Inflation Beliefs and the Indexation of Household Debt*

Otávio Braga† — *Duke University*

João Pugliese — *Stanford University*

Lucas Iten Teixeira — *Central Bank of Brazil*

(Job Market Paper)

November 2, 2026

(Click Here for the Latest Version)

Abstract

Economic theory dictates that inflation beliefs should shape household financial decisions, but establishing this link empirically has been a central challenge. We argue that such difficulty can be partly explained by the fact that many households do not understand their own financial exposure to inflation. To show this, we study households holding inflation-indexed debt, whose contracts establish a one-to-one relation between inflation and the growth of their nominal payments. In a survey experiment with Brazilian borrowers, we separately elicit forecasts for inflation and installments. Most respondents underestimate their own exposure, expecting payments below what their inflation forecasts imply. We cross-randomize two independent treatments informing respondents about (i) historical inflation and (ii) the payment increase implied by their own inflation forecast. We find that shifts in inflation beliefs affect households' financial planning only when paired with (ii), which makes their payment consequences explicit; (i) alone revises forecasts but leaves planning unchanged. Finally, we calibrate a consumption-savings model to show that borrowers tolerate inflation exposure in their indexed debt far more than the rational benchmark predicts, even after accounting for possibly biased inflation beliefs and the underestimation of their future debt burden.

Keywords: indexation, inflation, beliefs, household debt, RCTs.

*We are grateful to David Berger, Nuno Clara, Cosmin Ilut and Kyle Jurado for their guidance. We also thank Rafael Alves, Arthur Carvalho Brito, Bomi Kang, Matthias Kehrig, Paymon Khorrami, Andrea Lanteri, Fernando Lopes, Arnaud Maurel, Angelo Mendes, Michael Pollmann, David Robinson and Melanie Wallskog for helpful comments. We owe special thanks to Caio Castro at the FGV Cidades research center for his work on data collection and management. This study was approved by FGV CEPH (Protocol P.645.2025) and pre-registered at the AEA RCT Registry (#19465). The views expressed here are those of the authors and do not necessarily reflect those of the Central Bank of Brazil. All errors are our own.

†Corresponding author: otavio.braga@duke.edu

Introduction

There are several channels through which inflation can affect the economic decisions of a household. It changes the prices of the goods and services the household buys and, since income may adjust with a lag, it can distort the cost of consuming today relative to consuming later. It also erodes the real value of debt fixed in nominal terms, and affects the real returns on savings and investments. Benchmark models in economics and finance assume that household decision-making takes into account each of these channels, implying that beliefs about future inflation should be relevant drivers of consumption, savings, borrowing and investments. Naturally, an empirical question that arises is whether households actually do so and, more broadly, whether they understand these channels in the first place. This is distinct from questioning households' understanding of the inflation process itself, as even agents with perfect foresight still have to work out what their forecasts imply for their own budget and decisions. In fact, economists understand considerably more about the determinants of inflation beliefs than about their implications for decisions (Weber et al., 2022; D'Acunto et al., 2023b; D'Acunto and Weber, 2024; Coibion and Gorodnichenko, 2026), and results on these implications are mixed, often of modest magnitudes, and sometimes misaligned with theoretical predictions (Bachmann et al., 2015; Burke and Ozdagli, 2023; Duca-Radu et al., 2021; Coibion et al., 2023). A central challenge in this literature is obtaining precise measures of a household's exposure to inflation, which can operate through multiple components of the household budget (Binetti et al., 2024; Jaravel, 2021; Kaplan and Schulhofer-Wohl, 2017).

This paper addresses this challenge using a setting particularly well suited to it: a survey experiment with borrowers whose liabilities are indexed to inflation. Debt indexation creates a measurable link between expected inflation and expected nominal liabilities: when the underlying index rises by $x\%$ between readjustment dates, debt balances and monthly payments are revised upwards by the same $x\%$. By separately eliciting borrowers' beliefs about (i) future inflation and (ii) their own future nominal payments, we can measure how much of their inflation forecast households carry into the payments they expect. We summarize this with a pass-through measure θ, inspired by the work of Stango and Zinman (2009): a household that properly maps its own forecast into its future payments has $\theta = 1$, while $\theta < 1$ means that its projected payments understate the implications of its own inflation beliefs. We find widespread evidence of weak pass-through, indicating substantial underestimation of nominal exposure. This finding gives us the first part of our argument, that households do not understand their own exposure to inflation; the experiment then tests whether this is a reason why they fail to act upon their own inflation beliefs.

For our survey, we collected 783 quality-filtered responses from clients of a large financial firm in Brazil who hold credit contracts for real estate purchases. These are rotating finance contracts known as *consórcios*, a common alternative to mortgages in Brazil (Van Doornik et al., 2024a), and their monthly installments are periodically readjusted by an inflation in-

dex. They typically have long maturities, 17.6 years at origination on average in our sample, and the associated installments are economically meaningful, accounting for 14% of respondents' monthly income on average. In addition to demographics, financial characteristics, and detailed information about their contracts, the survey elicits respondents' beliefs about annualized inflation over the next five years and their conditional expectations of nominal monthly payments at that same horizon. Inflation beliefs were elicited under three scenarios: low, average and high. The average scenario is interpreted as the respondent's expected inflation. The other two scenarios aim to capture sentiment towards inflation risk: the high-inflation scenario was described as one in which inflation turns out substantially higher than the respondent expects, but still realistic rather than extreme, and the low-inflation scenario was defined analogously. No explicit probabilities were attached to these scenarios, so we read them as the range of inflation outcomes each respondent considers realistic.

This design allows us to conduct a randomized controlled trial (RCT) with interventions that separately target each of these two components: an *inflation treatment* (T_I) that shifts respondents' inflation beliefs, following closely the approach in the existing literature (Weber et al., 2025; Coibion and Gorodnichenko, 2026), and a *mapping treatment* (T_M) that directly targets the mapping of inflation beliefs into future nominal payments. The two treatments were cross-randomized ($T_I \perp T_M$), generating a 2×2 factorial design. Respondents assigned to the inflation treatment were shown historical benchmarks for each scenario, based on the distribution of annualized inflation in the index that adjusts their monthly installments: the historical mean was a reference for their average scenario, and the 5^{th} and 95^{th} percentiles as reference points for realistic low- and high-inflation outcomes. Their five-year inflation beliefs were measured before and after this information was displayed.[1] Respondents revised their answers towards these signals, with treatment-induced learning rates of 26% and 30% for the expected and high scenarios and 31% for upside risk, all significant at the 1% level. The mapping treatment was assigned right after the posteriors for inflation beliefs were measured, and its implementation is the main novelty in our design.

For the control group in the mapping treatment arm ($T_M = 0$), respondents were simply asked: "*if your expected scenario for annualized inflation in the next 5 years actually happens, what do you think your nominal installments will be 5 years from now?*". For the treatment group ($T_M = 1$), we used each respondent's own answers to inform her exactly how much installments would increase after five years if her expected scenario materialized. The more cognitively demanding task of compounding annualized inflation into nominal installments was thus reduced to an addition: treated respondents only had to add the displayed increase to the installment they had reported. Survey participants were very responsive to this mapping intervention and it led to a sizable difference in pass-through across groups: only 20.3%

[1] Usually, RCTs in related literature provide information about inflation forecasts. Our choice to provide historical data moments in our treatment was based on the fact that 5-year ahead professional forecasts are not publicly available and tend to be inaccurate, especially for the Brazilian context.

of control respondents have a pass-through sufficiently close to one,[2] against 53.9% in the treatment group, and 27.4% of control respondents have a pass-through below 0.5, against 17.5% of treated respondents. These patterns are qualitatively in line with previous work on exponential-growth bias in household finance (Stango and Zinman, 2009). Importantly, these results cannot be attributed to unawareness of indexation: 76% of respondents reported being aware of it before the survey, and every participant was then shown a clear explanation of the indexation rules. Under-projection is also widespread across demographic and financial profiles, and these characteristics, together with contract information and proxies for attention to the survey, explain little of the variation in θ.

We then analyze the impacts of these information treatments on financial planning outcomes. The most important results in our experiment come from the interaction between the two treatments and their effects on monthly savings and household preferences for inflation-indexed contracts. After T_I and T_M were assigned, our survey measured respondents' planned monthly savings (as a fraction of their income) and their willingness to switch to a non-indexed payment plan, that is, a conventional constant-payment contract instead of their current indexed contract. The latter question, which was hypothetical, imposed the following trade-off: a higher initial installment in exchange for payments that no longer rise with inflation. Each respondent was first asked whether they would switch to a fixed-payment contract with 20% higher installments, and then about the same switch with 10% and 30% increases.

Using a reduced-form specification that separates shifts in expected inflation from shifts in inflation risk—measured by the gap between the high and the average scenarios—we find that the inflation treatment T_I alone has no detectable effect on either outcome, even though it shifts inflation beliefs in all three scenarios. In other words, shifting beliefs about expected inflation and inflation risk did not affect individual intentions to save or their willingness to switch to a fixed-payment contract. These results control for individual perceptions of income indexation, self-reported liquidity buffers, and a long array of contract and household specific characteristics. This is at odds with what one would expect, since inflation beliefs should matter when planning savings under an indexed contract and when choosing between fixed and indexed contracts. This is especially relevant in the context of our sample, where individuals report to have significant degree of income stickiness. We interpret this as the attenuation result often seen in related literature and formalized by Yang (2025): consistent with the weak pass-through we document, individuals are not fully aware of how inflation affects their monthly payments and budgets, which prevents them from planning around the information provided by T_I.

When combined with T_M, however, T_I does shift financial planning. Since T_I moves beliefs up or down relative to the historical benchmarks, we measure treatment intensity by defining two treatment doses, one for expected inflation and one for our measure of inflation

[2] "Sufficiently close" is defined here as $|\theta - 1| \leq 0.05$.

risk. The intuition underlying our main results is simple: bad news about inflation, whether about its expected level or its risk, triggers precautionary behavior—households plan to save more and become more likely to prefer contracts without indexation. In the mapping arm ($T_M = 1$), telling a household that the historical mean of inflation is one percentage point above its own forecast raises planned saving by about 0.73 percentage points, and a comparable gap in inflation risk raises it by 0.81. The contract choice responds mostly to news about inflation risk: the same one-point gap raises the probability of paying a 30% premium to freeze installments by 3.06 percentage points, about a tenth of the 32% baseline. None of these responses appears without the mapping treatment. We again interpret this through the lens of Yang (2025): making the budget implications of inflation explicit sidesteps the low pass-through and turns inflation information into something the household can act on.

Our experiment identifies what moves these decisions, but it cannot tell us whether the resulting choices are large enough. Are individuals responding rationally to their inflation beliefs when deciding between indexed contracts and fixed-payment contracts? We answer that question with the consumption-savings model of Section 2, calibrated to each respondent's financial profile and to her own elicited beliefs. The model delivers γ_i^*, the permanent increase in the installment that would leave borrower i indifferent between her indexed contract and a fixed one, and it serves as a benchmark for what a rational household would do even under imprecise beliefs about inflation. Conditional on her own beliefs, the median respondent should be willing to pay up to 54% of her installment to fix her payments, far above any premium we offered. The model therefore predicts that 96%, 89% and 79% of respondents should accept our three offers, against 77%, 58% and 32% in the data. Correcting beliefs to the historical inflation process widens this gap rather than closing it, since the respondents whose thresholds rise are precisely those who held unusually low views on future inflation. Under-projection of installments accounts for part of the distance, but not for most of it. Replacing the contractual payment rule with each respondent's perceived one, using the θ_i we measure in the survey, lowers predicted acceptance, closing between a third and two-fifths of the gap. The model still predicts acceptance rates 13, 18 and 32 percentage points above the data.

Returning to our initial motivation on the puzzling relation between inflation beliefs and economic decisions of households, our findings suggests that economists are possibly missing less in how households form beliefs about inflation than in how households act on those beliefs. We study inflation in this debt indexation context, where it should drive financial decisions with major welfare implications over a household's lifetime, yet we still see strong limitations in how relevant information is incorporated into those decisions. Our experiment points to a way to improve the communication of macroeconomic fundamentals to households and how to build informational RCTs, but our calibrated model shows that much remains to be explained about the disconnect between household behavior in economic models and in the data.

The paper is organized as follows. The next section connects our findings to the related literature. Section 1 lays out the trade-offs of inflation-indexed debt, drawing on motivating evidence from the Brazilian mortgage market built from Central Bank of Brazil microdata. Section 2 develops our consumption-saving model, which we use both to discipline the survey objects and to run the quantitative exercises. Section 3 describes the survey and its experimental design, and Section 4 presents descriptive statistics. Section 5 shows how the treatments move beliefs about inflation and future payments, and Section 6 reports the effects of T_I and T_M on financial-planning outcomes. Section 7 presents results from the calibrated model, and Section 8 concludes.

Related Literature

Our main contribution relates to the literature studying the role of inflation beliefs in economic decisions. Influential papers have established a set of facts on how households form perceptions about inflation and how their actions connect to these beliefs (Coibion et al., 2022; Weber et al., 2022; D'Acunto et al., 2023b; D'Acunto and Weber, 2024), but the evidence on actions is mixed and sometimes misaligned with theoretical predictions (Bachmann et al., 2015; Duca-Radu et al., 2021; Burke and Ozdagli, 2023; Coibion et al., 2023). As formalized by Yang (2025), a plausible explanation is that these effects can be attenuated by uncertainty over the belief-action map: individuals do not know exactly how to map beliefs into actions. We bring evidence that, in the context of inflation, households can already fail at the early step of understanding how inflation creates exposure in their own budget constraints. Our setting also sidesteps a central measurement challenge in this literature, since exposure to inflation usually operates through multiple components of the household budget and is hard to pin down (Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2021; Binetti et al., 2024), whereas indexed debt makes one component of this exposure directly observable.

Our cross-randomized experimental design shows that the interaction between the inflation treatment and the mapping treatment induces economic behavior that inflation information cannot induce alone. This relates to the finding in Schnorpfeil et al. (2023), in which simply telling households about the erosion of their own nominal debt burden is effective in shifting consumption. Their treatment is close in spirit to our mapping treatment: they show how past inflation affected budgets, and that induced behavior. The broader idea, that making budget consequences explicit is what turns inflation information into action, should carry implications for the growing set of informational RCTs in macro and finance that provide inflation information and assess impacts on decisions (Coibion et al., 2022; D'Acunto et al., 2022; Coibion et al., 2023; Weber et al., 2025; Coibion et al., 2024; Coibion and Gorodnichenko, 2026).

This paper connects to several other strands of literature in economics and finance. Our

findings on the underestimation of future nominal payments speak directly to seminal papers on the welfare properties of inflation-indexed mortgage contracts (Friedman, 1974; Modigliani and Lessard, 1975; Campbell and Cocco, 2003). These papers argue that indexation can be welfare enhancing because it reduces the front-loading of real payments that nominal contracts impose on households. Our results suggest a possible downside: households may fail to project inflation into their future payments, and so underestimate the burden of an indexed contract in the future.

Our findings on the incomplete pass-through of inflation beliefs into expected nominal payments also speak to the literature on financial literacy (Lusardi and Mitchell, 2014), exponential-growth bias (Stango and Zinman, 2009; Almenberg and Gerdes, 2012) and on how cognitive ability shapes expectations and choices (D'Acunto et al., 2023a). Aside from literacy and cognitive factors, they also relate to behavioral economics work on how people approach questions that call for cognitive-costly calculations (Frederick, 2005; Kahneman, 2011). What is novel and perhaps surprising in our findings is that such a pattern happens among households that bear real indexation risk, with installments absorbing a meaningful share of their income, and that are mostly college educated.

The gap between respondents' preferences for fixed-payment contracts and the predictions of our consumption-savings model connects to a literature in household finance documenting household inaction, even when the gains from acting are large. In our model, switching to a fixed-payment contract is largely beneficial, but most respondents still decline the offer. Similar inaction has been documented for refinancing in mortgage markets (Keys et al., 2016; Andersen et al., 2020; Berger et al., 2021; Majer, 2026; Berger et al., 2025). Gaps of this kind also motivate a theoretical literature that formalizes the map from beliefs to actions, where agents solve their decision problems imprecisely rather than optimally (Woodford, 2020; Ilut and Valchev, 2023; Enke and Graeber, 2023). This discussion regarding preferences over fixed-payment versus inflation-indexed contracts also connects to papers discussing how borrowers choose between adjustable-rate and fixed-rate mortgages (Badarinza et al., 2018; Koijen et al., 2009) and payment structure (Clara et al., 2024).

At last, our work connects with other papers studying consórcio contracts. A recent literature spanning development, labor, and finance has exploited the random component of consórcio allocation to study the effects of credit access on employment and earnings (Van Doornik et al., 2024b) and on fertility and housing (Fazio et al., 2024). In related work, Braga et al. (2026) study how participants' beliefs about the allocation mechanism shape both the demand for and the supply of these contracts. Our paper sheds light on an underexplored feature of the consórcio contracts, which is the inflation-indexation of monthly installments.

1 Motivating Evidence

Indexation is the laboratory of this paper. It ties households' nominal expenses to inflation by contract, and this is what allows us to compare what borrowers believe about inflation with what they believe about their own liabilities and monthly budgets in the future. Before turning to the survey, this section shows why and when indexation matters to a household. We do so by analyzing a recent Brazilian episode that allows indexed and fixed-rate contracts to be directly compared: we look at differences in initial installments, their evolution during a period of ordinary inflation, and during an inflationary episode.

Relative to a contract with fixed rates, indexed debt imposes a straightforward tradeoff. A fixed-rate contract has its monthly installments constant over time[3], and their level is determined by the loan rate established at origination. These rates incorporate market expectations of inflation throughout the duration of the contract, a step that is not necessary in indexed contracts. By stripping the inflation premium out of loan rates, the indexed contract begins with a lower installment; in exchange, whenever the index rises by $x\%$, so does the outstanding debt balance. The design has long been discussed as an alternative design for household debt: beginning with Modigliani and Lessard (1975) and Friedman (1974), a large literature has argued that indexation improves risk sharing and removes the payment "tilt" that inflation imposes on conventional nominal loans (Baesel and Biger, 1980; Campbell and Cocco, 2003; Green and Wachter, 2005), while the real-world record has been considerably more mixed (Cukierman, 1975; Buckley et al., 1993). Whether indexation ends up being an advantage for households depends on how long the indexed installment stays below the fixed payment alternative and how far it then rises above it.

Following a regulatory change in August 2019, Brazil's largest public bank and then other major lenders began originating long-term household loans corrected by the IPCA[4], the country's headline consumer price index (Vivan, 2025), alongside the conventional contracts they already offered, whose nominal installments are essentially fixed.[5] Three facts make this setting interesting. First, the two contracts were offered side by side, by the same lenders and in the same months, so the gap in their initial terms can be observed. Second, the Central Bank's credit registry records the origination terms of virtually every contract. Third, borrowers signed when inflation was low and stable, with the IPCA closing 2019 at 4.31%, and then lived through the post-pandemic surge, which carried inflation into double digits by 2022.

Among the roughly 120 thousand contracts originated between September 2019 and February 2020 (16% of them indexed), our estimates indicate that indexed contracts carried

[3] Assuming a constant payment amortization schedule.

[4] The acronym stands for "*Índice de Preços ao Consumidor Amplo*." For a thorough description of the regulatory change, see Conselho Monetário Nacional (2019).

