Saving by Consuming: The Intertemporal Behavior of Hand-to-Mouth Households*

Mathieu Pedemonte
Inter-American Development Bank

Esteban Verdugo
University of Michigan
Job Market Paper

November, 2025

Latest Version

Abstract

This paper studies the forward-looking behavior of hand-to-mouth (HtM) households enabled by durable goods consumption. In a partial equilibrium model, HtM households use durables as imperfect substitutes for financial assets. Following an expectations shock, HtM households modify the composition of their consumption bundle of durables and non-durables to engage in intertemporal substitution. In contrast, financially unconstrained households rely on financial markets for consumption smoothing. We empirically test the model's predictions using a survey-based randomized controlled trial that shocks inflation expectations for HtM and non-HtM households. While both household types significantly revise their expectations, they differ in their intended reactions in consumption spending as predicted by the model, confirming that HtM households use durables as a consumption-smoothing mechanism. In a two-agent New Keynesian model, we find that durable goods reshape monetary policy transmission. While spending becomes more volatile, aggregate consumption stabilizes. The possibility of intertemporal substitution for constrained households implies that, when durability is sufficiently high, the aggregate dynamics of the heterogeneous-agent economy converge to those of the representative-agent framework.

Keywords: Consumption, Durable Goods, Hand-to-Mouth Households, and Inflation Expectations.
JEL Codes: C83 (Survey Methods · Sampling Methods), D12 (Consumer Economics: Empirical Analysis), D84 (Expectations · Speculations), E21 (Consumption · Saving · Wealth), and E31 (Price Level · Inflation · Deflation).

*Esteban is deeply indebted to Christopher House, John Leahy, and Stephen J. Terry for their invaluable advice, encouragement, and support. We are also grateful to Emilio Colombi, Joaquín Endara, Joshua K. Hausman, Edward S. Knotek II, Matthew D. Shapiro, Hiroshi Toma, Damián Vergara, Toni M. Whited, Triana Yentzen, and seminar participants at the University of Michigan and Universidad de Chile for their valuable comments and suggestions. Esteban gratefully acknowledges the Federal Reserve Bank of Cleveland for its hospitality and support during his 2023 AEA Summer Economics Fellowship. The experiment used in this project was registered in the American Economic Association's registry for randomized controlled trials under ID AEARCTR-0012835. The views expressed herein are those of the individual authors and do not necessarily reflect the official positions of the Inter-American Development Bank.
E-mails: Pedemonte: mathieupe@iadb.org. Verdugo: everdugo@umich.edu.

I Introduction

At the heart of the literature on consumption and financial constraints lies the *hand-to-mouth* (HtM) agent (Campbell and Mankiw, 1989). Excluded from financial markets, this agent consumes all of its income each period. As forward-looking behavior is associated with borrowing and saving opportunities, the HtM agent is myopic, confined to purely static decisions, and largely irrelevant for the study of expectation-driven dynamics. This view, however, is at odds with observed consumption patterns, where a sizable portion of consumption goods inherently contains an intertemporal dimension. Durable goods provide utility flows that extend beyond the purchase period, serving as a form of depreciating wealth that agents transfer over time. So all consumers make intertemporal choices when deciding how to allocate their spending across consumption goods. Since HtM households represent a large share of the population — roughly 40% in the U.S. (Aguiar et al., 2025) — understanding the intertemporal component of their consumption behavior is essential for assessing aggregate dynamics. This paper revisits the role of the HtM agent in macroeconomic models by incorporating durables into their consumption bundle. Combining experimental evidence with a theoretical framework, we demonstrate that once durables are accounted for, the HtM agent exhibits forward-looking behavior; thus, the expectations of HtM households are meaningful for both individual decisions and the aggregate economy.

The paper is structured in three parts. First, we propose a partial equilibrium model to illustrate how the HtM agent uses durables as a saving instrument to smooth consumption over time. The model predicts significant behavioral differences between the forward-looking HtM agent and a *non-hand-to-mouth* (non-HtM) agent — who enjoys access to financial markets. Second, we conduct a survey-based randomized controlled trial (RCT) and provide evidence consistent with the model's predictions. Finally, we move beyond the partial equilibrium framework and extend the standard two-agent New Keynesian (TANK) model by incorporating durable consumption to evaluate the aggregate implications of the HtM agent's intertemporal behavior. The possibility of intertemporal consumption smoothing for constrained households dampens their aggregate influence, implying that as durability increases, the aggregate dynamics of the heterogeneous-agent economy converge towards those of the representative-agent framework. Our results highlight the role of the HtM agent in expectations-driven dynamics. By showing that durables enable forward-looking behavior — and by providing causal evidence consistent with this mechanism — we refine the traditional characterization of the HtM agent.

