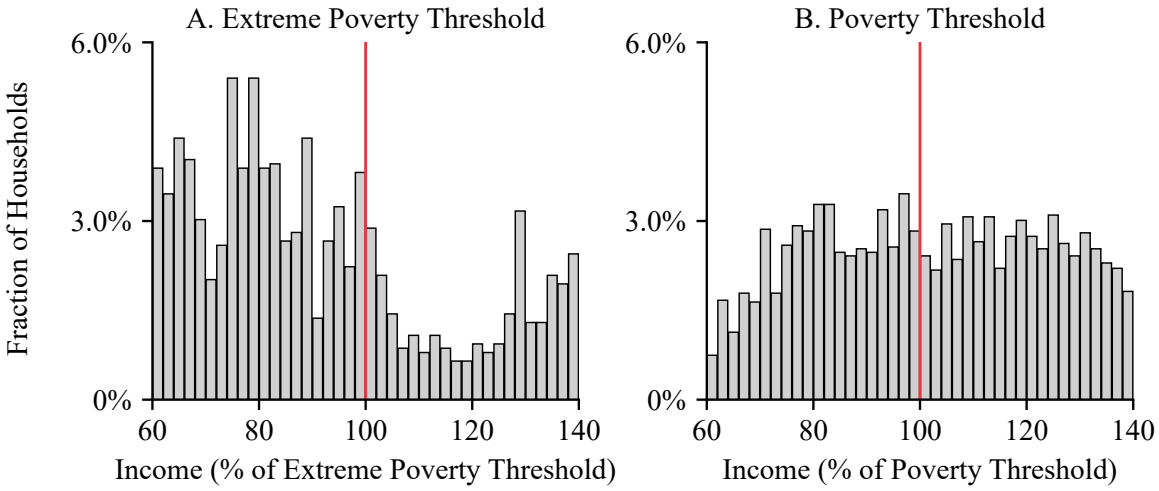


Figure IA.2: **Income Distributions of Non-Beneficiaries**

During 2013–2019, bunching concentrates just below the extreme-poverty threshold among beneficiary households (Figure 1). This figure plots households not enrolled in Bolsa Família, and no excess mass appears at either threshold in either period. Income is normalized to put the year-specific threshold at 100, and the vertical line marks the threshold. Data are from PNAD Contínua annual microdata.