[5] Formally, their balances are corrected by a reference rate, not a price index, that has remained near zero since 2017, which is why they behave as fixed-payment loans in practice (Appendix G.3).

rates 4.2 percentage points below comparable fixed-rate contracts from the same lender and month (Appendix G.3). To express this rate gap as an installment gap, we take the contract most familiar to the household finance literature: a thirty-year loan repaid in constant installments, the standard in the United States mortgage markets. At the average conventional rate of 8.37%, a rate 4.2 points lower means the indexed borrower starts out paying 65 cents for every dollar paid under the fixed contract (Appendix G).

Figure 1 follows this installment over time. When these contracts were signed inflation was at ordinary levels, and had it stayed there the payment gap between an indexed and a fixed contract would have survived for a longer period. Instead, as Panel (b) in Figure 1 shows, twelve-month inflation climbed from under 2% in 2020 to double digits in 2022 after the Covid-19 pandemic, and the indexed installment reached 94 cents on the dollar already in 2026: six and a half years into a thirty-year contract, most of the discount is gone.[6] A household deciding on such a contract therefore needed to anticipate inflation along two dimensions: its expected level, which determines how long the entry discount lasts, and its upside risk, which determines how costly the contract can become—and how costly it in fact became over these six years.[7]

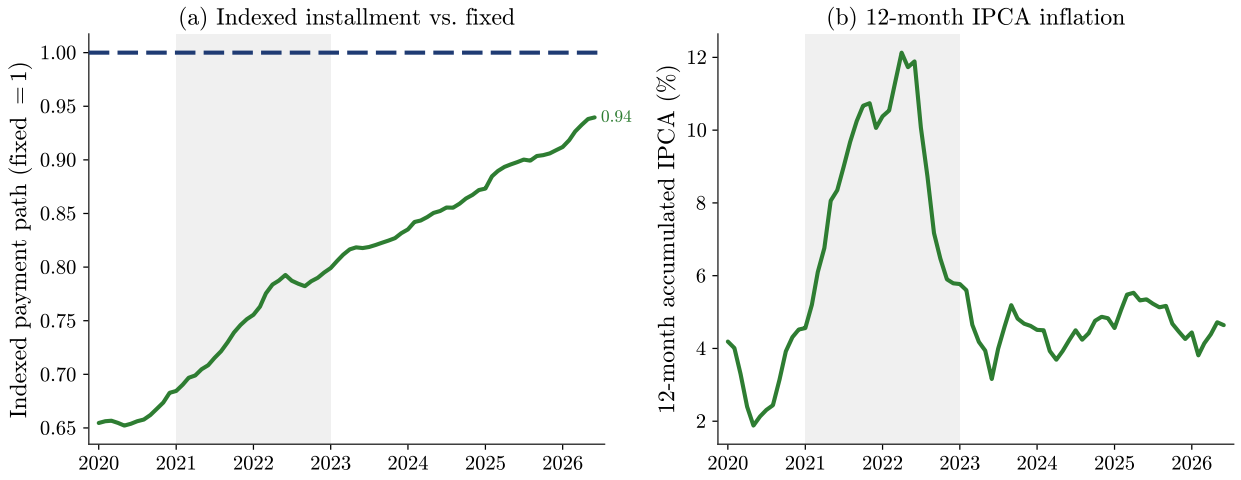


Figure 1. Indexed installment per dollar of the fixed installment. Thirty-year contracts signed in January 2020 and repaid in constant installments; the indexed contract carries a rate 4.2 percentage points below the fixed contract's 8.37%. Panel (a) follows the indexed installment under realized IPCA inflation, and panel (b) plots the twelve-month accumulated IPCA over the same period.

From the consumer perspective, making that bet on indexation requires inflation beliefs, and it also requires turning those beliefs into the installments the household will actually pay.

[6] Paths assume no prepayment or restructuring; see Appendix G. We also abstract from discussing debt balance implications, since such number is also adjusted by inflation and should affect prepayment decisions.

[7] We abstract from downside risk and deflationary scenarios because of how rare they are, especially for the Brazilian economy.

These are the two objects this paper measures, and the targets of its two separate treatments. Our survey places respondents in front of exactly this trade. As detailed in Section 3.3, we elicit their low, average, and high inflation scenarios and the nominal installment they expect to pay under their own average scenario, and then we assess relations with financial decisions.

2 Theoretical Framework

This section lays out a framework that disciplines our interpretation of the household problem when facing an inflation-indexed liability path. The model serves two purposes: (i) guiding the interpretation of the empirical objects elicited in the survey, and (ii) providing a normative benchmark for households' evaluation of alternative debt contracts. For the former, we present throughout the paper the model counterpart of each object we measure. For the latter, we compare the benchmark predicted by the model with survey responses on the willingness to switch contracts in Section 7.

We build a consumption-savings model in which households hold an inflation-indexed installment loan, so the nominal burden of their debt moves directly with realized inflation. Income is also exposed to inflation, since nominal wages adjust only periodically. Planning therefore requires the household to assess exposure on both sides of the budget. We then model preferences over an alternative contract without indexation.

2.1 Environment

Preferences. Time is discrete and the household is infinitely lived, with CRRA preferences over real consumption,

$$\tilde{\mathbb{E}}_t \sum_{s\geq 0} \beta^s \, u(c_{t+s}), \qquad u(c) = \frac{c^{1-\chi}}{1-\chi},$$

where χ is the coefficient of relative risk aversion. The operator $\tilde{\mathbb{E}}_t$ indicates expectations formed under the subjective beliefs of this household with respect to the exogenous processes that affect their decision problem, which we formalize below.

Budget Constraint. The household is repaying an installment loan that runs for T periods, owing the nominal installment PMT_t in each month $t \leq T$ and nothing thereafter. Each period it chooses real consumption c_t and end-of-period nominal assets A_{t+1}, held in a one-period bond paying the gross nominal return

$$1 + R_t = \exp(\pi_t + r).$$

This equation illustrates a one-for-one inflation pass-through to the nominal return and fixes the real return at r, so that saving is the household's only instrument for buffering payment

shocks. With P_t the price level, the budget constraint is

$$P_t c_t + A_{t+1} + \mathrm{PMT}_t \, \mathbf{1}\{t \le T\} \;=\; Y_t + A_t(1 + R_t), \qquad A_{t+1} \ge \underline{A}, \tag{1}$$

where Y_t is the stochastic nominal income, $\underline{A}$ is a borrowing limit, and the indicator $\mathbf{1}\{t \le T\}$ drops the installment once the loan is repaid. We assume that strategic defaults do not occur, something we further discuss below when defining value functions for this problem.

Inflation Process. Inflation is the object households hold beliefs about. It follows a first-order Markov process with perceived conditional distribution

$$\pi_{t+1} \sim \tilde{\Gamma}(\cdot \mid \pi_t), \tag{2}$$

the subjective law $\tilde{\Gamma}$ on which the household acts, which need not coincide with the physical process $\Gamma(\cdot \mid \pi_t)$ that actually generates inflation and, with it, the price level, $P_{t+1} = P_t \exp(\pi_{t+1})$. The household's decisions respond to $\tilde{\Gamma}$, not to the physical law. Our survey in Section 3 describes how we elicit $\tilde{\Gamma}$. We keep both laws generic here, but for the quantitative work in Section 7 we adopt a Gaussian AR(1) process.

Income Process. Nominal income adjusts only intermittently. Each month it resets with probability $1 - \psi$ and is left unchanged with probability ψ, so $\psi \in [0, 1]$ measures stickiness. At a reset, income catches up to a fraction $\rho_y \in [0, 1]$ of the inflation accumulated since the previous readjustment; between resets, it moves only with a real drift μ_y and an idiosyncratic shock $\xi_{t+1} \sim \mathcal{N}(0, \sigma_y^2)$, so real income erodes as prices rise. Tracking the catch-up requires one extra state, the cumulative inflation since the last reset S_t, and the pair (Y_t, S_t) evolves as

$$(Y_{t+1},\, S_{t+1}) = \begin{cases} \big(Y_t \exp(\mu_y + \rho_y(S_t + \pi_{t+1}) + \xi_{t+1}), \;\; 0 \big) & \text{with prob. } 1 - \psi \quad \text{(reset)}, \\ \big(Y_t \exp(\mu_y + \xi_{t+1}), \;\; S_t + \pi_{t+1} \big) & \text{with prob. } \psi \quad \text{(no reset)}. \end{cases} \tag{3}$$

Neither ψ nor ρ_y is directly observable, and estimating them for an individual household would require a long panel of nominal income. Our survey elicits a household's perception of both. Unlike inflation beliefs and the payment rule, these perceptions have no observable counterpart against which they could be scored as right or wrong. We do not see that as a limitation, since a household deciding whether to keep an indexed contract or switch to a fixed one weighs the rising payment path against how exposed it believes its income to be.

Payment Path. The contract is price-level adjusted: the installment tracks the price level one-for-one,

$$\mathrm{PMT}_{t+1} = \exp(\pi_{t+1}) \, \mathrm{PMT}_t. \tag{4}$$

Value Function. Let Ω_t denote the state vector. It collects wealth (A_t), income (Y_t), the cumulative inflation since last reset (S_t), the price level (P_t), current inflation (π_t), and

the installment (PMT_t),

$$\Omega_t = \big(A_t,\, Y_t,\, S_t,\, P_t,\, \pi_t,\, \text{PMT}_t\big). \tag{5}$$

Given the household's perceptions, its value function solves

$$\tilde{V}_t(\Omega_t) = \max_{c_t,\, A_{t+1} \geq \underline{A}} u(c_t) + \beta\, \tilde{\mathbb{E}}_t\Big[\tilde{V}_{t+1}(\Omega_{t+1})\Big] \quad \text{s.t.} \quad (1),\ (2),\ (3),\ (4). \tag{6}$$

The tilde on $\tilde{V}_t$ records that the value function is computed under the household's subjective inflation law $\tilde{\Gamma}$. Note that the household does not die when the loan ends at T: from $T+1$ onward it simply solves the same consumption–saving problem, now debt-free. The value of that stationary, infinite-horizon problem supplies the terminal condition $\tilde{V}_{T+1} = \tilde{V}^{\infty}$ for every contract below.[8] We also rule out strategic default: the household meets its installment whenever it can, defaulting only when its resources cannot cover the payment.[9]

2.2 Preferences over Alternative Contracts

We consider an alternative contract that has payments fixed in nominal terms, it freezes the nominal installment for the remainder of the term at a markup $\gamma \geq 0$ over the current one. Let $\tilde{V}_t^{\text{fix}}(\Omega_t; \gamma)$ denote the value of a household in state Ω_t that switches to this contract: it solves (6) with the payment rule (4) replaced by the constant path

$$\text{PMT}_{t+k} = (1+\gamma)\,\text{PMT}_t \qquad \text{for all } k \geq 0. \tag{7}$$

The household prefers to switch at markup γ when

$$D_t(\Omega_t; \gamma) \;\equiv\; \mathbb{1}\big\{\tilde{V}_t^{\text{fix}}(\Omega_t; \gamma) \geq \tilde{V}_t(\Omega_t)\big\} \;=\; 1. \tag{8}$$

$D_t(\Omega_t; \gamma)$ is what the survey measures. Respondents face hypothetical offers at 10%, 20% and 30% markups, so we have survey measurements of $D_t(\Omega_t; \gamma)$, with $\gamma \in \{0.1, 0.2, 0.3\}$. Switching trades the lower starting level of indexation for the removal of future payment uncertainty, so D_t is governed by risk aversion χ, by liquidity, and by the household's beliefs about inflation and about its own income exposure.

How much that matters depends on the income process. For a household whose nominal income does not adjust, the indexed path in Figure 1 also describes how the installment grows relative to income: an installment absorbing 14% of income today (the sample average in

[8] The continuation value solves $\tilde{V}^{\infty}(A, Y, S, P, \pi) = \max_{c,\, A' \geq \underline{A}} u(c) + \beta\, \tilde{\mathbb{E}}[\tilde{V}^{\infty}(A', Y', S', P', \pi')]$ subject to $Pc + A' = Y + A(1+R)$, the budget constraint (1) stripped of the installment.

[9] Numerically, default is triggered only when available resources are insufficient to pay the installment given the borrowing limit $\underline{A}$; the household then moves to the terminal debt-free phase $\tilde{V}^{\infty}$ but bears a large fixed utility cost. The penalty keeps default a last resort rather than a strategic option in our model, so contract rankings reflect payment burdens rather than the value of an option to walk away.

our survey, see Section 4.1) would absorb about 20% after ten years of 4% inflation. For a household whose income fully tracks prices, it is the indexed installment that stays constant relative to income, while the fixed one is deflated. This is something we further explore in our numerical results in Section 7.

2.3 Saving and Budgeting

Let $A^*_{t+1}(\Omega_t)$ be the optimal asset choice in problem (6). The saving policy function is the household's end-of-period assets relative to its current income,

$$\mathcal{S}_t(\Omega_t) \equiv \frac{A^*_{t+1}(\Omega_t)}{Y_t}, \tag{9}$$

optimal under the subjective inflation law $\tilde{\Gamma}$. The survey does not measure $\mathcal{S}_t$ exactly: it asks respondents what share of their monthly income they plan to save over the next twelve months to prepare for financial emergencies. We treat this answer as an attempt to measure the same object. Naturally, such object should be directly determined by inflation beliefs due to both income and liability exposures.

3 Survey: Context and Design

This section describes the construction of each survey block, with particular focus on the implementation of the two cross-randomized treatments. We also provide background on sample recruitment and the financial product held by respondents, the consórcios. Section 3.1 describes consórcio contracts, Section 3.2 brings details on sample recruitment, and Section 3.3 describes the design of the survey.

3.1 Background on Consórcios

Consórcios are Brazil's largest form of group-based finance and a major source of household credit for durable assets. Participants join a group run by a licensed administrator, contribute monthly installments to a common pool over a fixed term, and each month one or more members receive credit to purchase a durable good, allocated either by lottery or by an auction in which members bid to receive credit earlier (Van Doornik et al., 2024a; Braga et al., 2026). The financial firms that run consórcio groups are responsible for conducting these auctions and lotteries, with their operations regulated and supervised by the Central Bank. Clients only deal with the financial firm and are not required to interact with other participants. Real estate consórcios alone had 2.83 million active contracts in December 2025

and had R$30.24 billion in credit distributed during 2025,[10] which corresponds to roughly 9.58% of all real estate credit volume originated through conventional mortgage markets in Brazil in the same year.[11]

From the standpoint of the household, a real estate consórcio is similar to a mortgage with an uncertain credit disbursement date. Installment values and maturities are determined at origination; they cover the credit amount plus administrative fees, and are revised according to the contractual price index at fixed readjustment dates, so that a rise of $x\%$ in the index raises the outstanding balance and the monthly payment by $x\%$. For real estate consórcios this index is the INCC (*Índice Nacional de Custo da Construção*), a construction cost index computed by Fundação Getulio Vargas. Although it tracks input and labor costs in construction rather than a consumption basket, over long horizons it behaves much like Brazil's headline consumer price index, which is the index the Central Bank targets and the one households are most likely to have in mind when they think about inflation.[12] Participants begin paying installments from the first period of the contract. Once credit is received, a participant can use the resources to purchase (or remodel) a home and continues to pay the monthly installments until the end of the term, with the purchased property serving as collateral—as under a standard mortgage. Households are guaranteed to receive credit by the end of the contract. Thus, apart from the timing of credit disbursement, the contractual path of the liability is essentially a long-term stream of indexed monthly payments.

The relevant difference between the liability path in a consórcio and in an indexed mortgage is the possibility of restructuring the payment path when credit is allocated. In particular, participants who receive credit earlier after participating in their group auctions can deduct their winning bids from future payments. Depending on how the deduction is applied, this can reduce the value of subsequent installments or shorten the remaining term. Such restructuring is possible but not the norm, and the standard payment path is to continue the original schedule of installments. We account for the possibility of restructuring in the elicitation of future payments, as explained in Section 3.3. Moreover, regardless of any restructuring, indexation remains a permanent feature of the contract.

3.2 Logistics and Sample

The survey was fielded in August 2026 through a partnership with a large consórcio administrator in Brazil. Invitations containing a Qualtrics link were sent by the firm itself—

[10] Data on active contracts and distributed credit are from the *Anuário Estatístico do Sistema de Consórcios 2026*, published by the Associação Brasileira de Administradoras de Consórcios (ABAC, the Brazilian Association of Consórcio Administrators).

[11] Calculated using consórcio credit figures from ABAC (2026) and conventional mortgage market volume statistics from the Banco Central do Brasil (BCB), available at https://www.bcb.gov.br/estatisticas/mercadoimobiliario.

[12] Figure 17 plots the cumulative behavior of both indices.

through its official WhatsApp channel—to active clients holding real estate contracts.[13] The questionnaire contained approximately 30 questions, and the median respondent completed it in 13.8 minutes.[14]

Recruitment through the administrator's official channel is a central feature of our design. Respondents are individual holders of actual inflation-indexed credit contracts, with substantial shares of their monthly budgets at stake. The beliefs we elicit, therefore, concern a liability these households are currently paying, and our questions about future nominal payments refer to their own contracts rather than to a stylized scenario.

Around 15 thousand clients were contacted by the consórcio administrator; 1,463 consented and 961 completed the survey.[15] We then apply three quality filters: an attention check, an income–installment consistency check applied to respondents flagged as speeders, and a trim that drops the fastest 10% of respondents on the outcome-elicitation pages. This yields a final sample of 783 respondents.

3.3 Survey Design

Pre-treatment Blocks. The survey followed the sequence described in Figure 2. A consórcio client that clicked the Qualtrics link received through WhatsApp was directed to a welcome page, followed by a consent form. After providing consent, individuals faced the first block of questions regarding demographics (age, sex, race, education) and financial profile (income brackets, employment type, exposure of own income to inflation). They were then asked about their credit contract (maturity at origination, installment value, payment-to-income ratio, number of remaining installments, credit status) and their awareness of indexation and current inflation. These variables give us a direct connection between our data and the state vector Ω_t defined in the model, and the income exposure parameters ψ and ρ_y (see Section 2.1). After this pre-treatment block, all individuals were reminded and exposed to a careful explanation on the implications of indexation. The survey then moved to the inflation beliefs block.

Inflation Treatment and Beliefs. The inflation elicitation and treatment assignment closely follow the approach in Weber et al. (2025); Coibion and Gorodnichenko (2026). The logic is simple: we measure priors, randomize treatment assignment, then measure posteriors. There are, however, a few distinctions in our setup—which we explain below.

We first elicit prior beliefs about the INCC, the index used in real estate consórcio contracts. Respondents report the annualized inflation they expect over the next five years

[13] Survey respondents were eligible for a raffle offering one R$200 gift card for every 40 participants.

[14] The experiment was pre-registered at the AEA RCT Registry (AEARCTR-0019465), where the treatment arms, primary outcomes and sample restrictions are all pre-specified.