To study the forward-looking behavior of the HtM agent, we develop a dynamic two-agent, two-good partial equilibrium framework in which the consumption bundle consists of a durable and a non-durable good. Although the HtM agent's total spending remains constrained by their current income in each period, they can exercise intertemporal behavior by modifying the composition of their consumption bundle within each period. When they expect high future marginal utility of non-durable consumption, they can reduce current non-durable spending to invest in durables, thereby expanding their future durable stock and freeing future resources to increase non-durable consumption. This way, despite lacking access to financial markets, the HtM agent smooths consumption over time by managing their durable accumulation. Because of the financial friction, any durable acquisition necessarily comes at the expense of current non-durable consumption, which is traded off not only against immediate durable utility but also against the stream of future utility services that durability provides. As a result, the agent's beliefs about future scenarios directly influence their current consumption choices, highlighting the relevance of expectations for reasons different than the intertemporal substitution effect associated with the traditional Euler equation.

Building on this forward-looking mechanism, we explore how agents adjust current durable and non-durable consumption in response to changes in inflation expectations. We show that the response of the forward-looking HtM agent notably differs from that of the non-HtM agent and the myopic HtM agent featured in the canonical TANK literature. The optimal responses of the agents depend on the financial friction, the depreciation rate of the durable good, and the anticipated variation in key relative prices. We use the model to derive predictions of agents' consumption reaction in two scenarios: higher aggregate inflation expectations and higher durable-goods inflation expectations. Under plausible conditions, higher aggregate inflation expectations imply an expected fall in real wages, inducing the HtM agent to accelerate durable purchases — the HtM agent buffers against anticipated real income losses by reallocating current spending from non-durables to durables. In contrast, the non-HTM agent relies on financial markets, reacting also to the lower real interest rate. When the resulting intertemporal substitution effect is larger than the one associated with the expected fall in real wages, the non-HtM agent reduces their net saving position, thereby financing higher overall current consumption. Thus, while the financial friction leads the HtM agent to expand durable accumulation at the expense of non-durable consumption when aggregate inflation expectations rise, the non-HtM agent increases consumption of both goods.

We then empirically examine the consumption reactions of HtM and non-HtM households to changes in inflation expectations and evaluate the model's predictions. To this end, we introduced a novel module into the Federal Reserve Bank of Cleveland's Survey of Consumer Expectations (Knotek et al., 2020) and conducted an RCT to causally estimate the effects of changing respondents' inflation expectations. We identify HtM and non-HtM respondents, elicit their inflation expectations, and randomly assign information treatments to generate exogenous variation in those expectations following standard practices in the literature (Coibion et al., 2023a). Using this variation, we estimate the causal effect of exogenous changes in inflation expectations on respondents' intended consumption, with a focus on heterogeneous reactions across agent types. The results are consistent with the model's predictions, validating that HtM households exhibit forward-looking behavior, using durable spending as a consumption-smoothing mechanism. For instance, we find that a 1 percentage point (p.p.) increase in aggregate inflation expectations among non-HtM respondents leads to a statistically significant rise of about 0.1 p.p. in their current consumption across all categories. Among HtM respondents, the effects are concentrated only in durable and semi-durable categories, also with magnitudes of around 0.1 p.p., while their responses for services and non-durable goods are not statistically different from zero. These patterns are robust to alternative definitions of HtM status.