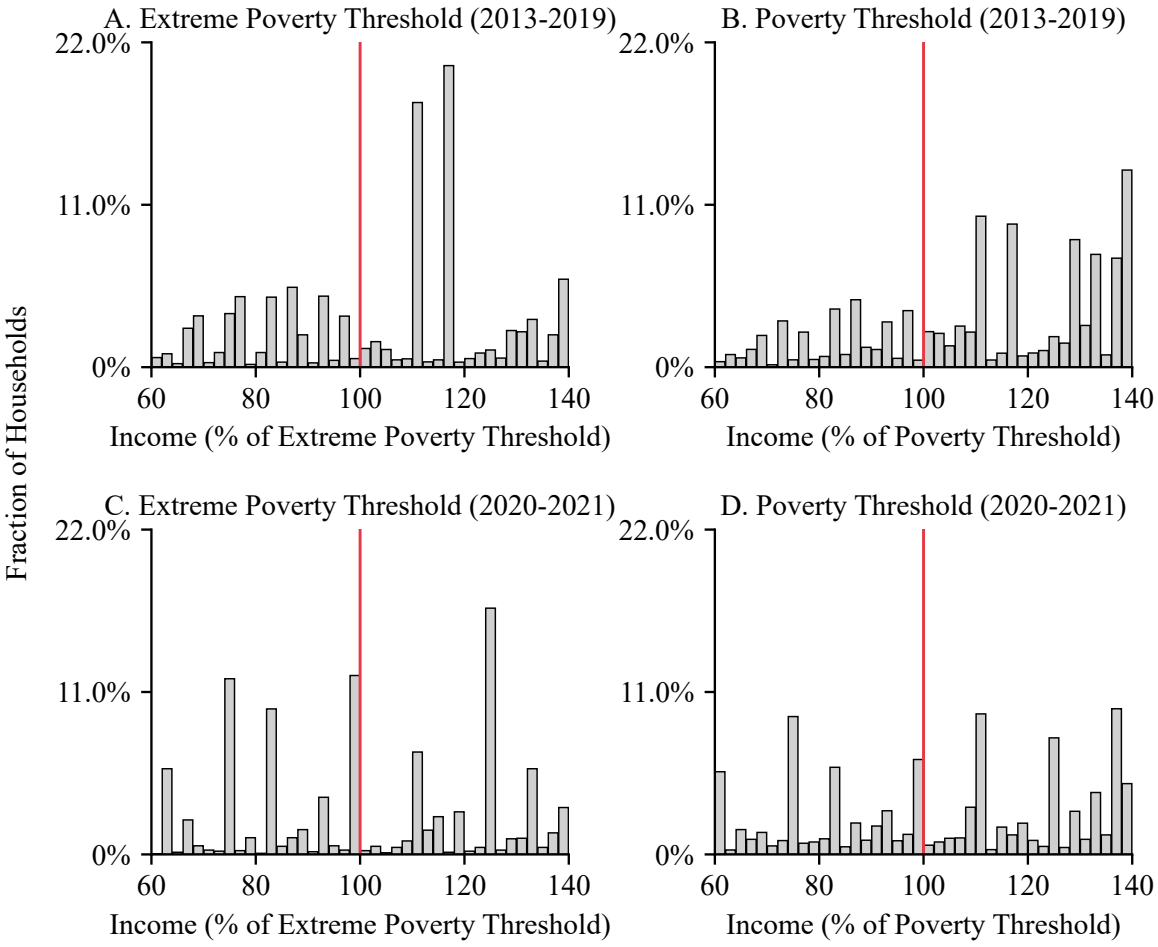


Table IA.1: **A Street Vendor Crossing the Threshold**

This table computes monthly benefits and total resources for a household with two children aged 0–15 under the 2018 Bolsa Família schedule, at different levels of reported per capita income. Below the extreme-poverty threshold (R$89) the household receives the basic benefit (R$89) plus variable benefits (R$41 per child); between R$89 and the poverty threshold (R$178) it receives the variable benefits only; above R$178 it receives nothing. Total resources are reported income plus the transfer. The final column compares total resources with those of a household reporting exactly R$89. In practice the two-year protection rule cushions the second notch at R$178; nothing cushions the first.

Reported income	Transfer	Total resources	vs. reporting R$89
70	171	241	−19
88	171	259	−1
89	**171**	**260**	**0**
90	82	172	−88
120	82	202	−58
150	82	232	−28
178	82	260	0
200	0	200	−60
260	0	260	0

Table IA.2: **Bunching at Eligibility Thresholds for Beneficiaries, 2020–2021**

This table re-estimates the excess-mass specification of Table 2 on the 2020–21 survey rounds, when Auxílio Emergencial payments did not depend on reported income near the thresholds. Income is normalized so that the threshold equals 100, and each row restricts to households within the stated exclusion window. Confidence intervals are symmetric, $b \pm 1.96$ SE, with standard errors from a bootstrap with 1,000 replications clustered at the household level. The implied elasticity is suppressed in every row because no estimate is statistically significant. Data are from PNAD Contínua annual microdata, 2020–2021.

Window	Excess mass b	t-stat	95% CI	Obs.
Panel A: Extreme-poverty threshold				
±8%	0.140	0.71	[−0.247, 0.526]	382
±10%	0.005	0.02	[−0.377, 0.386]	448
±12%	−0.054	−0.27	[−0.455, 0.347]	526
Panel B: Poverty threshold				
±8%	0.059	0.60	[−0.135, 0.253]	1,139
±10%	0.013	0.13	[−0.182, 0.207]	1,371
±12%	0.022	0.20	[−0.194, 0.238]	1,626

Table IA.3: **Bunching at Eligibility Thresholds for Non-Beneficiaries**

This table applies the excess-mass specification of Table 2 to households not enrolled in Bolsa Família, in the means-tested period (2013–2019) and in the Auxílio period (2020–2021). Income is normalized so that the threshold equals 100, and each row restricts to households within the stated exclusion window. Confidence intervals are symmetric, $b \pm 1.96$ SE, with standard errors from a bootstrap with 1,000 replications clustered at the household level. The implied elasticity is suppressed in every row. The prevailing negative estimates record a deficit of non-beneficiary mass near the thresholds, the complement of beneficiary bunching, since nearly every household close to the lines is enrolled. *, **, and *** denote significance at the 10%, 5%, and 1% level. Data are from PNAD Contínua annual microdata.

Window	Excess mass b	t-stat	95% CI	Obs.
Panel A: Extreme-poverty threshold, 2013–2019				
±8%	-0.751^{***}	−51.19	[−0.780, −0.723]	1,067
±10%	-0.846^{***}	−90.48	[−0.865, −0.828]	1,123
±12%	-0.449^{***}	−14.70	[−0.509, −0.389]	2,932
Panel B: Poverty threshold, 2013–2019				
±8%	-0.557^{***}	−41.05	[−0.583, −0.530]	5,539
±10%	-0.640^{***}	−57.68	[−0.662, −0.619]	6,688
±12%	-0.102^{***}	−2.91	[−0.170, −0.033]	11,538
Panel C: Extreme-poverty threshold, 2020–2021				
±8%	-0.334^{***}	−3.94	[−0.500, −0.168]	543
±10%	-0.550^{***}	−10.36	[−0.654, −0.446]	602
±12%	0.290	1.19	[−0.186, 0.766]	1,018
Panel D: Poverty threshold, 2020–2021				
±8%	-0.523^{***}	−17.96	[−0.580, −0.466]	1,539
±10%	-0.553^{***}	−21.55	[−0.603, −0.503]	1,999
±12%	0.108	1.16	[−0.074, 0.289]	3,263