[15] Most of the attrition comes from clients that do not even visualize these WhatsApp messages or do not open the Qualtrics link. It is impossible to assess any kind of selection on observables, as we have no data on these individuals.

under three scenarios: low, average, and high. The average scenario has a straightforward interpretation of an expected value. For purposes of the analysis, we interpret these scenarios as the 5^{th} percentile, the expected value, and the 95^{th} percentile of their perceived distribution of annualized variation of the price index between dates t and $t+5$. Let P_t denote the price level measured by the index at period t. Using a yearly frequency for t, the random variable for which we ask scenarios is denoted by $\Pi_t^{(t+5)}$, and its formal definition is given by:

$$\Pi_t^{(t+5)} \equiv \left(\frac{P_{t+5}}{P_t}\right)^{\frac{1}{5}} - 1 \ . \tag{10}$$

Using such definition, we can mathematically define the beliefs we elicited during the survey for a given participant i: Π_i^L as the perceived 5^{th} percentile of $\Pi_t^{(t+5)}$ (low scenario), Π_i^e as its expected value (average scenario), and Π_i^H as the 95^{th} percentile (high scenario). We formalize these definitions in Appendix A. When eliciting the Π_i^L and Π_i^H objects, however, we presented them using a much more heuristic definition. As pointed out in several papers that discuss belief elicitation, individuals tend to struggle with the understanding of basic statistical objects, which led us to believe that respondents would have a poor understanding of the concept of quantiles (Manski, 2004; Weber et al., 2022). Hence, when asking about the low and high scenarios, we described the object as a "a realistic high-inflation scenario where the index grows more than you expected". We also explicitly mentioned that we were not asking about "extreme" scenarios. Of course, individual respondents can differ substantially in their interpretations of what these scenarios mean, so we can also interpret the triple $(\Pi_i^L, \Pi_i^e, \Pi_i^H)$ as capturing respondents' perceptions of plausible inflation scenarios. These objects also allow us to build a bridge between our data and the model object $\tilde{\Gamma}$ referring to perceived law-of-motion for the inflation process (see Section 2.1).

The reason why we stick to the quantile interpretation is that it naturally links these beliefs to the randomized inflation treatment. After eliciting priors, while the control group was simply exposed to a brief reminder on the indexation rules, the treated group was exposed to a screen in which three different numbers were shown: the 5^{th} percentile, the mean, and the 95^{th} percentile of the historical data distribution of annualized inflation (referring to the time window between January 2008 and June 2026). A screenshot of the treatment page is available in Appendix B. We interpret these statistics as signals about the corresponding scenarios being elicited $(\Pi_i^L, \Pi_i^e, \Pi_i^H)$, with the historical mean serving as a central signal for the average scenario. A natural question that might arise is: why show historical data if we are asking about 5-year ahead forecasts? The reason why we chose this is because forecasts for long horizons can be highly uncertain, and few professional forecasters even attempt to produce precise high- and low-inflation scenarios at such horizons. Moreover, as the predictive content of current conditions declines with the forecast horizon, historical data can still be informative about long-term realizations.

Mapping Treatment and Nominal Payments. Next, all respondents report the monthly installment, in Brazilian reais, that they expect to pay five years ahead if their own average inflation scenario materializes.[16] This object is denoted by PMT_i^e and can be interpreted as individual i's expectation over 5-year ahead installments conditional on their own inflation belief Π_i^e

$$\text{PMT}_i^e \equiv \mathbb{E}_{it}\left[\text{PMT}_{t+5} \mid \Pi_t^{(t+5)} = \Pi_i^e\right] . \tag{11}$$

Because inflation is fixed at the respondent's own number, the correct answer involves no forecasting: it is the current installment compounded at that rate for five years. Whatever distance separates PMT_i^e from this benchmark therefore reflects a failure to map the respondent's own inflation belief into their own payment, rather than bad forecasts about inflation itself. In Section 5.2, we formalize a pass-through statistic that measures how much of their own belief respondents carry into their projected installment.[17]

The mapping treatment is randomized on this page eliciting nominal payments, immediately before the question and independently of the inflation treatment. The treatment performs this compounding for the respondent: we implemented an algorithm through Qualtrics that takes the current installment PMT_{it} they reported in the pre-treatment block, compounds it at their stated annual rate Π_i^e, and displays the implied increase in reais,

$$\Delta\text{PMT}_i \equiv \text{PMT}_{it}\left[(1+\Pi_i^e)^5 - 1\right] . \tag{12}$$

Control group respondents answer without this aid. The treatment carries no new information about inflation. It only converts a compounding problem into an addition: treated respondents need only add the number exposed to them (ΔPMT_i) to their self-reported installment (PMT_{it}). A screenshot of the page is in Appendix B.

Outcomes on Financial Planning. The last block measures two ways a household can prepare for rising installments. The first is saving: respondents report how much they plan to save as a percentage of their monthly income (over the next 12 months) to prepare for financial emergencies. The second is demand for a fixed payment schedule. After a screen explaining the tradeoff between an indexed path of payments and a initially higher but certain path, respondents face a hypothetical offer to switch their indexed contracts to a fixed installment schedule. The first offer involves a 20% higher (but fixed) path. Those who accept are then offered 30%, and those who reject are offered 10%. Answers are binary—"yes" or "no". Note that these choices are precisely the ones formalized in our theoretical benchmark (see Section 2.2 and Section 2.3). As illustrated in our motivating evidence using

[16] We ask them to assume the payment path goes through no restructuring, only inflation adjustments.

[17] By Jensen's inequality, the installment implied by the average inflation scenario is not exactly the expected installment five years ahead. At our horizon and inflation levels the two should be numerically close, and we read the question as eliciting the former.

the Brazilian Central Bank data (see Section 1), these choices regarding financial planning outcomes should be highly influenced by households' perceptions on the inflation process, and that is why we focus our analysis there.

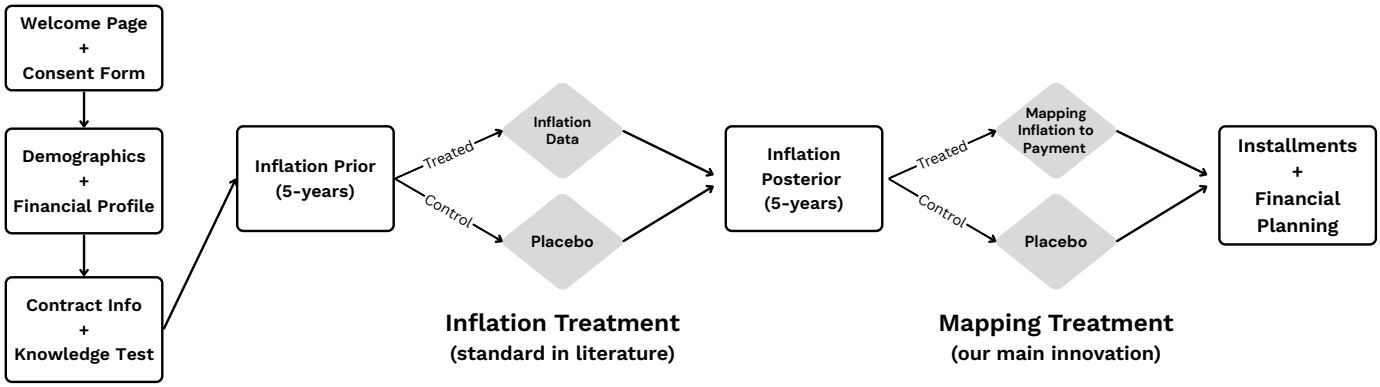


Figure 2. Survey layout and treatment allocation. Respondents move through the blocks from left to right. The two treatments are cross-randomized: the inflation treatment is assigned between the prior and posterior belief elicitations, and the mapping treatment immediately before the nominal installment question.

4 Descriptive Statistics

4.1 Demographics and Financial Profile

We start with the description of age, gender, race and education profiles in our sample. Respondents are predominantly male, white, middle-aged and college educated. Part of the results shown in Table 1 reflects the regional concentration of the survey respondents, who come mostly from Southern states with a predominantly white population: 50.7% of respondents were located in Paraná, 20.3% in São Paulo, 6.8% in Santa Catarina and 4.9% in Rio Grande do Sul.[18]

According to the 2022 Demographic Census (Instituto Brasileiro de Geografia e Estatística, 2023), the white population share in these states ranges from 57.8% in São Paulo to 78.4% in Rio Grande do Sul, against 67.9% of our respondents. Education differs sharply from the general population: 65.6% of participants hold a college degree, roughly half of them with a graduate degree, while only 17.4% to 25.0% of adults in these states hold one. Respondents also report higher income than the average household where they live. Figure 3 shows that 41.6% report household income above 10 thousand Brazilian reais, while the household income implied by official household surveys averages roughly 7 to 8 thousand

[18]The remaining 17.3% are spread thinly across the remaining states.

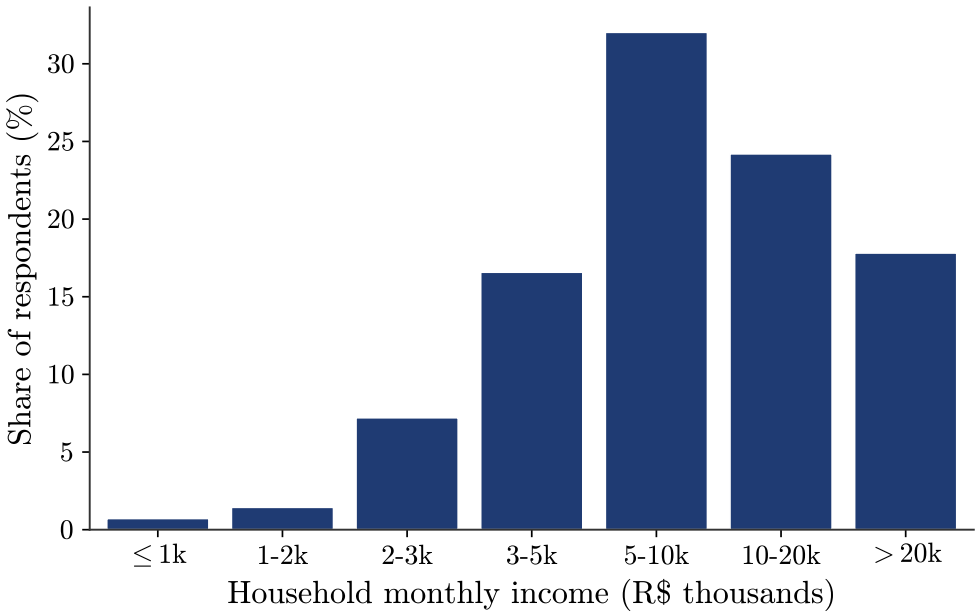


Figure 3. Household monthly income. Distribution of self-reported household monthly income across brackets, in thousands of reais. As a reference, the exchange rate was R$1,000 ≈ US$195 at the time of the survey (August 2026). Monthly minimum wages in Brazil at the time of the survey was R$1,621.

reais per month in these states.[19] Since this is an average rather than a median, and the Brazilian income distribution is heavily right-skewed, the comparison understates how much richer our respondents are than the typical household in these states.

The age of survey participants can be compared against the population that actually holds housing debt than against the population at large. In the 2022 Census, 9.1% of Brazilians lived in a home still being paid for, a share that peaks at 13.5% between ages 35 and 39 and declines steadily thereafter (Instituto Brasileiro de Geografia e Estatística, 2024). Our respondents average 42 years of age, placing them in the age range where households are most likely to be servicing a housing liability.

Table 1 also shows that respondents' consórcio contracts have long maturities, 17.6 years on average, and absorb a meaningful share of the budget, with a mean payment-to-income ratio of 14%. Before observing these descriptive patterns in our data, a natural concern is that respondents might be too poorly educated, too financially constrained, or too disengaged from their own contracts to understand their exposure. If anything, selection on observables runs the other way: these are households with unusual education and income, holding a real and long-dated indexed liability, so we have no evidence that the misperceptions documented in Section 5.2 are an artifact of an unsophisticated sample.

[19] In the household survey "PNAD Contínua", IBGE reports household income per resident rather than per household; we obtain the figures above by multiplying it by the average number of residents per household in each state.

	Mean	Std. dev.	P25	P50	P75
Demographics					
Age (years)	42.0	10.9	34.0	42.0	50.0
Male	0.644		–		
Race					
White	0.679		–		
Black	0.054		–		
Brown (*pardo*)	0.244		–		
Others	0.023		–		
Education					
Less than high school	0.065		–		
High school complete	0.278		–		
College complete	0.318		–		
Graduate	0.338		–		
Contract					
Current installment (R$)	1320	2067	349	792	1498
Payment-to-income ratio (%)	14.0	12.1	6.0	10.0	20.0
Maturity at origination (months)	211	39	204	222	240
Remaining months	182	51	161	200	218
Received credit	0.116		–		
Household finances & employment					
Prepared for financial emergencies?					
Not at all prepared	0.079		–		
Slightly prepared	0.243		–		
Moderately prepared	0.430		–		
Very prepared	0.201		–		
Completely prepared	0.047		–		
Employment type					
Entrepreneur / self-employed	0.446		–		
Wage-earner, formal (registered)	0.356		–		
Public-sector worker	0.130		–		
Wage-earner, informal (unregistered)	0.043		–		
Unemployed	0.024		–		
Observations		783			

Table 1. Sample descriptive statistics. Continuous variables report the mean, standard deviation and the 25th, 50th and 75th percentiles; binary indicators report the share only. Race follows the Brazilian census classification, with *pardo* translated as "brown" and "Others" pooling the Asian, Indigenous and residual self-classifications. Payment-to-income is the self-reported share of monthly household income committed to the installment.

4.2 Perceptions of Income Erosion

Within the financial profile block of the survey, we included a few questions designed to capture individual perceptions on income erosion. Since our paper assesses the role of inflation beliefs in households' decision making, a crucial step is to understand how income can be affected by price-level indices—or at least how households perceive that phenomenon. Based on the theoretical income process described in Section 2.1, we designed one question to measure the income pass-through parameter (ρ_y) and another one to separately measure the adjustment frequency (ψ). On top of that, we also asked a qualitative question regarding overall perception of income erosion. Responses for all of the questions show widespread perceptions of income erosion by inflation.

For the qualitative question in Table 2, only 5.7% of respondents reported to have income rising by about the same as inflation. In the quantitative question about pass-through, the share of responses implying a full pass-through is higher, standing at 15.2%. With respect to frequency of adjustment after an inflationary shock of 10%, only 8.9% of respondents report to have income catching up within 6 months. It is also important to highlight that a significant share of respondents report some degree of uncertainty about how inflation affects their own income, which is shown by the last alternative listed for each question in Table 2.

These patterns line up closely with what has been documented elsewhere. In representative US surveys, Stantcheva (2024) finds that the predominant reason households dislike inflation is the belief that it erodes their purchasing power, with around 80% of respondents reporting that prices systematically rise faster than wages—a view Binetti et al. (2024) also document. The finding is durable and not specific to the recent inflationary episode: Shiller (1997), surveying respondents in the United States, Germany and Brazil, reports that non-economists in all three countries hold a sticky-wage view in which wages do not respond to inflationary shocks.

To incorporate these perceptions into our empirical analyses, we combine the pass-through and adjustment-frequency answers into a single continuous income indexation score per respondent, $\zeta_i \in [0, 1]$, which measures how much of inflation a respondent expects to reach their income within a year. Appendix D describes the construction and shows that the score varies across employment types in the expected direction and also with income.

4.3 Indexation Awareness and Attention to Inflation

Were individuals at all aware of their own contract indexation? If so, then respondents could possibly fail to map inflation to installments just because they did not know about contract rules. Therefore, before implementing any of the treatments or gauging inflation forecasts, we asked individuals if they were aware of the indexation feature. 76.4% of respondents report to be aware. Along with the fact that all individuals are explicitly reminded of indexation features right after answering this question, this number suggests that eventual

Qualitative: "Does your income keep up with inflation?"	
Yes, rises by about the same amount	5.7%
Yes, but rises less than prices	40.6%
No, stays the same regardless	39.6%
It varies	14.0%
Pass-through: "If inflation = 10%, how much passes to income?"	
5 p.p. or less	37.3%
More than 5 p.p., less than 10 p.p.	26.3%
10 p.p. or more	15.2%
I do not know	21.2%
Adjustment Frequency: "How long for income to catch up?"	
Less than 6 months	8.9%
6 to 12 months	26.9%
12 to 24 months	28.2%
More than 24 months	19.8%
I do not know	16.1%

Table 2. Perceived income exposure to inflation. Shares of the main sample ($N = 783$) choosing each response. The three questions map to the income process of Section 2.1: the second elicits the pass-through parameter ρ_y and the third the adjustment frequency ψ. Pass-through answers are stated in percentage points of income for a hypothetical 10% inflation, to distinguish them from the inflation rate itself. No respondent selected "I do not know" on the qualitative question, so that option is omitted there.

mistakes in understanding inflation exposure related to the consórcio contract cannot be attributed to ignorance about contract rules.

Despite the fact that the majority of respondents were aware of the indexation of their contracts, we found significant heterogeneity in respondents estimates of past realized inflation. The survey asked people what they thought was the realized inflation between August 2025 and July 2026, the 12 months before the survey implementation. This was public information. Among those who declared to be aware of indexation, 32.9% of respondents gave answers within one percentage point of the true value (12 month inflation was sitting at 6.76%), and 55.2% were within two percentage points. For the group that was unaware of indexation, those shares dropped to 20.0% and 41.1%. Dispersion of answers was also smaller for aware individuals, with a standard deviation of 3.44 percentage points compared to 4.64 for the unaware. Even so, the differences are modest in absolute terms: around half of all respondents miss the realized rate by more than two percentage points, whether or not they knew their installments were indexed.

These differences, however, do not carry over to the object we care about. Appendix

E shows that the awareness indicator is uninformative for our individual perceived pass-through measure, and that knowledge of past inflation does not predict the level of that measure either. It does weakly predict how far respondents land from full pass-through: those who recall higher past inflation, and those whose recall is closer to the true value, tend to have nominal projections more consistent with their expected inflation. The magnitudes are small, though, and they leave the bulk of the dispersion in our pass-through measure unexplained. Knowing the contract rule and following inflation are evidently not enough to translate one's own beliefs into one's own future payments.

5 Treatment Effects on Beliefs

5.1 Inflation Treatment and Belief Revisions

The inflation treatment is designed to exogenously shift beliefs about the inflation distribution. As discussed in Section 3 and shown in Appendix B, individuals that receive the inflation treatment are exposed to signals about each inflation scenario: low, average and high. We refer to them as s^L_Π, s^e_Π and s^H_Π, respectively. In this section we will formally test shifts in beliefs and show that our exposure to historical moments of the distribution significantly influences households' long-term forecasts.

We now distinguish pre- and post-treatment inflation beliefs. We let $(\Pi^L_{i,\text{pre}}, \Pi^e_{i,\text{pre}}, \Pi^H_{i,\text{pre}})$ denote priors and $(\Pi^L_{i,\text{post}}, \Pi^e_{i,\text{post}}, \Pi^H_{i,\text{post}})$ be the posteriors. We also define a measure of upside risk of inflation, given by $\sigma^H_i \equiv \Pi^H_i - \Pi^e_i$. This is the distance in percentage points between the perceived high-inflation scenario and the expected scenario, which can also be measured before and after our interventions. Analogously, we define our measure of downside risk $\sigma^L_i \equiv \Pi^e_i - \Pi^L_i$. The signals for these two measures follow from the signals for the scenarios, $s^H_\sigma \equiv s^H_\Pi - s^e_\Pi$ and $s^L_\sigma \equiv s^e_\Pi - s^L_\Pi$.

The cross-section of inflation beliefs. In our setup, the measurement of inflation forecasts differs from related papers in two features: the three-scenario elicitation and the longer horizon (5-year ahead) of the forecasts. Our findings, however, are largely aligned with the stylized facts of the related literature, namely the wide disagreement across households and expectations that sit above realized inflation (Weber et al., 2022; D'Acunto et al., 2023b). We find large cross-sectional variation in each of the three scenarios. The average respondent expects annualized inflation of 8.4% over the next five years, with a standard deviation of 4.2 percentage points across respondents, and the low and high scenarios average 5.3% and 12.7%, with even larger dispersion in the high one (Table 5 in Appendix C). We also find significant differences between the moments coming from the historical data and the households' forecasts. The average expected scenario is 1.4 percentage points above the historical mean of roughly 7% shown in the treatment, and also above the 6.76% accumulated

by the index in the twelve months before the survey. In our setup, however, we cannot necessarily call this a bias, since deviations of forecasts from historical moments do not imply bias and we do not have a clear forecasting benchmark to compare the results with.