After empirically establishing that the consumption choices of HtM households are forward-looking, we examine the aggregate implications of this behavior. To achieve this, we extend the standard TANK model by allowing agents to engage in durable consumption. We relax the traditional zero-net-supply constraint on bonds by introducing an endowment-owner agent that serves as the counterpart to households' savings, enabling non-HtM households to effectively save and respond differently than HtM households. Using this framework, we analyze the responses of the aggregate economy, as well as those of HtM and non-HtM households, when facing a monetary policy shock. Introducing durables as an intertemporal substitution channel fundamentally reshapes the transmission of monetary policy in heterogeneous-agent economies. Although durables amplify the volatility of households' spending — particularly among financially constrained households that front-load durable purchases following an expansionary monetary policy shock — it stabilizes aggregate consumption as HtM engage in consumption smoothing across periods. The possibility of durable consumption relaxes the consumption-income linkage typically characterizing the behavior of financially constrained households. More importantly, durables significantly dampen the impact of households' liquidity constraints on aggregate outcomes. As durability increases — i.e., the depreciation rate of durables decreases — the aggregate dynamics of the heterogeneous-agent economy increasingly resemble those of a single representative-agent framework with a household enjoying financial market access. Thus, durability introduces an intertemporal margin of consumption smoothing for financially constrained households, mitigating the traditional amplification result of monetary policy shocks typically attributed to these households, and making the heterogeneous-agent economy behave more like its representative-agent counterpart.

Related Literature. This paper contributes to several strands of the literature. First, we contribute to the literature exploring the intersection of financial constraints and consumption heterogeneity in shaping aggregate dynamics. Campbell and Mankiw (1989) documents the empirical failure of the representative agent Euler equation and introduces the idea that aggregate dynamics reflect a mixture of HtM and non-HtM households as a simple way to match observed high marginal propensities to consume. Building on this idea, Galí et al. (2007) extends the New Keynesian framework to allow for HtM consumers, showing that the constrained agent amplifies the real effects of fiscal policy. Bilbiie (2008) further explores the consequences of incorporating HtM households in a dynamic general equilibrium setting, showing that their presence affects the sensitivity of aggregate demand to the interest rate nonlinearly, significantly modifying monetary policy prescriptions. Although we abstract from relevant features emphasized in heterogeneous-agent New Keynesian (HANK) models (McKay et al., 2016; Kaplan et al., 2018) — such as idiosyncratic risk — TANK models provide a tractable tool to examine aggregate policy transmission while also accounting for key dimensions of household heterogeneity. The traditional TANK model captures relevant margins of aggregate consequences associated with household heterogeneity when evaluating monetary and other policies. Debortoli and Galí (2024) argue that, at least to some extent, a suitably specified TANK model — without idiosyncratic income risk — can capture the main channels through which heterogeneity shapes macroeconomic outcomes. This paper is also related to the extensive literature on durable consumption. Building on Hall (1978), early studies explore durable goods within the framework of the permanent income hypothesis (Mankiw, 1982; Bernanke, 1985; Caballero, 1990). In a seminal contribution, Grossman and Laroque (1990) develops an analytical model that involves durable consumption with fixed transaction costs, introducing the notion of infrequent durable purchases. This insight is later expanded in the subsequent literature on lumpy adjustment in durable goods (Bar-Ilan and Blinder, 1992; Caballero, 1993; Eberly, 1994; De Gregorio et al., 1998). Among these, Leahy and Zeira (2005) notably analyzed the cyclical implications of the timing of lumpy durable goods purchases within a general equilibrium setting. In a related but different strand, Barsky et al. (2007) showed that in sticky-price models, incorporating durables can generate strong real effects from nominal rigidities, while sufficiently long-lived flexibly priced durables can lead the monetary policy to have little to no effects on aggregate output. More recent contributions to this literature advance the field by emphasizing the extensive margin of durable adjustment and its nonlinear aggregate consequences, while considering general equilibrium feedback (Berger and Vavra, 2015; McKay and Wieland, 2021). By relying on a TANK model, our paper contributes to the literature by incorporating durables into the agents' consumption bundle and demonstrating that even financially constrained households can be forward-looking, thereby reshaping the role of expectation-driven dynamics in TANK environments. We highlight a novel channel through which durable goods provide an intertemporal margin of adjustment for HtM households, enabling consumption smoothing over time and significantly mitigating the amplification mechanism typically attributed to these households.

Finally, this paper contributes to the more recent empirical literature on inflation expectations and consumption, a relationship for which the evidence remains mixed (Binder and Brunet, 2021; Dräger and Nghiem, 2021; Burke and Ozdagli, 2023). For instance, some studies find little systematic link between inflation expectations and households' willingness to spend on durable goods (Bachmann et al., 2015), whereas others find that households expecting stable prices are significantly less likely to acquire durables compared to those anticipating higher inflation (Vellekoop and Wiederholt, 2019; Andrade et al., 2023). Our paper is closely related to the growing research using survey-based RCTs to identify causal effects of macroeconomic expectations (Armantier et al., 2016; Cavallo et al., 2017; Roth and Wohlfart, 2020), particularly those applying this approach to address the endogeneity between inflation expectations and households' consumption decisions (Coibion et al., 2022a, 2023b; Aktug and Atesagaoglu, 2024; Jiang et al., 2024). Recent work highlights important sources of heterogeneity in these responses. Candia (2024) concludes that consumption responses depend on the inflation environment, with households increasing durables purchases in high-inflation periods and reducing them in low-inflation periods. By concluding that higher inflation expectations lead households to increase durable spending once uncertainty is accounted for, Georgarakos et al. (2024) revises the opposite finding initially reported in Coibion et al. (2023a). Their interpretation aligns with emerging evidence that households believe nominal wages are sticky, implying that higher expected inflation translates into lower expected real income (Hajdini et al., 2022; Jain et al., 2024; Stantcheva, 2024). In related work, Roth et al. (2023) documents that non-HtM households respond to monetary policy news consistent with the Euler equation. We contribute to this literature by providing experimental evidence on heterogeneous consumption reactions between HtM and non-HtM households. Our results reveal that HtM households adjust their consumption plans in response to changes in inflation expectations, incorporating forward-looking considerations. Our empirical results highlight the relevance of nominal wage rigidities and the role of durable spending as a consumption-smoothing mechanism.

II A Two-Agent, Two-Good Economy

To uncover the mechanisms driving households' consumption responses when a forward-looking component is introduced through a consumption option, we begin by developing a partial equilibrium model featuring two agents and two consumption goods. As in the canonical TANK model, the only source of heterogeneity across agents is an exogenous differential access to financial markets (Galí et al., 2007; Bilbiie, 2008). This simplified setup provides a transparent environment to derive theoretical predictions, which we later confront with the empirical results of our experiment. In Section IV, we extend this partial equilibrium framework to a general equilibrium setting where we explore the aggregate implications of forward-looking behavior among HtM households. The partial equilibrium model thus forms the basis for the household block of the general equilibrium analysis.

II.A The Environment

Time is discrete and indexed by t. We consider two representative households with identical preferences, indexed by $i = \{\text{H}, \text{N}\}$, and characterized by heterogeneous participation in financial markets. We refer to the household that is exogenously and entirely excluded from financial markets as HtM ($i = \text{H}$), while the remaining saving household is the non-HtM ($i = \text{N}$).[1] At the end of each period, these infinitely lived households derive utility from consuming a bundle composed of non-durable goods and the services of the stock of durables they possess. For any given period t, we denote household i's consumption of non-durable goods as $c_{i,t}$ and the stock of durable goods this household holds at the beginning of the period as $d_{i,t-1}$. During this period, this stock depreciates at a rate $\delta \in (0, 1)$, and we denote the household i's net acquisitions of durables as $z_{i,t}$. The stock of

[1] In the literature, agents exogenously excluded from financial markets are also referred to as *rule-of-thumb* (Campbell and Mankiw, 1989), *non-traders* (Alvarez et al., 2001), *non-Ricardians* (Galí et al., 2004), and *non-asset holders* (Bilbiie, 2008), among others.

durables each household holds at the end of the period — which, together with non-durable consumption $c_{i,t}$, provides utility services to the household — is determined according to the following law of motion

$$d_{i,t} = z_{i,t} + (1-\delta)\, d_{i,t-1} \tag{1}$$

In every period, each household supplies an inelastic labor unit in exchange for the nominal wage $P_{w,t}$. Assuming costless reversibility for durables, and letting $P_{c,t}$ and $P_{d,t}$ denote the nominal market prices of non-durable and durable goods, the period t budget constraint of household i is

$$P_{c,t}c_{i,t} + P_{d,t}z_{i,t} = P_{w,t} + (R_{t-1}B_{t-1} - B_t) \cdot \mathbb{1}_{i=\text{N}} \tag{2}$$

where the left side represents the household's total consumption expenditure, while the right side represents the household's net income. The term $\mathbb{1}_{i=\text{N}}$ is an indicator function equal to one if the household has access to financial markets ($i = \text{N}$), and zero otherwise. The economy features a fully liquid financial instrument consisting of a one-period, non-contingent financial vehicle B_t offering a gross nominal return of R_t.