The three scenarios also allow us to describe the risk that respondents perceive around their expected scenario, which is not possible with point forecasts. Perceived downside risk averages 3.1 percentage points, almost the same as the 3 points implied by the historical distribution. Perceived upside risk, on the other hand, averages 4.3 percentage points, against 6 points in the historical data. Relative to history, our respondents expect a higher level of inflation but a smaller distance between their expected and their high-inflation scenario. This is relevant for the type of liability they hold since, as discussed in Section 1, the high-inflation scenario is the one that determines how costly an indexed contract can become.

Belief revisions. To measure how much respondents learn from the treatment, we follow Coibion et al. (2022) and Weber et al. (2025) and compare how treated and control respondents revise their beliefs as a function of how far their priors were from the signal. For each belief measure $x \in \{\Pi^L, \Pi^e, \Pi^H, \sigma^H, \sigma^L\}$ with signal s_x, we estimate

$$x_{i,\text{post}} - x_{i,\text{pre}} = \alpha + b\,(s_x - x_{i,\text{pre}}) + \delta\, T_{I,i} + \beta\,(s_x - x_{i,\text{pre}}) \times T_{I,i} + \varepsilon_i \ , \tag{13}$$

where $T_{I,i}$ indicates assignment to the inflation treatment. The coefficient b captures how control respondents revise their answers when asked a second time without receiving any information, and β is the additional share of the initial gap that treated respondents close. We refer to β as the learning rate induced by the inflation treatment: $\beta = 0$ means that the information has no effect, and $\beta > 0$ means that treated respondents revised their beliefs towards the signal more strongly than the control group.

Figure 4 plots revisions against the distance between prior and signal, for the expected scenario and for upside risk (for $x \in \{\Pi^e, \sigma^H\}$). These are the belief measures in which we focus in the paper, but the results hold for all measures (see Appendix C for results regarding $x \in \{\Pi^L, \Pi^H, \sigma^L\}$). A flat line at zero would mean no revision, and the dashed 45-degree line full adoption of the signal. The control line is not flat: respondents who saw no information also move toward the signal, revising down when their prior was high and up when it was low. This is common in these experiments and reflects noise and mean reversion in repeated answers (Coibion et al., 2023; Weber et al., 2025), which is why learning is measured by the gap between the two lines. The treated line is visibly steeper and closer to the 45-degree line, and the gap widens as priors move away from the signal. The estimated learning rate is 0.26 for the expected scenario and 0.31 for upside risk, both significant at the 1% level (see Appendix C for details on learning rate estimation).

These learning rates are in line with what the literature finds for economies with moderate to high inflation. Weber et al. (2025) run this type of regression for a large set of experiments and show that treatment effects are smaller when inflation is higher, since households are

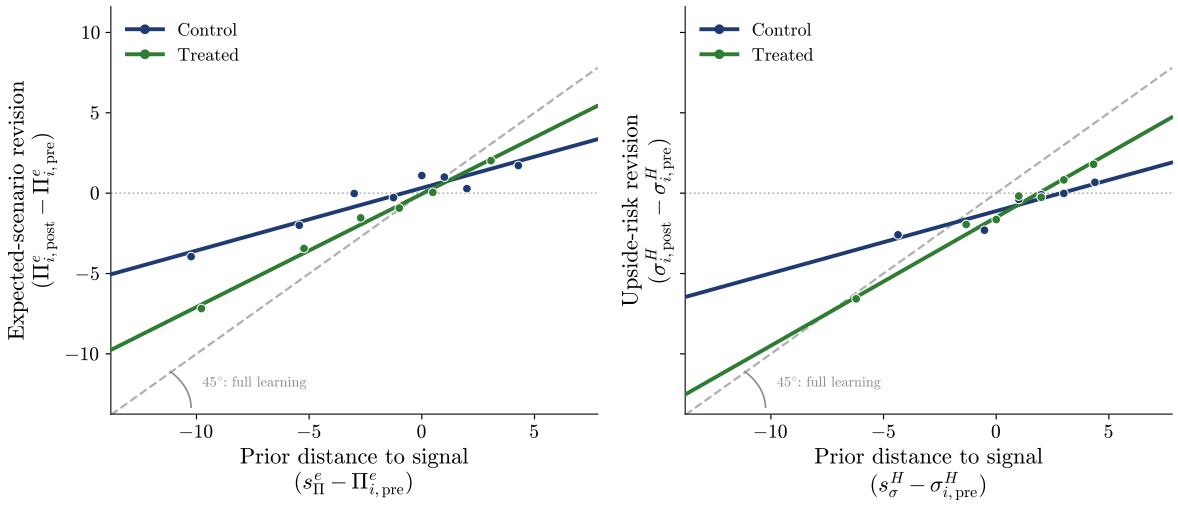


Figure 4. Belief revisions and prior distance to the signal. Binned scatter plots of belief revisions (posterior minus prior) against the distance between the signal and the prior, for the expected scenario (left) and for upside risk (right), with fitted lines by inflation-treatment arm. The dashed 45-degree line corresponds to full updating to the signal. All values in percentage points.

then already more informed. For U.S. households, the weight that treated respondents put on their priors fell relative to the control group by about 75% in 2018, when inflation was close to 2%, and by about 25% in 2021–2022. For euro-area households the corresponding number was about 20% at the end of 2021, when inflation was around 5%, and their pooled estimates imply effects of about 30% at an inflation rate like the 6.76% observed for our index, reaching zero when inflation is close to 11.5%. Our estimates are comparable to these numbers and, if anything, somewhat above them.[20]

Two features of our treatment should be kept in mind in this comparison. First, the experiments above provide a single number about the level of inflation, such as the most recent inflation rate, the central bank's target or a professional forecast, and measure beliefs about the following twelve months. We ask about a five-year horizon and show moments of the historical distribution of the index, for the reasons discussed in Section 3.3. Households that follow current inflation closely may still know little about its long-run distribution, which could explain why learning rates remain sizable in a setting where inflation is well above the levels of advanced economies. Second, because the signal contains a low and a high scenario in addition to the historical mean, it also moves the risk that respondents perceive around their expected scenario. In this respect our treatment is closer to experiments that provide information about second moments (Coibion et al., 2024; Kumar et al., 2023) than to the usual inflation treatment. The learning rate for upside risk is similar in size to the one for

[20] Weber et al. (2025) scale the change in slope by the weight that the control group puts on its prior. Our control respondents also revise toward the signal, so scaling our estimates in the same way makes them larger than the numbers reported here.

the expected scenario, so the treatment generates variation in both objects, which we use in Section 6.

5.2 Mapping Treatment and Pass-Through Parameters

The mapping treatment—thoroughly described in Section 3 and Appendix B—targets households' beliefs over their future nominal payments. Our question regarding such payments has the objective of verifying if beliefs about future installments are consistent with beliefs about future inflation. To verify such consistency and properly measure the pass-through of one object to the other, we define $\theta_i \in \mathbb{R}$ such that

$$\text{PMT}_i^e = \text{PMT}_{it}\,(1 + \theta_i\ \Pi_i^e)^5 \quad , \tag{14}$$

where PMT_i^e represents the belief about nominal installments in 5 years and Π_i^e is the expected annualized inflation defined in Section 3. The parameter θ_i is inspired by the exponential-growth bias literature in household finance (Stango and Zinman, 2009), and it should be interpreted as a measurement of perceived pass-through of inflation to installments at the individual level. It can be isolated from Equation 14 and is computed as a residual based on responses for PMT_{it} and Π_i^e—we formalize this on Appendix A. When $\theta_i = 1$, individual i is acting rationally and properly mapping beliefs about inflation to her payment projections. When $\theta_i < 1$, this implies an incomplete pass-through and that this individual is under-projecting their own installments (conditional on their own beliefs). Finally, when $\theta_i > 1$, we have a case of over-projection.

Pass-Through in the Control Group. Table 3 reports the distribution of θ_i by treatment arm, and Figure 5 plots its cumulative distribution. Among control respondents the median pass-through is 0.898, close to the rational benchmark, but this average hides rich patterns: only 20.3% of them meet the condition of $\theta_i \in [0.95, 1.05]$—a narrow interval around full pass-through. The distribution is wide, with an interquartile range of 0.608, and asymmetric. A quarter of control respondents report θ_i below 0.438, meaning they carry less than half of their own expected inflation into their own installment, and 27.4% fall below 0.5.

Two things are worth emphasizing about this result. First, under-projection is not a failure to understand the contract. Awareness of indexation was measured in the pre-treatment block, and every respondent—including those who reported not knowing their installments were indexed—was given an explicit explanation of the indexation rule before any of these questions were asked. Respondents who under-project know that inflation raises their payments, state a number for how much inflation they expect, and still report a payment inconsistent with their own answer. Second, we do not think this pattern has a single cause, and our design cannot separate the possible ones. A natural candidate is financial literacy, which

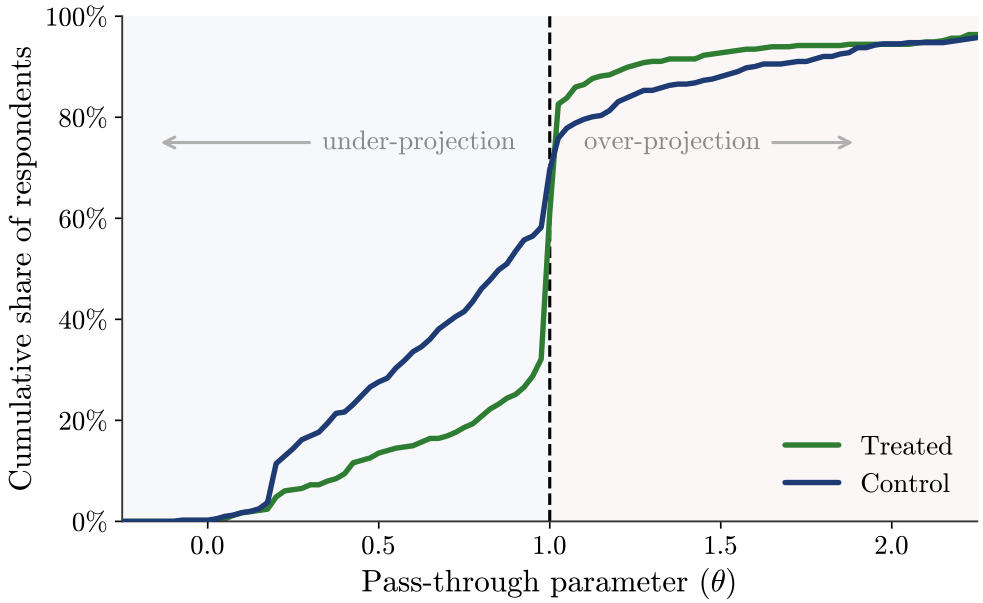


Figure 5. Cumulative distribution of the belief pass-through parameter θ. Distribution of θ_i, defined in equation (14), by mapping-treatment arm. Values below one indicate under-projection: the respondent carries less than the full amount of their own expected inflation into their projected installment. The treated distribution steps sharply at $\theta = 1$, while the control distribution climbs gradually across the whole support.

goes deeper than knowing how to perform mathematical compounding. Even a respondent who ignores compounding altogether, and simply adds five years of inflation to the current installment, would report a pass-through of about 0.85 to 0.9 at the inflation rates our respondents expect. A quarter of control respondents fall below 0.44, implying that many of them do not reach even this simple calculation, which points to a more basic difficulty with how adjustments accumulate over time. This is in line with the exponential-growth bias documented in household finance (Stango and Zinman, 2009), with its link to financial literacy (Almenberg and Gerdes, 2012), and with the broader evidence on how literacy shapes household decisions (Lusardi and Mitchell, 2014). Cognitive ability may also play a role: D'Acunto et al. (2023a) show that it shapes how households form expectations and act on them, and we do not measure it.

It is reasonable to suspect, however, that literacy might not explain all of it. Our sample is mostly educated, with two thirds of respondents holding a college degree, and pre-treatment covariates explain little of the variation in θ_i (Appendix E). Education is associated with smaller errors, but most of the dispersion remains unexplained. We believe that part of the incomplete pass-through patterns may come from how respondents approach the question. A household asked about its own installment five years ahead can answer quickly, starting from the payment it already knows, instead of stopping to work through the calculation. This is close to the distinction between fast, intuitive answers and slower, deliberate ones

	Control ($T_M = 0$)	Treated ($T_M = 1$)	All ($T_M \in \{0, 1\}$)
Pass-through θ (p.p.)			
P10	19.5	25.5	19.9
P25	43.8	84.4	59.1
P50	89.8	100.0	99.5
P75	104.6	100.0	100.4
P90	176.8	121.1	154.1
IQR	60.8	15.6	41.3
Share of respondents (%)			
$\theta \in [0.95, 1.05]$	20.3	53.9	37.5
$\theta < .50$	27.4	17.5	22.4
$\theta > 1.05$	25.1	15.9	20.4
N	379	395	774

Table 3. Distribution moments of the belief pass-through parameter θ. θ_i is defined in equation (14) and reported in percentage points, so 100 denotes full pass-through of a respondent's own inflation belief into their projected installment. Treated respondents were shown the compounded implication of their own average inflation scenario before reporting their expected installment; control respondents were not. Around nine (9) respondents had invalid answers for future installments and, therefore, had a θ mathematically undefined. We drop those observations to compute the results in this table.

(Kahneman, 2011): people who are able to solve a problem often give the intuitive answer when they do not stop to check it (Frederick, 2005), and estimates that start from a familiar value tend to stay too close to it (Tversky and Kahneman, 1974). We see these explanations as complementary and our setup does not allow us to properly separate them. What our data do show is that knowing the contract rule is not enough for respondents to properly pass their own inflation beliefs into their own perceptions over future payments.

Most importantly, the novelty in our findings is not the fact that individuals fail to do compounding arithmetic or are inattentive with regards to their own financial obligations. The most interesting feature of our setup is that the respondents of our survey are real households with real credit contracts of very long term maturities, with these indexed installments taking 14% of their monthly income on average. Even in such a context, under-projection is pervasive and widespread.

Pass-through in the Treatment Group. The mapping treatment induces a larger mass of respondents reporting a pass-through close to one (i.e., the rational benchmark). Among treated respondents the median pass-through is exactly 1.0 and the interquartile range falls from 0.61 to 0.16, a 74% reduction in dispersion. The share within five percentage points of one rises from 20.3% to 53.9%, and the share below 0.5 falls from 27.4% to 17.5%.

Figure 5 shows this as a near-vertical step at $\theta = 1$ in the treated distribution, against the gradual climb of the control.

This effect created by the mapping treatment is a crucial step to understand how the interaction between the mapping treatment and the inflation treatment can induce household behavior. The mapping arm delivers a large, exogenous shift in how households carry their inflation beliefs into their own payments, holding those beliefs fixed. It therefore lets us ask whether inflation information moves behavior only once its nominal implications are made visible: are households who correctly perceive their exposure more responsive to the inflation treatment than those who do not? Section 6 takes up this question.

6 Treatment Effects on Financial Decisions

In this section, we present our causal evidence on the role of our information treatments in explaining financial decisions. We leverage our cross-randomization of the two treatment arms to show how shifts inflation beliefs are only significantly affecting financial planning for the subsample that was exposed the nominal implications of their own beliefs. In our setup, financial planning is measured by two types of decisions: planned savings and willingness to switch to a fixed payment contract.

6.1 Empirical Strategy

The inflation treatment moves the whole subjective distribution, so a single measure of "how much beliefs shifted" would mix the level of expected inflation with the risk perceived around it. We therefore split the exogenous shift into two continuous measures of treatment intensity (which we refer to as "treatment doses"), built from the same prior-to-signal distances that drive the belief revisions of Section 5.1:

$$Z_i^e \equiv T_{I,i}\left(s_{\Pi}^e - \Pi_{i,\text{pre}}^e\right), \qquad Z_i^\sigma \equiv T_{I,i}\left(s_{\sigma}^H - \sigma_{i,\text{pre}}^H\right). \tag{15}$$

Z_i^e is how far the treatment pushes the expected scenario toward the historical mean, and Z_i^σ is the analogous push on perceived upside inflation risk. Both are zero for respondents who received no information, and among the treated they are weakly correlated (correlation of 0.13), so their coefficients can be read as two separate channels. Section 5.1 is their first stage: the learning rates reported there are the slopes of posterior revisions on these same distances.

Let $P_i \equiv \left(T_{I,i},\ \Pi_{i,\text{pre}}^e,\ \sigma_{i,\text{pre}}^H\right)'$ collect the inflation treatment and the two priors from which the doses are built. Let $T_{M,i}$ denote the mapping treatment. We estimate

$$Y_i = \alpha + \delta\, T_{M,i} + \sum_{k\in\{e,\sigma\}} \Big[\beta_k Z_i^k + \Delta_k\, Z_i^k \times T_{M,i}\Big] + P_i'\rho + \big(P_i \times T_{M,i}\big)'\phi + X_i'\lambda + \varepsilon_i, \quad (16)$$

where Y_i is the outcome of interest that could either be a switching decision $D_t(\Omega_t;\gamma)$ or the planned saving rate $\mathcal{S}_t$ (both defined in Section 2), and X_i collects the remaining controls. The coefficient β_k is the slope of behavior on dose k among respondents who did not receive the mapping treatment ($T_{M,i} = 0$), $\beta_k + \Delta_k$ is the slope among those who did ($T_{M,i} = 1$). Both are average treatment effects of the information, expressed per percentage point of distance between the respondent's prior and the signal (Haaland et al., 2023), and Δ_k is the main object of interest: how much more of their own belief shift households carry into their decisions once they have been shown what it implies for their own monthly payments. Standard errors are heteroskedasticity-robust throughout.

The vector X_i of control variables contains information on demographics, financial profile, contract information, among other features collected in the pre-treatment blocks. Demographics are age, sex, race and college education. The financial profile includes income bracket, employment type, preparedness for financial emergencies, and our income indexation measure described in Appendix D. Contract information includes the remaining months, whether the individual has alredy received credit through her consórcio contract, the log of current installment, payment-to-income ratio, and whether the respondent reported to be aware of contract indexation. We also control for the absolute distance between the respondent's forecast and the realized inflation rate. Unlike P_i, these variables enter only in levels and are the same in every specification.

Our identification relies on two features of the treatment assignment. First, the inflation treatment is randomly assigned, so conditional on the prior the dose Z_i^k is exogenous: respondents with the same prior differ in their dose only because some were shown the historical moments ($T_{I,i} = 1$) and others were not ($T_{I,i} = 0$). This is why P_i enters in levels—the dose is a function of the prior, and without the prior as a control β_k would confound the effect of the information with whatever else distinguishes respondents who started far from the signal. Second, the two treatments are cross-randomized, so $T_{M,i}$ is independent of both the dose and the prior. The mapping arms therefore differ only in the mapping treatment itself, and Δ_k compares two causal slopes rather than two groups of respondents who differ in who they are. Letting P_i interact with $T_{M,i}$ allows the priors and the inflation treatment to enter behavior differently across arms, so that Δ_k captures the change in the dose slope alone. Because $T_{M,i}$ is a randomly assigned treatment rather than a characteristic of the respondent, Δ_k is an interaction between two randomized treatments and not heterogeneity across respondents, and (16) is the fully interacted specification that makes it interpretable (Muralidharan et al., 2025).