Given this setup, the control and endogenous state variables are specific to each household. For the HtM household, the control variables are $\{z_{\text{H},t}, c_{\text{H},t}\}$ and their only endogenous state variable is $d_{\text{H},t-1}$. For the non-HtM household, these variables are $\{z_{\text{N},t}, c_{\text{N},t}, B_t\}$ and $\{d_{\text{N},t-1}, B_{t-1}\}$, respectively. Therefore, given a path for the set of exogenous state variables $\{P_{d,k}, P_{c,k}, P_{w,k}, R_k\}_{k\geq t}$, and each household's initial holdings of durable goods and bonds $\{d_{i,t-1}, B_{t-1} \cdot \mathbb{1}_{i=\text{N}}\}$, the optimization problem each household faces corresponds to

$$\max_{\{z_{i,k}, c_{i,k}, B_k \cdot \mathbb{1}_{i=\text{N}}\}_{k\geq t}} \sum_{k\geq t} \beta^{k-t} \mathbb{E}_t \left[u\left(x_{i,k}\right) \right] \tag{3}$$

subject to the durables law of motion and the budget constraint. The parameter $\beta \in (0,1)$ is the households' common and exogenous discount factor, and $\mathbb{E}_t\,[\bullet]$ denotes the full information rational expectations operator conditional on the information available on period t. At the end of each period, each household receives utility from consuming a bundle composed of the non-durables acquired during the period ($c_{i,t}$) and the stock of durable goods held at the end of the period ($d_{i,t}$). We denote this consumption bundle as $x_{i,t} = x(c_{i,t}, d_{i,t}) : \mathbb{R}^2_+ \to \mathbb{R}_+$ and define the utility agents receive using the twice continuously differentiable, strictly concave function $u_{i,t} = u(x_{i,t}) : \mathbb{R}_+ \to \mathbb{R}$, which is assumed to satisfy Inada conditions. In this section, we assume that preferences are separable between non-durable and durable consumption, although most conclusions remain robust when this assumption is relaxed.

We use the non-durable good as the economy's numeraire. In consequence, the economy's real wage is $w_t = P_{w,t}/P_{c,t}$, and the relative price of durables in terms of non-durables is $q_t = P_{d,t}/P_{c,t}$. The real interest rate of the fully liquid bond is given by $r_{t-1} = R_{t-1}/\pi_{c,t}$, where $\pi_{c,t} = P_{c,t}/P_{c,t-1}$ denotes the gross inflation rate of the numeraire, and the saving household's real asset holdings are $b_t = B_t/P_{c,t}$. To rule out arbitrage between financial assets and durable goods, we impose the condition $r_{t-1} \geq (1-\delta)\pi_{q,t}$, where $\pi_{q,t} = q_t/q_{t-1}$ denotes the gross inflation rate of the economy's relative price between both consumption goods. This condition prevents the non-HtM agent from generating unbounded profits by borrowing to invest in durables.

II.B Economics of Durable Consumption

This section builds intuition on how durables introduce a forward-looking dimension into the behavior of the HtM household, partly offsetting their financial friction. To illustrate this mechanism clearly, we consider a simple two-state example. Starting in period t, households know the real value of their labor income w_t. However, in period $t+1$, a one-time shock occurs that leads the economy to one of two possible states of the world, indexed by $y = \{h, l\}$, where the only exogenous variable that changes across states is the households' real permanent labor income. There is a high-income state ($y = h$), which occurs with probability p, where households supply an inelastic labor unit in exchange for the real wage $w_k = w_h$ with $k \geq t+1$. There is a low-income state ($y = l$) — which occurs with the complementary probability $(1-p)$ — where the real wage is $w_k = w_l$, such that $w_h > w_l$. This uncertainty about future scenarios helps to illustrate the forward-looking component in the HtM behavior, and how their beliefs about the future shape their current consumption choices. For simplicity, we assume that the goods' relative price is one in every period. The households' optimization problem is represented by equation (3), where, in every period $k \geq t+1$, they face two distinct budget restrictions that must hold once the state of the world is revealed, not just in expectations.