Related literature tends to work with two-stage least squares specifications, but our design sticks to the reduced form results (Coibion and Gorodnichenko, 2026; Coibion et al., 2022; Coibion and Gorodnichenko, 2025). With about 783 respondents split across four treatment cells, instrumenting belief revisions with the doses gives confidence intervals too wide to be informative, and the weights that such estimators place on different respondents in passive-control designs can also complicate the interpretations of coefficients (Vilfort and Zhang, 2025; Haaland et al., 2023). The reduced form does not need these choices, although it does require us to interpret coefficients differently.

6.2 Planned Savings

Our first outcome is planned savings. In the final block of the survey, after both treatments and switching decisions, respondents answered the following question:

> *What percentage of your monthly income do you plan to save, on average, over the next 12 months? (Including savings for emergencies and future installment adjustments)*

Answers were given on a slider from 0 to 100. These are self-reported plans and not actual savings, which we do not observe. We treat the answer as the survey counterpart of the saving policy $\mathcal{S}_t$ of Section 2.3, and we use it only to compare respondents across treatment arms.

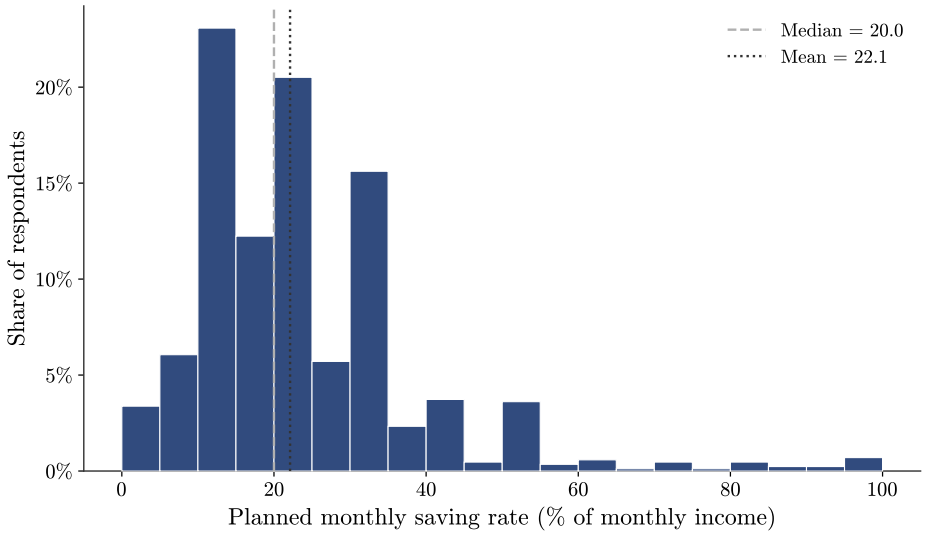


Figure 6. Planned saving rate. Distribution of the planned monthly saving rate as a share of monthly income over the next twelve months. Dashed and dotted lines mark the median and the mean.

Figure 6 shows the distribution of answers. The median respondent plans to save 20% of monthly income and the mean is 22.1%. The distribution is right-skewed, with strong

bunching at multiples of ten and a few respondents reporting rates above 60%. These levels are high and does not necessarily represent households' actual savings; these are stated intentions.

We estimate equation (16) with the planned saving rate as the outcome, so that coefficients are in percentage points of the saving rate per percentage point of dose. Our results show that informing individuals about inflation can only affect their planned savings when they are also exposed to the mapping treatment. Without showing the budget implications of their own inflation beliefs, we found no significant effects on planned savings.

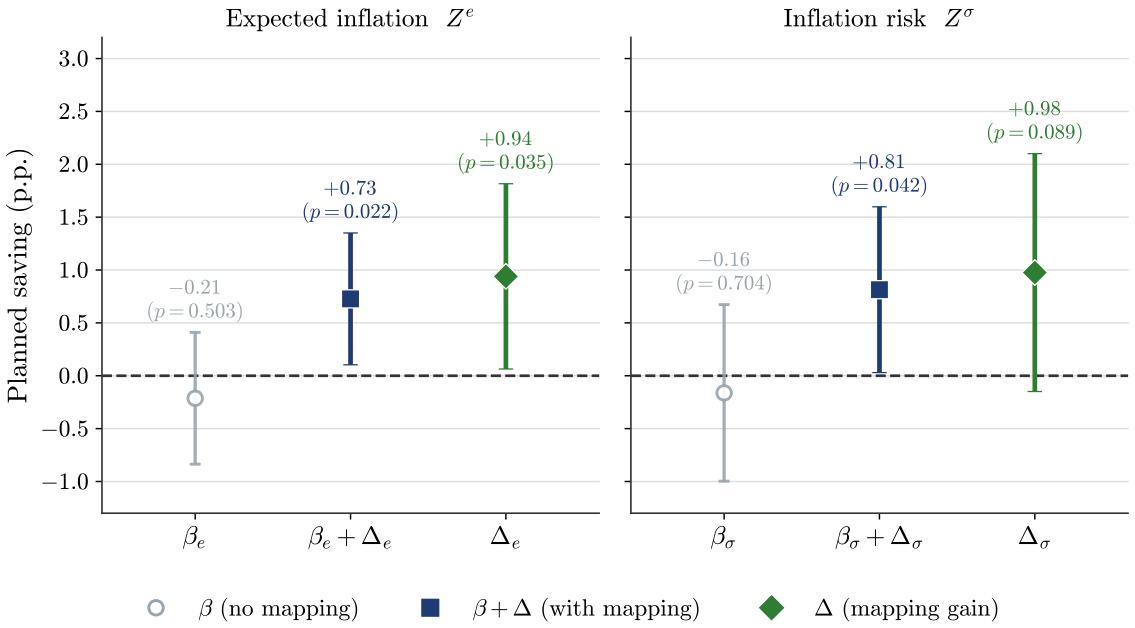


Figure 7. Treatment effects on planned saving. OLS estimates of equation (16) for the expected-inflation dose (left) and the upside-risk dose (right). Within each panel we plot the dose slope without the mapping treatment (β_k), the slope with it ($\beta_k + \Delta_k$), and the difference (Δ_k). Vertical bars are 95% confidence intervals from heteroskedasticity-robust standard errors; point estimates and p-values are reported above each marker.

Figure 7 shows that the dose effects are indistinguishable from zero without mapping ($p = 0.50$ and $p = 0.70$) but significant once combined with $T_{M,i} = 1$: for both doses the slope $\beta_k + \Delta_k$ for those exposed to the mapping treatment and the mapping gain Δ_k are significant at are significant at the 5% level, except the mapping gain on inflation risk, which is significant at the 10% level. Telling a household with $T_{M,i} = 1$ that the average historical inflation is 1 percentage point higher than its own forecast raises the their planned savings by 0.73 percentage points (this is given by the sum $\beta_e+\Delta_e$). To gauge the economic magnitude of such effect, a one-standard-deviation increase in the expected-inflation dose among treated households (about 5 percentage points) raises planned saving by roughly $0.73 \times 5 \approx 3.6$ percentage points—close to one-quarter of a standard deviation of the planned saving rate.

We also find such effect for inflation risk. Telling a household with $T_{M,i} = 1$ that historical upside risk is 1 percentage point higher than its own forecast raises their planned savings by 0.81 percentage points (the sum $\beta_\sigma + \Delta_\sigma$). In terms of magnitude, a one-standard-deviation increase in the inflation-risk dose among treated households (about 3.8 percentage points) raises planned saving by roughly $0.81 \times 3.8 \approx 3.1$ percentage points—around one-fifth of a standard deviation of the planned saving rate. Once again, inflation beliefs seem to matter, but only when their implications to monthly budgets are made explicit.

6.3 Switching to a Fixed-Payment Contract

Our second outcome is the decision to leave the indexed contract. Before the question, respondents were shown a visual description of how their contract works: as in Figure 1, they are exposed to a plot showing rising payments and a fixed (but initially higher) installment path alternative[21]. They were then offered the hypothetical choice of changing their current indexed contract to a fixed one, just as shown in the plot. The offer was stated in both nominal values (in Brazilian reais) and percentage terms, with the nominal amount computed from the respondent's self-reported current installment:[22]

> *Suppose that stopping the adjustments raised your installment to R$ _____ (an increase of about 30%), starting next month. Would you prefer that over your current contract?*

The premium was elicited by a bisection procedure. Every respondent was asked about +20% first. Those who accepted were then asked about +30%; those who declined were asked about +10%. Assuming monotonic switching decisions, we are able to recover answers for willingness to accept fixed payment contracts that are 10%, 20% and 30% higher than current installments. We read this survey question as the counterpart of the switching decision formalized in Section 2.2.

In the full sample, 77% accept at +10%, 58% at +20% and 32% at +30%. These rates seem low: the time remaining in contract is long for most respondents, and even the +30% premium buys certainty over the entire remaining term. We revisit this point in Section 7, using our theoretical model as a benchmark.

We estimate equation (16) with the switch-to-fixed indicator at the +30% offer as the outcome, scaled by 100 so that coefficients are in percentage points of the switching probability per percentage point of dose. As with planned saving, inflation information moves the switching decision only when households are also shown its budget implications. Figure 9 reports the estimates: without mapping the dose effects are indistinguishable from zero

[21] The plot is stylized and merely a visual resource to explain the tradeoffs in indexation. No numbers are included in the explanation.

[22] The nominal values in the offer are customized based on the respondents' answers throughout the survey.

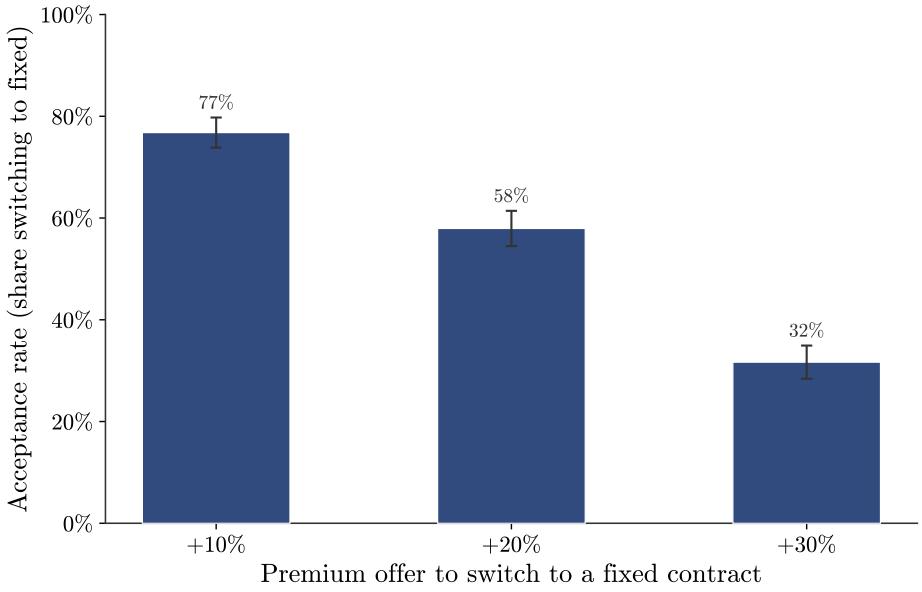


Figure 8. Acceptance of the switch to a fixed contract. Share of respondents preferring the fixed contract at each premium offer, in the full sample and by mapping arm.

($p = 0.81$ for Z^e and $p = 0.46$ for Z^σ), and the inflation-risk effect becomes positive and significant once combined with $T_{M,i} = 1$.

At the +30% offer the expected-inflation dose has no significant effect on switching, with or without mapping ($\beta_e + \Delta_e = -1.51$, $p = 0.13$; $\Delta_e = -1.28$, $p = 0.35$). We do find, however, significant effects for inflation risk. Telling a household with $T_{M,i} = 1$ that historical inflation risk is 1 percentage point higher than its own forecast raises the probability of switching by 3.06 percentage points (the sum $\beta_\sigma + \Delta_\sigma$, $p = 0.01$), and we can reject the null that $\Delta_\sigma = 0$ ($p = 0.03$). Relative to the 32% baseline acceptance rate, this is about a 10% increase in the likelihood of switching. The +30% offer carried the highest premium and the lowest baseline acceptance, yet this is exactly where beliefs move behavior—but only for the subgroup shown the nominal implications of their beliefs ($T_{M,i} = 1$), and not for those who were not ($T_{M,i} = 0$). As with planned saving, inflation beliefs seem to matter, but only when their implications to monthly budgets are mad explicit.

6.4 Discussion on Causal Estimates

Our results speak to a literature that has attempted (and often struggled) to connect inflation beliefs to household decisions. Survey correlations between inflation expectations and spending can be weak or unstable (Bachmann et al., 2015; Burke and Ozdagli, 2023), and information experiments that move expectations by large amounts produce behavioral responses that are often modest or different from our theoretical predictions (Coibion et al., 2023,0; Duca-Radu et al., 2021; Coibion and Gorodnichenko, 2026). Our inflation treatment

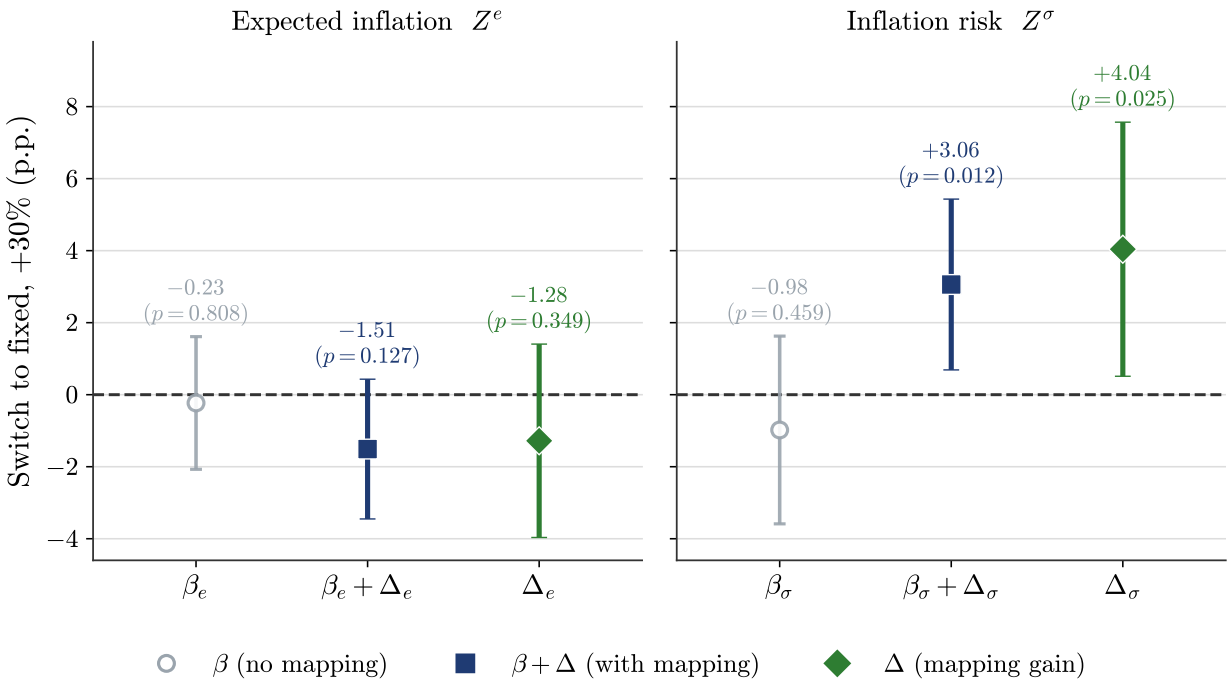


Figure 9. Treatment effects on the switch-to-fixed decision, for the 30% markup offer. OLS estimates of equation (16) for the expected-inflation dose (left) and the upside-risk dose (right). Within each panel we plot the dose slope without the mapping treatment (β_k), the slope with it ($\beta_k + \Delta_k$), and the difference (Δ_k). Vertical bars are 95% confidence intervals from heteroskedasticity-robust standard errors; point estimates and p-values are reported above each marker.

alone reproduces this pattern. It moves inflation beliefs, and it does not move switching decisions or savings in the arm that received no other information ($T_{I,i} = 1, T_{M,i} = 0$). Yang (2025) argues that a belief can be held with confidence and still not reach a decision when the household does not know how it maps into the choice at hand. Our mapping treatment eases part of that step by showing how the households' nominal payments are exposed to inflation. It adds no information about inflation, only the implication of the respondent's own beliefs for her own installment, and that is what turns the belief shift into a change in behavior. We believe that this carries a lesson for the design of information treatments: showing households inflation is not the same as showing them what inflation does to them.

The second lesson concerns which moment of the belief distribution matters. Coibion et al. (2024) show that, holding first moments fixed, higher perceived uncertainty lowers household spending. Our distinction between expected inflation and inflation risk lets us ask a similar question in the context of debt indexation. Telling a household with $T_{M,i} = 1$ that historical inflation risk is 1 percentage point higher than its own forecast raises the probability of paying a 30% premium to freeze its installments by 3.06 percentage points, and raises its planned saving rate by 0.81 percentage points. Both are precautionary responses to the

scenario in which installments rise the most: the fixed-payment contract buys insurance against it, and saving builds a buffer for it. Without the interaction with the mapping treatment, none of these results hold.

7 Rational Benchmark on Choices to Switch

This section brings an exercise using the consumption-savings model in Section 2. We answer a question that our empirical estimates are not able to handle: conditional on their own (possibly biased) beliefs about inflation, should households be willing to switch to fixed contracts? We answer this by calibrating our theoretical model using the survey measurements on household financial profiles and their inflation beliefs. Our results serve as a benchmark on what would a rational agent decide even under imperfect information about inflation.

7.1 Indifference Between Fixed and Indexed Contracts

We now define the object that we measure using our model. We do so at the respondent level. By extending our framework in Section 2.2, we define the indifference thresholds between fixed and indexed contracts. This is performed by defining the γ^* object, which denotes the maximum installment increase a given household with indexed liabilities is willing to take to switch to a fixed payment contract. From our model environment, we have $V_t(\Omega_t)$ denoting the households perceived value under an indexed liability path, and $\tilde{V}_t^{\text{fix}}\big(\Omega_t;\gamma\big)$ is the perceived value of a fixed payment alternative with a markup of $\gamma > 0$. Formally,

$$\gamma^*(\Omega_t) \equiv \sup\Big\{\gamma \geq 0 \;:\; \tilde{V}_t^{\text{fix}}\big(\Omega_t;\gamma\big) \;\geq\; \tilde{V}_t(\Omega_t)\Big\}. \tag{17}$$

This is directly related to the switching decisions $D_t(\Omega_t;\gamma)$. It is trivial to show that, under a no-deflation scenario, we get $D_t(\Omega_t;0) = 0$. In other words, if deflation has zero perceived probability of occurring, a household will always choose to have fixed payments over an indexed path if they do not have to pay a premium for that.[23] It is also trivial to check that a higher γ monotonically decreases the value of alternative $\tilde{V}_t^{\text{fix}}\big(\Omega_t;\gamma\big)$. The object γ^* finds the threshold for γ such that households are indifferent between that fixed alternative and their current value.

7.2 Calibration

Our survey is tightly connected to the model. Each part of the survey was designed to be interpretable through the lens of the model, which means we have enough empirical objects

[23] Throughout our survey, our answer alternatives did not allow for individuals to report deflation scenarios. Therefore, our data and also our modeling choices always imply a no-deflation scenario—which we see as a very mild and reasonable assumption, especially for the Brazilian context.

to calibrate the state vector Ω_t and most of the parameters governing the perceived inflation and income process.

We solve the model at a quarterly frequency, aggregating the monthly primitives of Section 2.1. Table 4 lists every input. The coordinates in Panel A are set separately for each respondent from her survey answers, and the parameters in Panel B are common to all respondents.

Perceived inflation follows a Gaussian AR(1) in quarterly log units, with persistence ρ_π, mean μ_q and innovation standard deviation σ_q. For each respondent, we choose (μ_q, σ_q) so that the implied distribution of annualized inflation over the next five years has her average scenario $\Pi^e_{i,\text{post}}$ as its mean, and a 90% band as wide as the distance between her low and high scenarios, consistent with the interpretation in Section 3.3. These are posterior beliefs, elicited after the inflation treatment. Current inflation π_t is the realized 6.76% for respondents in the inflation treatment ($T_{I,i} = 1$), who were shown it, and the respondent's own estimate of the index over the last twelve months for those with $T_{I,i} = 0$. We also assume that all individuals fully understand the implications of inflation for their own contracts, so that the perceived installment path follows the law of motion in Equation (4) exactly, since this is a normative benchmark.

The remaining term T is the number of remaining months divided by three, and PMT_0 is the reported installment-to-income ratio. Initial assets A_0, in quarters of income, map the five answers on preparedness for financial emergencies to 0.15, 0.5, 1, 2 and 4.

The income process uses estimates from Engbom et al. (2022). They report a standard deviation of one-year residual log earnings changes of 0.380 among formal-sector workers in Brazil and 0.652 among informal-sector workers, which give quarterly innovations of $\sigma_y = 0.190$ and $\sigma_y = 0.326$. We assign the formal value to formal wage-earners and public-sector workers and the informal value to all other respondents, following their sector definition, which places the self-employed in the informal sector. The real drift μ_y corresponds to real income growth of 1% per year, slightly below the median real earnings growth they document between 1995 and 2018. The income pass-through and stickiness are respondent-specific: $\rho_{y,i}$ is the pass-through component of the income indexation score of Appendix D, and ψ_i maps the answer on how long income takes to catch up with inflation into a quarterly probability of no reset.

	Description	Value or source
Panel A. Respondent-specific, from the survey		
$\mu_q,\ \sigma_q$	Perceived inflation: mean, innovation s.d.	Three inflation scenarios (posteriors)
π_t	Current inflation	6.76% if $T_{I,i} = 1$; own estimate if $T_{I,i} = 0$
$\rho_{y,i}$	Income pass-through at a reset	Pass-through answer
ψ_i	Probability of no income reset (quarterly)	Speed-of-adjustment answer
T	Remaining term (quarters)	Remaining months / 3
PMT_0	Installment-to-income ratio	Reported share
A_0	Liquid buffer (quarters of income)	Emergency preparedness, 0.15 to 4
σ_y	Income volatility	0.190 formal / 0.326 other
Panel B. Common to all respondents		
χ	Relative risk aversion	2
β	Discount factor (annual)	0.96
r	Real interest rate (annual)	5%
ρ_π	Persistence of inflation (quarterly)	$0.55^3 \approx 0.17$
μ_y	Real income drift (quarterly)	0.00249

Table 4. Calibration. Panel A is set separately for each respondent from her survey answers; Panel B is common to all respondents. The model is solved at a quarterly frequency. σ_y and μ_y following other papers Engbom et al. (2022).

7.3 Results and Counterfactuals

We start by solving for γ_i^* under two sets of beliefs, both with full pass-through of inflation to installments: first the survey beliefs reported by each respondent, and second corrected beliefs, fixed at the physical parameters of the INCC process. Figure 10 plots both distributions. Under survey beliefs, the median respondent would accept a permanent increase of up to 54% of her installment to fix her payments, and the mean is 77%. Correcting beliefs trims the right tail, which is driven by respondents expecting very high inflation, and the median falls to 49%.

These thresholds sit far above the premiums we offered. The first three bars at each offer in Figure 11 compare the acceptance rates implied by γ_i^* with the ones observed in the survey. Under survey beliefs and full pass-through, the model says that 96% of respondents should accept the +10% offer, 89% the +20% and 79% the +30%; in the data, the rates are 77%, 59% and 33%. Correcting beliefs widens the gap further, to 98%, 94% and 86%, because the respondents whose thresholds move up are the ones who held unusually low views on future inflation.

The direction of this result is intuitive once we look at the numbers behind it. Few respondents report being financially constrained or carrying unusually high payment-to-income ratios, and most have a long time remaining in their contracts (see Table 1). On

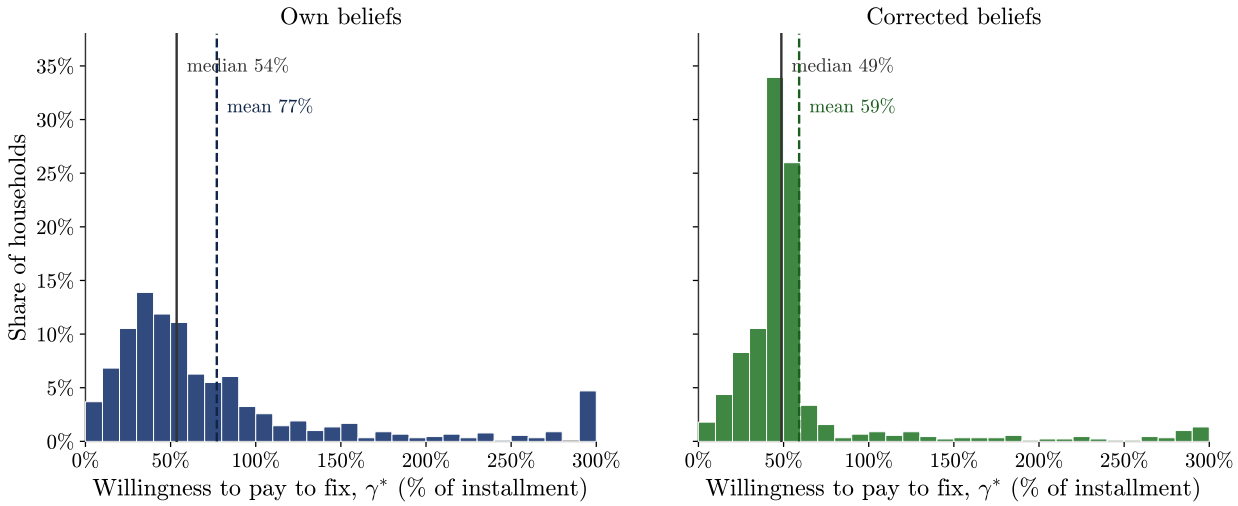


Figure 10. Willingness to pay to fix the installment. Distribution of γ_i^* across respondents under survey beliefs (left) and corrected beliefs (right), both with full pass-through. Solid and dashed lines mark the median and the mean. The bar at 300% collects respondents at the upper bound.

top of that, the historical mean of annualized INCC is close to 7%. Against a liability that compounds at that rate for another decade (the median respondent has 16.6 years left in their contract), a one-time increase of 10% is a great deal, and so are 20% and 30%. Beliefs do create dispersion in γ_i^*, as Figure 10 shows, but under almost any belief in our sample the offers fall below the threshold.

7.4 Incorporating Incomplete Pass-Through

Can we explain switching rates by incorporating the pass-through parameters θ_i measured in Section 5.2? We answer this by replacing the payment rule in Equation (4) with each respondent's perceived law of motion for her indexed installment,

$$\mathrm{PMT}_{t+1} = \exp(\theta_i \, \pi_{t+1}) \, \mathrm{PMT}_t, \tag{18}$$

where θ_i is set at the value we measure for each respondent. A household with $\theta_i < 1$ perceives an indexed installment that rises more slowly than the contract implies, so the fixed contract looks less valuable to her.

The last bar at each offer in Figure 11 reports the acceptance rates under survey beliefs and survey pass-through. They fall to 90% at +10%, 76% at +20% and 64% at +30%, against 96%, 89% and 79% under full pass-through. The under-projection documented in Section 5.2 therefore matters for contract choice: it closes about a third of the distance to the data at +10% and +30%, and about two-fifths at +20%. It is not enough, however. The model still predicts acceptance rates 13, 17 and 31 percentage points above the data, and

the gap is largest at the premium where households in our survey are least willing to switch.

The data point in the same direction. The mapping treatment moves most treated respondents close to full pass-through (Table 3), yet acceptance rates are nearly identical across the two mapping arms (Figure 8). If under-projection were what kept households in their indexed contracts, the treated arm should switch more, and it does not.

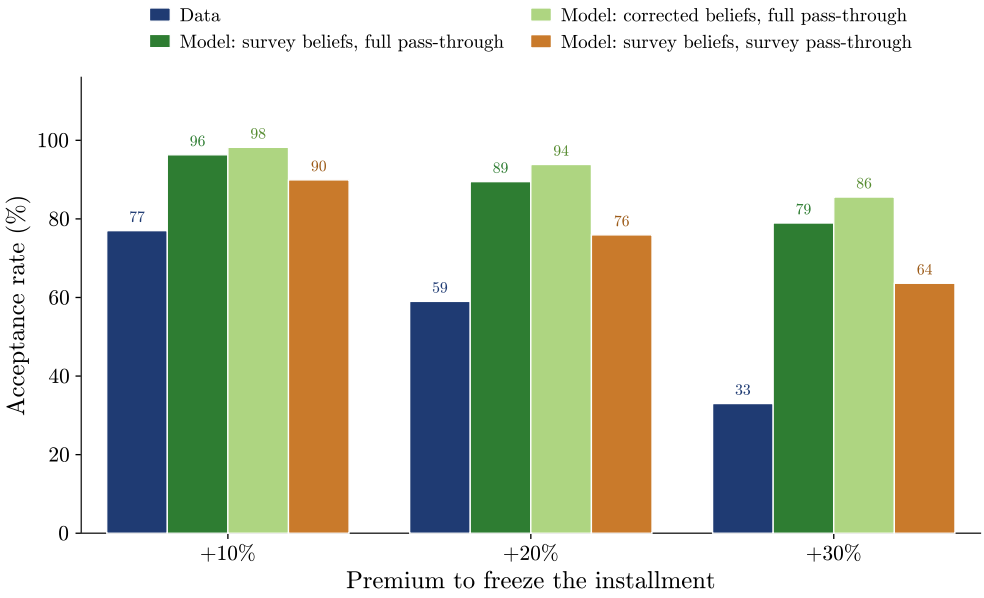


Figure 11. Acceptance of the switch to a fixed contract, model and data. Share of respondents accepting each premium offer in the survey, and implied by the model ($\gamma_i^* \geq$ premium) under three specifications: survey beliefs with full pass-through, corrected beliefs with full pass-through, and survey beliefs with the pass-through θ_i measured in the survey.

7.5 Discussion

We read this gap between data and model-predicted switching rates as a measure of how much remains to be understood about these decisions. Our experiment shows that inflation beliefs can shift contract preferences once their nominal implications are made explicit; this benchmark shows that the resulting choices are still far from what the beliefs themselves would justify. The pattern is familiar from household finance: borrowers leave large gains from refinancing on the table (Keys et al., 2016; Andersen et al., 2020; Berger et al., 2021), and choices between adjustable and fixed rates do not incorporate the full term structure (Koijen et al., 2009; Badarinza et al., 2018). Modeling the decision rules behind this behavior requires departures from the rational benchmark, in which agents attend to only part of the problem (Sims, 2003; Gabaix, 2014), perceive its elements with noise (Woodford, 2020), are uncertain about their own reasoning (Enke and Graeber, 2023), think through it narrowly (Lian, 2021), or solve it imperfectly (Ilut and Valchev, 2023). Our evidence suggests that

the gap between measured beliefs and normative predictions is where those models have the most to add.

One pattern in our data points at what such a decision rule might look like. Figure 16 (Appendix F) plots acceptance rates against the respondent's own projected nominal installment growth: the correlation is very strong at every offer. Respondents seem to decide based on the nominal number in front of them rather than on the beliefs behind it, which leaves room for money illusion (Fisher, 1928; Shafir et al., 1997; Campbell and Vuolteenaho, 2004) and simple nominal heuristics (Stango and Zinman, 2009) to explain more of the decision than beliefs do. Our design, however, cannot separate the role of money illusion from the other channels, and we leave this question to future work.

8 Concluding Remarks

This paper combines a consumption-savings model, an administrative dataset and, most importantly, a survey experiment to show how the context of inflation-indexed debt is a rich environment to study the role of inflation beliefs in households' decision making. We started by using the example of inflation-indexed mortgages originated in Brazil to show how initial installments are cheaper at origination, but exposure to inflation can quickly end such benefit during inflationary episodes. We then showed how our survey respondents—households with actual inflation-indexed liabilities—largely underestimate how inflation affects their future nominal payments.

Our main contribution comes from the cross-randomization of two treatments that target (i) inflation beliefs, and (ii) the pass-through of inflation beliefs into future nominal payments. Our empirical results show that providing information about inflation alone (the inflation treatment) did not change household financial planning. Combining the inflation treatment with the treatment targeting the pass-through (the mapping treatment), however, did. In other words, news about inflation only induced behavior when combined with explicit information about how inflation affects households' monthly budgets. We believe that our findings carry important implications for how information RCTs study the relation between macroeconomic beliefs and individual economic decisions.

Finally, we use our consumption-savings model to study household preferences for inflation-indexed versus fixed-payment contracts. We show that, even accounting for possibly biased inflation beliefs and for the under-projection of future payments, a rational benchmark fails to explain these preferences: respondents decline fixed-payment contracts even when they are clearly advantageous through the lens of the model. Our intervention targeted one observable component of the belief-to-action map, the pass-through of inflation into nominal payments, but accounting for it in the model closes only part of the gap. Further empirical and theoretical investigation on the belief-to-action mapping remains a natural direction for future research.

References

Almenberg, J. and Gerdes, C. (2012). Exponential growth bias and financial literacy. *Applied Economics Letters*, 19(17):1693–1696.

Andersen, S., Campbell, J. Y., Nielsen, K. M., and Ramadorai, T. (2020). Sources of inertia in household finance: Evidence from the Danish mortgage market. *American Economic Review*, 110(10):3184–3230.

Bachmann, R., Berg, T. O., and Sims, E. R. (2015). Inflation expectations and readiness to spend: Cross-sectional evidence. *American Economic Journal: Economic Policy*, 7(1):1–35.

Badarinza, C., Campbell, J. Y., and Ramadorai, T. (2018). What calls to ARMs? International evidence on interest rates and the choice of adjustable-rate mortgages. *Management Science*, 64(5):2275–2288.

Baesel, J. B. and Biger, N. (1980). The allocation of risk: Some implications of fixed versus index-linked mortgages. *Journal of Financial and Quantitative Analysis*, 15(2):457–468.

Berger, D., Milbradt, K., Tourre, F., and Vavra, J. (2021). Mortgage prepayment and path-dependent effects of monetary policy. *American Economic Review*, 111(9):2829–2878.

Berger, D. W., Milbradt, K., Tourre, F., and Vavra, J. S. (2025). Optimal mortgage refinancing with inattention. *American Economic Review: Insights*, 7(4):497–515.

Binetti, A., Nuzzi, F., and Stantcheva, S. (2024). People's understanding of inflation. *Journal of Monetary Economics*, 148:103652.

Braga, O., Pugliese, J., and Teixeira, L. I. (2026). Consumer misperceptions and redistribution in financial products: Evidence from brazilian consórcios. Working Paper.

Buckley, R., Lipman, B., and Persaud, T. (1993). Mortgage design under inflation and real wage uncertainty: The use of a dual index instrument. *World Development*, 21(3):455–464.

Burke, M. A. and Ozdagli, A. (2023). Household inflation expectations and consumer spending: Evidence from panel data. *The Review of Economics and Statistics*, 105(4):948–961.

Campbell, J. Y. and Cocco, J. a. F. (2003). Household risk management and optimal mortgage choice. *Quarterly Journal of Economics*, 118(4):1449–1494.

Campbell, J. Y. and Vuolteenaho, T. (2004). Inflation illusion and stock prices. *American Economic Review*, 94(2):19–23.

Clara, N., Boutros, M., and Kartashova, K. (2024). The value of mortgage choice: Payment structure and contract length. SSRN Working Paper No. 5018311.

Coibion, O., Georgarakos, D., Gorodnichenko, Y., Kenny, G., and Weber, M. (2024). The effect of macroeconomic uncertainty on household spending. *American Economic Review*, 114(3):645–677.

Coibion, O., Georgarakos, D., Gorodnichenko, Y., and van Rooij, M. (2023). How does consumption respond to news about inflation? Field evidence from a randomized control trial. *American Economic Journal: Macroeconomics*, 15(3):109–152.

Coibion, O. and Gorodnichenko, Y. (2025). Inflation, expectations and monetary policy: What have we learned and to what end? NBER Working Paper 33858, National Bureau of Economic Research.

Coibion, O. and Gorodnichenko, Y. (2026). *Expectations Matter: The New Causal Macroeconomics of Surveys and Experiments*. Princeton University Press, Princeton, NJ.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2022). Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy*, 130(6):1537–1584.

Conselho Monetário Nacional (2019). Resolução nº 4.739, de 19 de agosto de 2019. Altera a Resolução nº 4.676, de 31 de julho de 2018. Banco Central do Brasil, Brasília.

Cukierman, A. (1975). Price-level-adjusted mortgages in Israel. In Modigliani, F. and Lessard, D. R., editors, *New Mortgage Designs for Stable Housing in an Inflationary Environment*, number 14 in Federal Reserve Bank of Boston Conference Series, pages 159–176. Federal Reserve Bank of Boston, Boston, MA.

D'Acunto, F., Hoang, D., Paloviita, M., and Weber, M. (2023a). Iq, expectations, and choice. *Review of Economic Studies*, 90(5):2292–2325.

D'Acunto, F., Hoang, D., and Weber, M. (2022). Managing households' expectations with unconventional policies. *The Review of Financial Studies*, 35(4):1597–1642.

D'Acunto, F., Malmendier, U., and Weber, M. (2023b). What do the data tell us about inflation expectations? In Bachmann, R., Topa, G., and van der Klaauw, W., editors, *Handbook of Economic Expectations*, chapter 5, pages 133–161. Elsevier, Amsterdam.

D'Acunto, F. and Weber, M. (2024). Why survey-based subjective expectations are meaningful and important. *Annual Review of Economics*, 16:329–357.

Duca-Radu, I., Kenny, G., and Reuter, A. (2021). Inflation expectations, consumption and the lower bound: Micro evidence from a large multi-country survey. *Journal of Monetary Economics*, 118:120–134.

Engbom, N., Gonzaga, G., Moser, C., and Olivieri, R. (2022). Earnings inequality and dynamics in the presence of informality: The case of Brazil. *Quantitative Economics*, 13(4):1405–1446.

Enke, B. and Graeber, T. (2023). Cognitive uncertainty. *Quarterly Journal of Economics*, 138(4):2021–2067.

Fazio, D., Ramadorai, T., Skrastins, J., and Van Doornik, B. (2024). Housing and fertility. Technical report, Working paper. Available at SSRN 5046571.

Fisher, I. (1928). *The Money Illusion*. Adelphi, New York.

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4):25–42.

Friedman, M. (1974). *Monetary Correction: A Proposal for Escalator Clauses to Reduce the Costs of Ending Inflation*. Occasional Paper 41. Institute of Economic Affairs, London.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129(4):1661–1710.

Green, R. K. and Wachter, S. M. (2005). The american mortgage in historical and international context. *Journal of Economic Perspectives*, 19(4):93–114.

Haaland, I., Roth, C., and Wohlfart, J. (2023). Designing information provision experiments. *Journal of Economic Literature*, 61(1):3–40.