Durable goods feature two key qualities that are essential to our analysis. Durables are storable, allowing households to accumulate them over time and effectively serving as a form of depreciating wealth that households transfer across periods, transforming the static HtM behavior into a dynamic one by introducing an intertemporal margin of adjustment. Second, they provide a stream of consumption services that directly contribute to the owner's utility, revealing their dual nature as an asset and a consumption good. Both elements are captured by the households' optimality conditions, which are summarized for the HtM agent as

$$mu^c_{\mathrm{H},t} \quad = \quad mu^d_{\mathrm{H},t} + \beta\left(1-\delta\right)\left[p \times mu^c_{\mathrm{H},t+1}\left(w_h\right) + \left(1-p\right) \times mu^c_{\mathrm{H},t+1}\left(w_l\right)\right] \tag{4}$$

where $mu^{\vartheta}_{i,t} = \partial u_{i,t}/\partial \vartheta_{i,t}$ denotes the marginal utility of the consumption good $\vartheta = \{d, c\}$. This optimality condition states that the cost of marginally reducing current non-durable consumption is offset by two benefits. First, the reduction in non-durable expenditure allows the household to marginally increase current durable consumption. Second, the additional durable stock serves as a storable asset with a negative gross return due to depreciation. The undepreciated portion of this additional stock is transferred over to the subsequent period, providing a discounted expected future consumption utility. The agent's choices reflect a balance between the intertemporal marginal costs and benefits of allocating resources across goods: current non-durable consumption is traded off against immediate durable utility, but also against durable future utility services. In the special case when durables fully depreciate every period ($\delta = 1$), they no longer allow intertemporal wealth transfer, effectively losing their asset-like characteristics. In this scenario, the optimization problem of the HtM agent becomes purely static and the optimality condition reduces to a purely intratemporal tradeoff between both consumption goods, i.e., $mu^c_{\mathrm{H},t} = mu^d_{\mathrm{H},t}$.

Equation (4) also reveals that the intertemporal benefits of durability raise the HtM household's optimal consumption of this good. Higher durability amplifies the intertemporal benefits by raising the discounted expected value of future consumption that the undepreciated portion of the durables stock provides. This widens the optimal gap between the current marginal utilities of non-durable and durable consumption, leading the HtM household to optimally favor durable consumption at the expense of non-durable expenditure. Additionally, because of the household's desire for

Figure 1: Optimal consumption bundle of the HtM household in the stripped-down model.

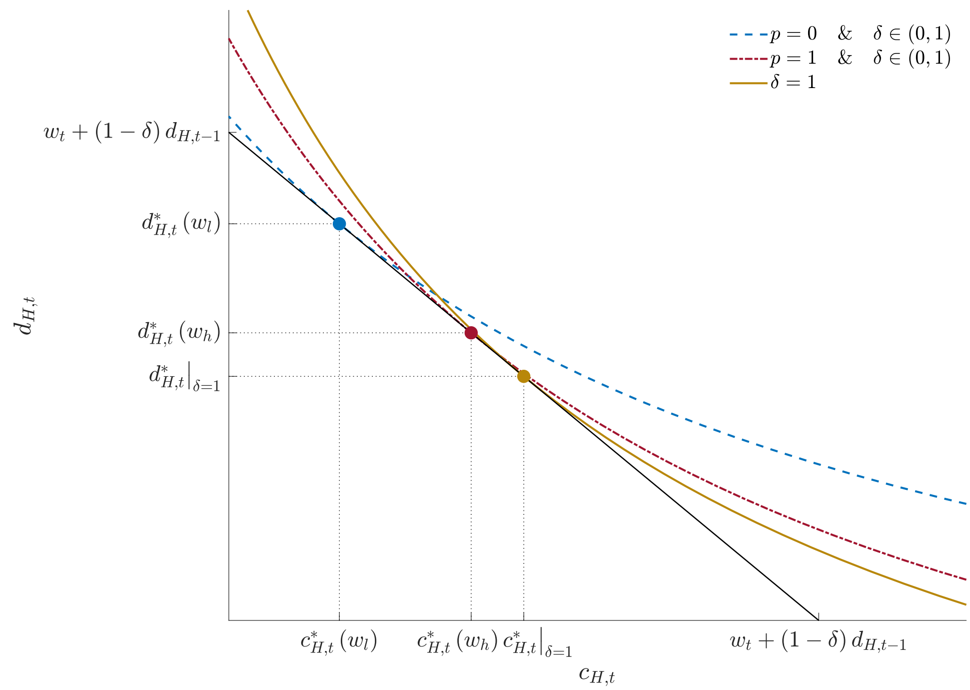

Note: The figure illustrates the optimal consumption choices of the HtM household in period t under three different scenarios. The household can access durable goods in the first two scenarios. One represents the low-income state (blue dashed indifference curve), and the other represents the high-income state (red dash-dotted indifference curve). The third scenario (yellow solid indifference curve) assumes that the durable good fully depreciates each period ($\delta = 1$).