Ilut, C. and Valchev, R. (2023). Economic agents as imperfect problem solvers. *The Quarterly Journal of Economics*, 138(1):313–362.

Instituto Brasileiro de Geografia e Estatística (2023). Censo demográfico 2022: Identificação étnico-racial da população — resultados do universo. Technical report, IBGE, Rio de Janeiro.

Instituto Brasileiro de Geografia e Estatística (2024). Censo demográfico 2022: Características dos domicílios — resultados preliminares da amostra. Technical report, IBGE, Rio de Janeiro.

Jaravel, X. (2021). Inflation inequality: Measurement, causes, and policy implications. *Annual Review of Economics*, 13:599–629.

Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York.

Kaplan, G. and Schulhofer-Wohl, S. (2017). Inflation at the household level. *Journal of Monetary Economics*, 91:19–38.

Keys, B. J., Pope, D. G., and Pope, J. C. (2016). Failure to refinance. *Journal of Financial Economics*, 122(3):482–499.

Koijen, R. S. J., Van Hemert, O., and Van Nieuwerburgh, S. (2009). Mortgage timing. *Journal of Financial Economics*, 93(2):292–324.

Kumar, S., Gorodnichenko, Y., and Coibion, O. (2023). The effect of macroeconomic uncertainty on firm decisions. *Econometrica*, 91(4):1297–1332.

Lian, C. (2021). A theory of narrow thinking. *The Review of Economic Studies*, 88(5):2344–2374.

Lusardi, A. and Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature*, 52(1):5–44.

Majer, T. (2026). When discounted rates end: The costs of taking action in the mortgage market. *The Review of Financial Studies*, 39(6):1700–1750.

Manski, C. F. (2004). Measuring expectations. *Journal of Economic Literature*, 42(4):1329–1359.

Modigliani, F. and Lessard, D. R. (1975). New mortgage designs for an inflationary environment. In Modigliani, F. and Lessard, D. R., editors, *New Mortgage Designs for an Inflationary Environment*, Conference Series No. 14. Federal Reserve Bank of Boston.

Muralidharan, K., Romero, M., and Wüthrich, K. (2025). Factorial designs, model selection, and (incorrect) inference in randomized experiments. *Review of Economics and Statistics*, 107(3):589–604.

Schnorpfeil, P., Weber, M., and Hackethal, A. (2023). Households' response to the wealth effects of inflation. NBER Working Paper 31672, National Bureau of Economic Research.

Shafir, E., Diamond, P., and Tversky, A. (1997). Money illusion. *The Quarterly Journal of Economics*, 112(2):341–374.

Shiller, R. J. (1997). Why do people dislike inflation? In Romer, C. D. and Romer, D. H., editors, *Reducing Inflation: Motivation and Strategy*, pages 13–70. University of Chicago Press, Chicago.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

Stango, V. and Zinman, J. (2009). Exponential growth bias and household finance. *The Journal of Finance*, 64(6):2807–2849.

Stantcheva, S. (2024). Why do we dislike inflation? *Brookings Papers on Economic Activity*, Spring:1–46.

Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131.

Van Doornik, B., Fazio, D., Ramadorai, T., and Skrastiņš, J. (2024a). Housing and fertility. CEPR Discussion Paper 19805, Centre for Economic Policy Research.

Van Doornik, B., Gomes, A., Schoenherr, D., and Skrastins, J. (2024b). Financial access and labor market outcomes: Evidence from credit lotteries. *American Economic Review*, 114(6):1854–1881.

Vilfort, V. and Zhang, W. (2025). Interpreting TSLS estimators in information provision experiments. *American Economic Review: Insights*, 7(3):376–395.

Vivan, G. F. A. (2025). Estabilizando a prestação nominal em contratos imobiliários com correção pela inflação. Nota técnica, Banco Central do Brasil.

Weber, M., Candia, B., Afrouzi, H., Ropele, T., Lluberas, R., Frache, S., Meyer, B., Kumar, S., Gorodnichenko, Y., Georgarakos, D., Coibion, O., Kenny, G., and Ponce, J. (2025). Tell me something I don't already know: Learning in low- and high-inflation settings. *Econometrica*, 93(1):229–264.

Weber, M., D'Acunto, F., Gorodnichenko, Y., and Coibion, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives*, 36(3):157–184.

Woodford, M. (2020). Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics*, 12:579–601.

Yang, J. (2025). On the decision-relevance of subjective beliefs. Working paper, University of California Santa Barbara. Updated September 2025.

A Appendix: Formalizing Belief Elicitation

This appendix formalizes the beliefs that are elicited throughout the survey. Let P_t denote the price level measured by the index at period t. Using a yearly frequency for t, the random variable for which we ask scenarios is denoted by $\Pi_t^{(t+5)}$, given by:

$$\Pi_t^{(t+5)} \equiv \left(\frac{P_{t+5}}{P_t}\right)^{\frac{1}{5}} - 1 \ .$$

This is the annualized variation of the price index between dates t and $t+5$. Our survey asks each respondent their beliefs about:

$$\underbrace{\mathbb{E}_t\left[\Pi_t^{(t+5)}\right]}_{\text{expected}} \quad , \quad \underbrace{\mathcal{Q}_t^{0.05}\left(\Pi_t^{(t+5)}\right)}_{\text{low}} \quad , \quad \underbrace{\mathcal{Q}_t^{0.95}\left(\Pi_t^{(t+5)}\right)}_{\text{high}} \ ,$$

where $\mathcal{Q}_t^{\tau}(x)$ measures the quantile τ of a random variable x under the information set in period t. Analogously, $\mathbb{E}_t[x]$ would represent its expected value conditional on information available at t. These are the objects that individuals are asked about when eliciting priors and posteriors for inflation. We, thus, index their beliefs by i and define:

$$(\Pi_i^L, \Pi_i^e, \Pi_i^H) \equiv \left(\mathcal{Q}_{it}^{0.05}\left(\Pi_t^{(t+5)}\right), \mathbb{E}_{it}\left[\Pi_t^{(t+5)}\right], \mathcal{Q}_{it}^{0.95}\left(\Pi_t^{(t+5)}\right)\right) \ .$$

The next step is to measure beliefs about future nominal payments. We ask individuals about their expected nominal payments 5 years ahead conditional on the annualized inflation being their own expected scenario Π_i^e. Using i to index each respondent, our survey asks individuals about:

$$\text{PMT}_i^e \equiv \mathbb{E}_{it}\left[\text{PMT}_{t+5} \mid \Pi_t^{(t+5)} = \Pi_i^e\right] \ .$$

which is their nominal projection in case their expected scenario for annualized inflation materializes. Our pass-through measure comes from the comparison between PMT_i^e and the proper compounding of Π_i^e. Note that the information set of individual i at period t contains their current installment PMT_{it}. The measure θ_i is obtained from the equation:

$$\text{PMT}_{i,t+5} = \text{PMT}_{it}\left(1 + \theta_i \ \Pi_i^e\right)^5 \quad \Longrightarrow \quad \theta_i = \frac{\left(\frac{\text{PMT}_i^e}{\text{PMT}_{it}}\right)^{\frac{1}{5}} - 1}{\Pi_i^e} \ .$$

A belief consistent answer would give $\left(\frac{\text{PMT}_i^e}{\text{PMT}_{it}}\right)^{\frac{1}{5}} = 1 + \Pi_i^e$ and, therefore, $\theta_i = 1$.

B Appendix: Survey Treatment Screens

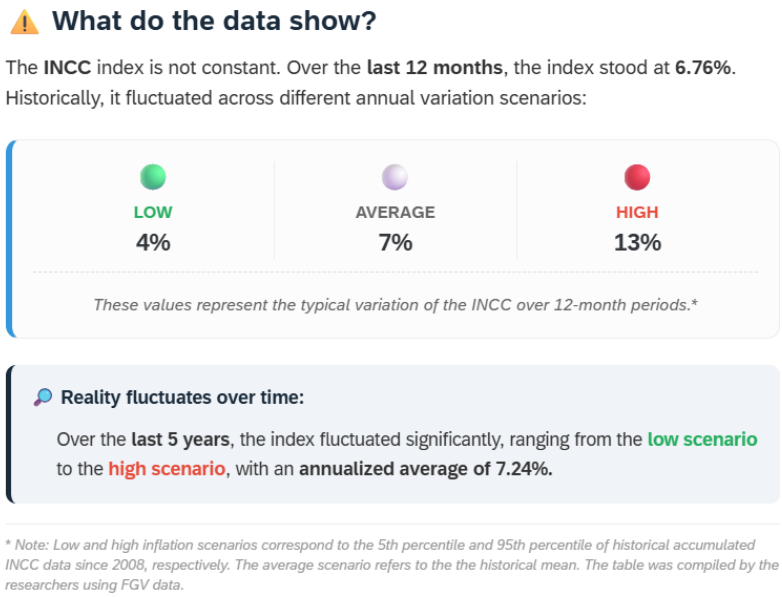


Figure 12. Inflation treatment. Historical INCC inflation is presented through low, average, and high annualized scenarios corresponding to the 5th percentile, median, and 95th percentile of the historical distribution. Original text was in Brazilian Portuguese.

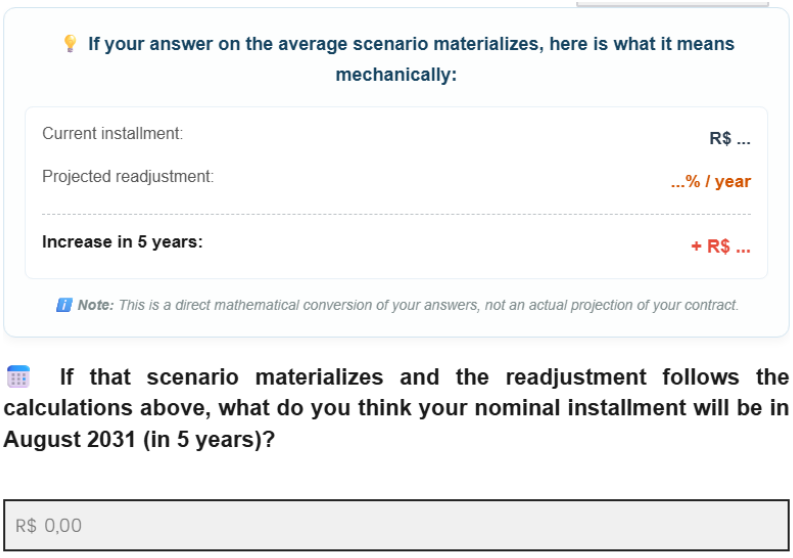


Figure 13. Mapping treatment. Respondents are shown the mechanical implication of their stated average inflation scenario for their current installment and then asked to report the nominal installment they expect to pay in five years.

C Appendix: Inflation Beliefs and Learning Rates

This appendix reports the descriptive statistics of prior beliefs and the learning rates discussed in Section 5.1. Table 5 describes the priors for the three scenarios and for the two risk measures, together with their distance to the signals shown in the inflation treatment. The signals are the 5th percentile, the median and the 95th percentile of the historical distribution of annualized INCC inflation, equal to 4%, 7% and 13%, so the implied signals for upside and downside risk are 6 and 3 percentage points. The distance to the signal is defined as the signal minus the prior. A negative average distance therefore means that priors are on average above the historical moment, which is the case for the low and expected scenarios. Since the signals are the same for all respondents, the dispersion of the distance is the same as the dispersion of the prior.

	Prior belief			Distance to signal		
	Mean	SD	Median	Mean	SD	Median
Scenario beliefs						
Low scenario Π^L	5.29	3.41	5.00	−1.29	3.41	−1.00
Expected scenario Π^e	8.42	4.23	8.00	−1.42	4.23	−1.00
High scenario Π^H	12.71	5.74	12.00	0.29	5.74	1.00
Risk measures						
Upside risk σ^H	4.29	3.34	4.00	1.71	3.34	2.00
Downside risk σ^L	3.14	2.62	3.00	−0.14	2.62	0.00

Table 5. Prior inflation beliefs and distance to the signal. Main sample ($N = 783$), in percentage points. Distance to signal is the signal minus the prior. Signals are 4, 7 and 13 for the low, expected and high scenarios, and 6 and 3 for upside and downside risk.

Table 6 reports the learning rate $\hat{\beta}$ from equation (13) for each belief measure. The first column uses the main sample. The second keeps only respondents with a college or graduate degree. Outside the low scenario, learning rates range from 0.26 to 0.39 in the full sample and from 0.21 to 0.36 among college-educated respondents, all significant at least the 5% level. The low scenario is the exception: its estimate is small and imprecise in the main sample ($\hat{\beta} = 0.15$), though it is significant among college-educated respondents.

	Full	College+
Scenario beliefs		
Low scenario Π^L	0.15	0.28**
	(0.10)	(0.12)
Expected scenario Π^e	0.26***	0.27***
	(0.08)	(0.10)
High scenario Π^H	0.30***	0.21***
	(0.06)	(0.08)
Risk measures		
Upside risk σ^H	0.31***	0.36***
	(0.09)	(0.11)
Downside risk σ^L	0.39***	0.33**
	(0.10)	(0.14)
Observations	783	514

Table 6. Learning rates by belief measure. Estimates of β in equation (13), the additional share of the initial gap to the signal closed by treated respondents relative to control. The College+ column restricts to respondents with a completed college degree or more. Robust standard errors in parentheses. Observations are the control and treated counts for the expected scenario. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D Appendix: Income Indexation Score

Our regressions use a single measure of how much of inflation reaches a respondent's income within a year. We build it from the pass-through and adjustment-frequency questions of Table 2:

$$\zeta_i = \rho_{y,i} \times \psi_i, \qquad \zeta_i \in [0, 1], \tag{19}$$

where $\rho_{y,i}$ is the share of inflation that eventually reaches income and ψ_i is the share of that adjustment realized within twelve months. A respondent who expects 75% of inflation to reach her income, but only after one to two years, gets $\zeta_i = 0.75 \times 0.40 = 0.30$.

Pass-through $\rho_{y,i}$, quantitative question	
10 p.p. or more	1.00
More than 5 p.p., less than 10 p.p.	0.75
5 p.p. or less	0.35
I do not know	qualitative question
Pass-through $\rho_{y,i}$, qualitative question (fallback)	
Yes, rises by about the same amount	1.00
Yes, but rises less than prices	0.50
It varies	0.50
No, stays the same regardless	0.15
Speed of adjustment ψ_i	
Less than 6 months	1.00
6 to 12 months	0.75
12 to 24 months	0.40
More than 24 months	0.20
I do not know	0.50

Table 7. Scoring rules for the income indexation score. Value assigned to each answer of the questions in Table 2. The score is $\zeta_i = \rho_{y,i} \times \psi_i$.

Table 7 lists the value assigned to each answer. The pass-through component comes from the quantitative question, and the qualitative one is used only for the 20.0% of respondents who answered "I do not know" there. We do not let the qualitative answer override the quantitative one because the two often disagree: among the 321 respondents who said their income stays the same regardless of prices, 71% went on to report a number on the quantitative question, and 28% reported pass-through above 5 percentage points. We read the qualitative question as picking up dissatisfaction with inflation more than actual pass-through, which is also why "stays the same" is scored at 0.15 rather than at zero. For the

speed component, "I do not know" is set at the midpoint. No respondent is left without a score.

Figure 14 shows the distribution of the two components and of the score. The score has a mean of 0.30, a median of 0.25 and a standard deviation of 0.23, with half of the respondents between 0.14 and 0.40. Only about 3% exceed 0.75, so very few respondents see their income as nearly fully indexed within the year.

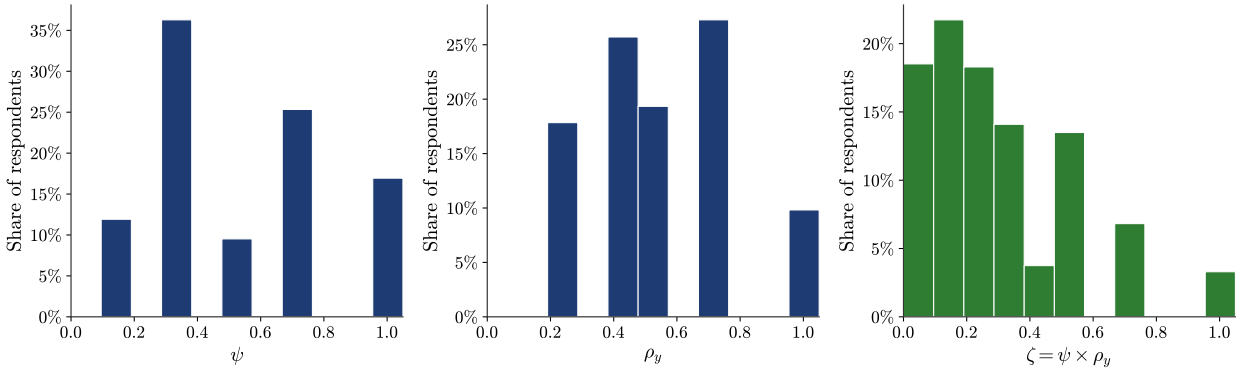
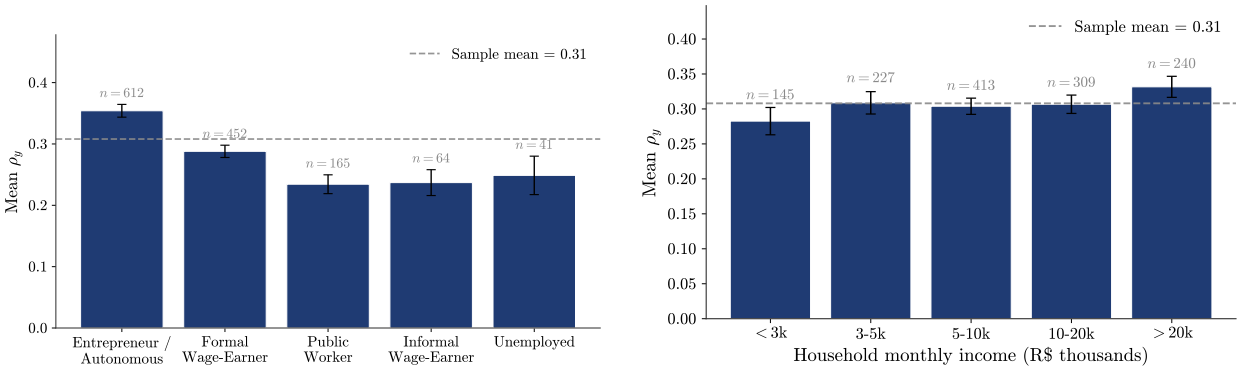


Figure 14. Income indexation score and its components. Distribution of the pass-through component $\rho_{y,i}$, of the speed component ψ_i, and of the score ζ_i defined in equation (19).



(a) By employment type **(b)** By household income

Figure 15. Mean indexation score by employment type and household income. Bars show the mean of ζ_i in each group, error bars one standard error, and the dashed line the sample mean. The lowest income brackets are grouped and the top bracket is "R$20,001 or more".

We check the score against two characteristics it was not built from. Across employment types (Figure 15, panel a) the ordering is what one would expect in Brazil. Entrepreneurs and self-employed workers, whose income follows the prices they charge, have the highest mean score (0.34). Formal wage-earners, whose wages are adjusted once a year through collective bargaining, come next (0.28). Public-sector workers have the lowest (0.23), with informal wage-earners at 0.21; the unemployed sit at 0.27, but with only 19 respondents.

Across income brackets (panel b) the score rises only mildly, from 0.27 in the lowest bracket to 0.34 in the highest. Perceived income indexation varies with how income is earned much more than with its level, so ζ_i is not standing in for income in our regressions.