consumption smoothing, their demand for durables is also driven by the future marginal utility of consumption they expect. This way, even without credit access, the HtM agent makes intertemporal choices by reallocating spending across goods, with their beliefs about future scenarios directly influencing current consumption decisions. Since marginal utility is higher in the low-income state than in the high-income state, when the HtM household is certain that they will receive a low-income ($p = 0$), the incentives to accumulate durables are stronger compared to when this household is certain that they will receive a high-income ($p = 1$). Pushed by a wealth effect and constrained by the financial friction, the HtM agent sacrifices current non-durable expenditure to accumulate a larger stock of durables when anticipating a future with a low income level. This way, the HtM household uses durables as a saving vehicle, smoothing consumption despite the lack of bond access.

Figure 1 illustrates the HtM household's optimal consumption choices in period t, considering both possible realizations of the state of the world and under scenarios with and without access to durable consumption. In the figure, the x-axis represents the household's non-durable consumption, while the y-axis represents the stock of durables the household holds at the end of the period. The budget constraint denotes all feasible combinations of non-durable and durable consumption that exhaust the HtM household's total available wealth $w_t + (1-\delta)d_{\mathrm{H},t-1}$. The yellow solid indifference curve represents the scenario where the durable good fully depreciates within the period ($\delta = 1$). In this case, the slope of the indifference curve reflects the standard marginal rate of substitution, capturing the purely intratemporal tradeoff existing when optimizing a consumption bundle with no intertemporal considerations. The HtM agent's optimization problem has no forward-looking component, and the current optimal allocations are independent of future states of the world, rendering the expectations of this household and the probability parameter p irrelevant for their current decisions. Introducing durability flattens the indifference curve, reflecting a valuation of the asset-like properties of the good. As implied by the optimality condition (4), this slope reduction is more pronounced when the durability is higher and when households anticipate higher future marginal utility. Figure 1 reveals a higher optimal consumption level of the durable good when $\delta < 1$. Given the household's desire to smooth consumption over time, this higher durable consumption is even larger when the household anticipates higher future marginal utility. As a result, durable consumption increases when the household is certain that they will be in the low-income state — represented by the blue dashed indifference curve in the figure — compared to when they are certain about being in the high-income state — represented by the red dash-dotted indifference curve. Durability thus embeds a forward-looking component into the HtM behavior, and the curvature of the indifference curve captures how durability and the household's expectations jointly shape their optimal current expenditure allocations.

II.C The Role of Inflation Expectations

This section examines agents' responses to shifts in expectations when households are allowed to engage in intertemporal substitution through durable consumption. We focus exclusively on inflation expectations and their influence on current consumption choices. We first present a set of propositions that illustrate the key mechanisms at play, then derive two predictions of the model that we evaluate against our empirical results. Throughout, we assume that prices in the absence of disturbances to inflation expectations are constant across periods. Using this unperturbed, constant path of prices and the corresponding households' initial optimal allocations as a benchmark, we characterize each agent's optimal adjustments in durable and non-durable consumption when facing a one-time change in inflation expectations. We observe that the responses of the forward-looking HtM household notably differ from those of the non-HtM household and the myopic HtM agent of the canonical TANK model. In the following, we assume that households expect inflation changes one period in advance, that exogenous state variables remain constant during this anticipation period, and, for simplicity, that both households begin with zero holdings of durable goods — and, in the case of the non-HtM household, with zero bond holdings as well.

Lemma 1. Optimal consumption responses to changes in inflation expectations. *Consider a one-time, ceteris paribus, exogenous change in the current expected inflation rate* $\mathbb{E}_t[\pi_{j,t+1}]$ *associated with the nominal price* $P_{j,t}$*, where* $j = \{d, c, w\}$ *and* $\pi_{j,t} = P_{j,t}/P_{j,t-1}$ *denotes the corresponding price-specific gross rate of inflation. Conditional on the initial path of exogenous state variables and the households' corresponding optimal allocations, the optimal responses of the households' current durable acquisition* $(z_{i,t})$ *and non-durable consumption* $(c_{i,t})$ *satisfy*