E Appendix: Explaining the Pass-Through Parameter

Section 5.2 argues that incomplete pass-through is not reducible to what our observable measures of sophistication capture. This appendix reports the regressions behind that claim. We focus on the mapping-control arm ($T_M = 0$), since we already know that the mapping treatment is explicitly targeting the pass-through measurement. Continuous regressors are standardized, so coefficients read as the change in pass-through associated with a one-standard-deviation increase. We report two dependent variables: θ_i itself, which captures the direction of the projection error, and $|\theta_i - 1|$, which captures its magnitude regardless of sign.

The blocks are added in the following order: demographics (age, race), then financial and contract characteristics (emergency buffer, formal employment, installment-to-income share, income indexation, flexible contract, remaining months), then measured inflation knowledge together with Qualtrics timing metadata (page-submit times, total duration, number of fast pages). The relevance of this last block comes from the fact that it proxies for effort when responding to survey questions—it tracks how long a respondent actually spent on the page. It is the single largest contributor to fit, and it still leaves the adjusted R^2 at 0.086 for θ_i and 0.126 for $|\theta_i - 1|$. Pre-treatment covariates, contract exposure and engagement jointly leave the bulk of the dispersion in Figure 5 unexplained.

Education predicts the magnitude of the error but not its direction: college-educated respondents sit closer to full pass-through, by about 0.13 in $|\theta_i - 1|$, while their θ_i is statistically indistinguishable from that of other respondents. The same holds for gender and for having already been awarded credit.

Two further patterns also arise. First, respondents with larger installments under-project more and those in higher income brackets under-project less. Second, the distance between a respondent's prior and the historical median inflation rate predicts $|\theta_i - 1|$ strongly and consistently across columns: respondents whose inflation beliefs sit far from the historical benchmark are also further from mapping those beliefs into their payments.

	(1)	(2)	(3)	(4)
Dependent variable: Pass-through θ				
log(installment)	-0.532^{***}	-0.537^{***}	-0.646^{***}	-0.647^{***}
	(0.182)	(0.186)	(0.220)	(0.222)
Income bracket	+0.116	+0.117	+0.203	+0.209
	(0.109)	(0.100)	(0.143)	(0.140)
Prior inflation (%)	−0.068	−0.064	−0.089	−0.008
	(0.117)	(0.119)	(0.119)	(0.120)
|prior − 6.76|	−0.016	−0.012	−0.007	−0.036
	(0.125)	(0.128)	(0.128)	(0.128)
College+	−0.076	−0.080	−0.105	−0.099
	(0.200)	(0.202)	(0.203)	(0.210)
Already contemplated	+0.044	+0.066	+0.007	−0.042
	(0.126)	(0.123)	(0.173)	(0.176)
Female	+0.018	+0.022	+0.033	+0.040
	(0.166)	(0.167)	(0.161)	(0.157)
R^2	0.114	0.115	0.133	0.152
Adj. R^2	0.095	0.091	0.095	0.097
Dependent variable: $\|\theta - 1\|$				
log(installment)	-0.358^{*}	-0.370^{*}	-0.457^{**}	-0.458^{**}
	(0.188)	(0.190)	(0.228)	(0.228)
Income bracket	−0.034	−0.048	+0.024	+0.029
	(0.105)	(0.094)	(0.140)	(0.133)
Prior inflation (%)	+0.069	+0.083	+0.063	$+0.169^{*}$
	(0.097)	(0.099)	(0.096)	(0.098)
|prior − 6.76|	−0.057	−0.064	−0.058	−0.114
	(0.108)	(0.112)	(0.111)	(0.109)
College+	−0.002	−0.018	−0.051	−0.052
	(0.183)	(0.184)	(0.183)	(0.189)
Already contemplated	−0.140	−0.114	−0.051	−0.081
	(0.119)	(0.113)	(0.150)	(0.150)
Female	+0.113	+0.124	+0.121	+0.115
	(0.153)	(0.154)	(0.148)	(0.142)
R^2	0.084	0.088	0.112	0.149
Adj. R^2	0.065	0.064	0.072	0.094
Control blocks included (both panels)				
Treatment arm (T_I)	✓	✓	✓	✓
Demographics	–	✓	✓	✓
Financials & contract	–	–	✓	✓
Knowledge & timing meta.	–	–	–	✓
Observations	379	379	379	379

Table 8. Who mis-projects? Pre-treatment correlates of the pass-through parameter θ. OLS, mapping-control arm only ($N = 379$). Columns add control blocks cumulatively. Continuous regressors are standardized; θ is winsorized at the 1st and 99th percentiles. Robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

F Appendix: Reduced-Form Extensions

	+10% (1)	+20% (2)	+30% (3)
$\beta_e + \Delta_e$ (pooled, no split)	-1.850**	-1.865*	-1.194
	(0.941)	(1.059)	(1.009)
Panel A. Lower income bracket			
$Z^e \times \mathbf{1}$[lower income]	-0.654	-1.182	-0.705
	(0.869)	(0.982)	(0.947)
$\beta_e+\Delta_e$, lower-income half	-2.110**	-2.335**	-1.474
	(0.981)	(1.114)	(1.068)
$\beta_e+\Delta_e$, upper-income half	-1.456	-1.153	-0.769
	(1.104)	(1.233)	(1.188)
Panel B. High installment share			
$Z^e \times \mathbf{1}$[high paym./inc.]	-0.557	-1.541	-0.338
	(0.853)	(0.997)	(0.908)
$\beta_e+\Delta_e$, high-share half	-2.090**	-2.529**	-1.339
	(0.999)	(1.113)	(1.076)
$\beta_e+\Delta_e$, low-share half	-1.533	-0.988	-1.002
	(1.076)	(1.232)	(1.147)
Panel C. High income indexation			
$Z^e \times \mathbf{1}$[high ρ_y]	+0.891	-0.131	-0.405
	(0.837)	(0.966)	(0.932)
$\beta_e+\Delta_e$, high-ρ_y half	-1.318	-1.944	-1.436
	(1.087)	(1.221)	(1.171)
$\beta_e+\Delta_e$, low-ρ_y half	-2.209**	-1.812	-1.031
	(0.977)	(1.119)	(1.066)
Panel D. Low emergency buffer			
$Z^e \times \mathbf{1}$[low buffer]	+1.128	+1.363	+2.002**
	(0.845)	(0.968)	(0.913)
$\beta_e+\Delta_e$, low-buffer group	-1.218	-1.101	-0.071
	(1.070)	(1.187)	(1.087)
$\beta_e+\Delta_e$, higher-buffer group	-2.346**	-2.464**	-2.073*
	(0.991)	(1.130)	(1.097)
N	702	702	702

Table 9. Heterogeneity of the expected-inflation dose in the switch-to-fixed decision. Each panel adds one interaction of Z^e with a binary moderator to equation (16), one regression per panel, in the attentive sample. Moderators are median splits, except the buffer, which is the two lowest categories of emergency-fund preparedness. The moderators enter the controls in levels, so each interaction is identified without the moderator's main effect; none is interacted with T_M, so the two with-mapping slopes differ only by the interaction term. HC1 standard errors in parentheses. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Planned	Switch to fixed		
	saving	+10%	+20%	+30%
	(1)	(2)	(3)	(4)
δ (mapping main effect, T_M)	+1.45	-5.21	-1.74	-0.54
	(1.48)	(4.41)	(5.17)	(4.93)
Panel A. Expected-inflation dose Z^e				
β_e (slope, no mapping)	-0.21	-1.18	-0.50	-0.23
	(0.32)	(0.87)	(1.03)	(0.94)
Δ_e (mapping gain)	+0.94**	-0.87	-1.34	-1.28
	(0.45)	(1.25)	(1.46)	(1.37)
$\beta_e+\Delta_e$ (slope, w/ mapping)	+0.73**	-2.05**	-1.83*	-1.51
	(0.32)	(0.89)	(1.02)	(0.99)
Panel B. Inflation-risk dose Z^σ				
β_σ (slope, no mapping)	-0.16	-0.22	-1.17	-0.98
	(0.43)	(1.27)	(1.51)	(1.33)
Δ_σ (mapping gain)	+0.98*	-1.54	+0.72	+4.04**
	(0.57)	(1.75)	(1.95)	(1.80)
$\beta_\sigma+\Delta_\sigma$ (slope, w/ mapping)	+0.81**	-1.76	-0.46	+3.06**
	(0.40)	(1.22)	(1.28)	(1.21)
Controls, P and $P \times T_M$	Yes	Yes	Yes	Yes
Mean of dep. var.	22.01	76.79	57.95	31.67
R^2	0.148	0.076	0.062	0.073
N	780	780	780	780

Table 10. Full mapping-interaction estimates: planned saving and switch-to-fixed. Each column is a separate regression of equation (16) on the full sample. Column (1) is an OLS of the planned saving rate (percentage points of monthly income); columns (2)–(4) are linear probability models with outcome $100 \times \mathbf{1}\{\text{switch to fixed}\}$ at each premium offer, so their coefficients are percentage points of the switching probability per point of dose. In addition to the dose slopes, the table reports δ, the mapping main effect (the coefficient on T_M that is not interacted with the treatment doses). For each dose Z^k it reports β_k, the slope without the mapping treatment; Δ_k, the mapping gain ($Z^k \times T_M$); and $\beta_k + \Delta_k$, the slope with mapping (joint variance from the fitted covariance). All columns carry the pre-period vector $P = (T_I, \Pi^e_{\text{pre}}, \sigma^H_{\text{pre}})$, its interaction $P \times T_M$, and the full demographic control set. HC1 standard errors in parentheses. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

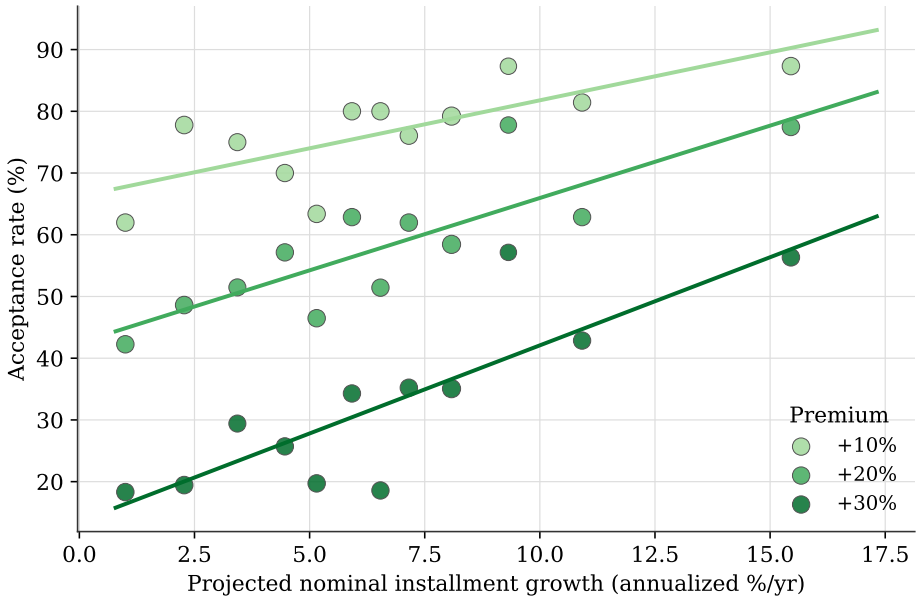


Figure 16. Acceptance rates and nominal installment projections. This binscatter reports the relation between acceptance rates and individual projections of nominal installment growth reported in the survey. There is a robust and strong correlation between the nominal projection and likelihood to accept switching, at every offer level.

G Appendix: Central Bank Data Analysis

G.1 Inflation Indices

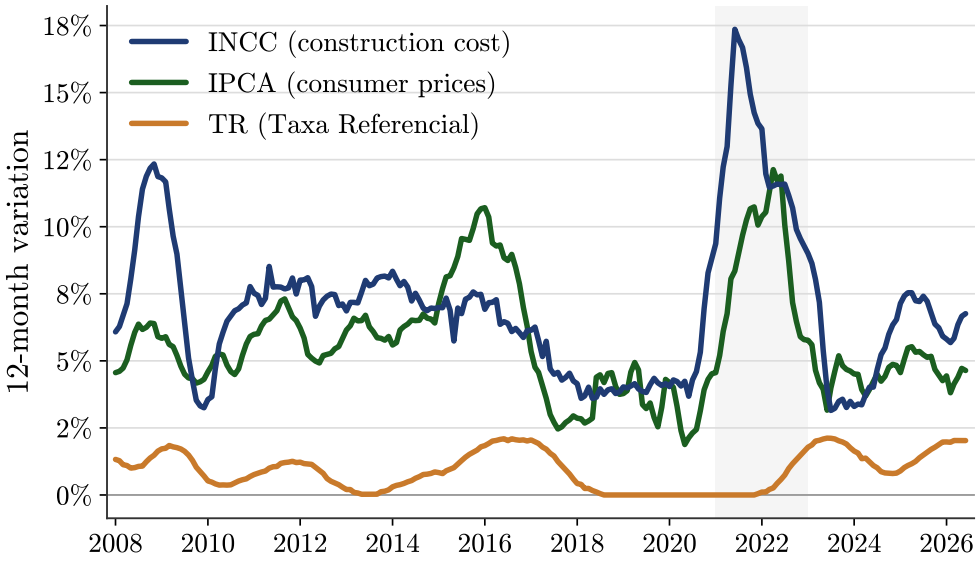


Figure 17. 12M Cumulative variation by index. This figure shows historical variation in the INCC index (used to adjust real estate consórcio contracts), the IPCA index (main consumer-price index in Brazil, used for mortgage contracts), the TR index (used to adjust the quasi-fixed rate mortgage contracts).

G.2 Sample construction

The extraction covers all credit operations reported to the SCR and classified as housing loans to individuals, originated between September 2019 and February 2020, funded through the SBPE segment of the SFH. The resulting cohort contains 119,906 contracts, which are classified as IPCA- or TR-indexed from the correction-index field of the registry.

Table 11 reports origination terms by contract type. The rate distributions barely overlap: TR contracts carry a mean effective rate of 8.37% per year over the index, while IPCA contracts cluster at a mean real coupon of 4.94% with a cross-sectional standard deviation of 0.03 percentage points — operationally a single posted rate.

G.3 Data and the initial installment gap

Our data come from the *Sistema de Informações de Crédito* (SCR), the supervisory credit registry maintained by the Central Bank of Brazil. All financial institutions operating in Brazil report to the SCR at the level of the individual credit operation for every exposure above a certain threshold, so coverage of housing credit is effectively universal. In our data,

	All	Inflation-Indexed	TR
Loans (N)	119,906	18,827	101,079
Effective rate (% p.a.)	7.83	4.94	8.37
(s.d.)	(1.49)	(0.03)	(0.88)
Installment share (% of balance)	1.047	0.837	1.086
median	0.920	0.707	0.944
Loan amount (R$ '000)	227.0	220.0	228.3
Maturity (months)	334.4	321.8	336.7
SAC amortization (%)	60.4	81.0	56.6
Price amortization (%)	8.4	2.6	9.5
Other amortization (%)	31.2	16.4	34.0

Table 11. Origination characteristics by indexation (SFH/SBPE cohort, Sept 2019–Feb 2020)

we observe the contract terms set at origination: loan amount, contractual interest rate, first scheduled installment, maturity, amortization system, and the index used for correction.[24]

In Brazil, most housing loans are made within the Sistema Financeiro da Habitação (SFH), which draws its funding mainly from regulated savings deposits. The standard SFH contract adjusts the outstanding balance by the Taxa Referencial (TR). The TR is a reference rate built from short-term deposit yields, it is not a price index (see Figure 17). It has remained close to zero since 2017, so in practice these loans work much like nominal loans with fixed payments. When we refer to indexed contracts, we mean those whose balance is adjusted by the IPCA instead. The time window in our sample includes a period in which both types of contracts were offered simultaneously.

We compare the two contracts after accounting for differences in loan composition by estimating

$$y_i = \alpha + \beta\,\mathrm{IPCA}_i + X_i'\gamma + \mu_{c(i)} + \lambda_{\ell(i)} + \varepsilon_i, \tag{20}$$

where y_i is either the contractual rate of loan i (in percentage points) or the log of its first scheduled installment. IPCA_i equals one for IPCA-indexed contracts. The vector X_i collects loan-level controls (amortization system, log loan amount, and maturity), and $\mu_{c(i)}$ and $\lambda_{\ell(i)}$ are origination-month and lender fixed effects. Standard errors are robust to heteroskedasticity.

Table 12 reports the unconditional gap and the gap under the full specification. IPCA contracts were originated at rates 3.4 percentage points below TR contracts unconditionally, and 4.2 points below comparable contracts from the same lender in the same month; the within-lender gap is larger because the lenders most active in the IPCA segment charged above-average TR rates. The corresponding first installment is roughly 45 log points—about

[24]Loan and borrower identifiers are cryptographically hashed; all analysis was conducted on Central Bank premises, and only aggregate statistics and regression coefficients leave the secure environment.

36%—lower than that of a comparable TR loan.

	(1) Unconditional	(2) Full specification
Effective rate (pp)	−3.430	−4.189
	(0.003)	(0.004)
Log first installment	−0.148	−0.453
	(0.005)	(0.003)
Origination-month FE		✓
Amortization, amount, maturity		✓
Lender FE		✓

Table 12. The origination gap: IPCA vs. TR contracts. Each cell reports the coefficient on the IPCA indicator from a separate regression estimate; $N = 119{,}906$. Robust standard errors in parentheses. All coefficients significant at the 1% level.

Table 12 also shows how the rate gap increases after accounting for contract features. Holding principal, maturity, and amortization system fixed, a coupon gap of 4.19 percentage points mechanically implies a first installment about 42 log points lower under either amortization schedule—close to the 45 log points we estimate. The installment advantage of the indexed contract is therefore almost entirely its rate advantage: borrowers were not being offered different loan sizes or maturities in disguise, but the same contract with the inflation premium stripped from the coupon.

Two limitations are worth flagging in our exercise. First, the comparison is descriptive: borrowers chose their index and lenders screened applicants, so β is a conditional mean difference in offered terms, not the causal effect of indexation on any individual borrower's price. Second, the extraction contains no loan-to-value ratio, borrower income, or region, so we cannot rule out that IPCA borrowers differ in unobserved risk. Neither caveat threatens the fact we document in these regressions, and the reason for that is visible in the raw data: IPCA rates are essentially uniform across the segment, and the unconditional gap is already 3.4 percentage points, so no configuration of omitted controls could overturn its sign or materially alter its magnitude.

G.4 Installment paths

The installment paths discussed in Section 1 consider a representative contract of each type on the constant-payment schedule, with a maturity of thirty years, a conventional contract priced at the cohort mean rate of 8.37%, and an indexed contract priced 4.19 percentage points below it, the within-lender gap of Table 12. These terms imply an indexed installment that begins at 65% of the conventional one. Given these origination terms, the

installment path requires no further data or assumptions: the indexed installment moves one-for-one with its correction index, so the path is pinned down by the realized history of the IPCA, absent any prepayment or restructuring.

Two remarks on scope. First, the conventional installment is held at its initial value. In the data it was essentially flat, drifting up mildly, by about 7% by mid-2026, only after the monetary tightening of 2021–22 lifted short rates enough to make the TR positive. Second, most contracts in the cohort follow the constant-amortization schedule, under which installments rise less than one-for-one with the index from a declining baseline; the qualitative picture—an indexed installment that converges toward, and under sustained inflation overtakes, its conventional counterpart—is the same.