$$\frac{\partial z_{i,t}}{\partial \mathbb{E}_t[\pi_{j,t+1}]} = \mathcal{D}_i\left(\frac{\partial \mathbb{E}_t[q_{t+1}]}{\partial \mathbb{E}_t[\pi_{j,t+1}]}, \frac{\partial \mathbb{E}_t[w_{t+1}]}{\partial \mathbb{E}_t[\pi_{j,t+1}]}, \frac{\partial \mathbb{E}_t[r_t]}{\partial \mathbb{E}_t[\pi_{j,t+1}]} \cdot \mathbb{1}_{i=N} \,\middle|\, \Theta\right) \tag{5}$$

$$\frac{\partial c_{i,t}}{\partial \mathbb{E}_t[\pi_{j,t+1}]} = \mathcal{C}_i\left(\frac{\partial \mathbb{E}_t[q_{t+1}]}{\partial \mathbb{E}_t[\pi_{j,t+1}]}, \frac{\partial \mathbb{E}_t[w_{t+1}]}{\partial \mathbb{E}_t[\pi_{j,t+1}]}, \frac{\partial \mathbb{E}_t[r_t]}{\partial \mathbb{E}_t[\pi_{j,t+1}]} \cdot \mathbb{1}_{i=N} \,\middle|\, \Theta\right) \tag{6}$$

where $i = \{H, N\}$*,* Θ *collects control and state variables, and* $\mathcal{D}_i(\bullet), \mathcal{C}_i(\bullet) : \mathbb{R}^3 \to \mathbb{R}$.

This lemma establishes that, following a change in inflation expectations, each agent's optimal responses in current consumption hinge on the variation in the anticipated value of future relative prices, with policy rules specific to each household type. Within our partial equilibrium framework, the main distinction across households is the indifference of the HtM household to variations in the expected return of the fully liquid bond: because of the financial friction, this household has no direct sensitivity to the economy's real interest rate. The following proposition illustrates the benchmark case of the myopic HtM household commonly studied in the literature.

Proposition 1. HtM unresponsiveness when the durable good fully depreciates. *When the durable good fully depreciates each period* $(\delta = 1)$*, the optimal response of the HtM household in current durable acquisition* $(z_{H,t})$ *and non-durable consumption* $(c_{H,t})$ *satisfy*

$$\left.\frac{\partial z_{H,t}}{\partial \mathbb{E}_t\left[\pi_{j,t+1}\right]}\right|_{\delta=1} = \left.\frac{\partial c_{H,t}}{\partial \mathbb{E}_t\left[\pi_{j,t+1}\right]}\right|_{\delta=1} = 0 \tag{7}$$

where $j = \{d, c, w\}$. By contrast, the non-HtM household remains responsive to changes in inflation expectations even in the absence of durable goods, adjusting the consumption of every good in the same direction.

The proposition reveals that the presence of a durable good is essential for expectations about future variables to directly influence the HtM household's current expenditure allocation. When the consumption bundle consists only of short-lived goods ($\delta = 1$), the beliefs of this entirely financially constrained agent become irrelevant, as their consumption choices involve no intertemporal trade-off. In contrast, given their ability to engage in intertemporal substitution through financial markets, the expectations of the non-HtM household have a prominent role in their current consumption, even in the absence of durable goods. When the durable good fully depreciates every period, a change in inflation beliefs leads the non-HtM household to adjust the consumption of both available goods in the same direction, with the magnitude of this adjustment depending on the current relative price between these goods and the curvature of the household's utility function. Using a variation in the expected inflation of a specific nominal price — the economy's nominal wage — the following proposition illustrates how durability introduces a forward-looking component in the behavior of the HtM household.

Proposition 2. Heterogeneous consumption responses to changes in expected wage inflation. *Given a one-time, ceteris paribus, exogenous change in the current expected inflation of the economy's nominal wage, $\mathbb{E}_t\left[\pi_{w,t+1}\right]$, the HtM household's optimal responses in current durable acquisition ($z_{H,t}$) and non-durable consumption ($c_{H,t}$) satisfy the following inequalities*

$$\frac{\partial z_{H,t}}{\partial \mathbb{E}_t\left[\pi_{w,t+1}\right]} < 0 \quad \& \quad \frac{\partial c_{H,t}}{\partial \mathbb{E}_t\left[\pi_{w,t+1}\right]} > 0 \tag{8}$$

In contrast, under the no-arbitrage condition, the optimal responses of the non-HtM household in durable acquisition ($z_{N,t}$), non-durable consumption ($c_{N,t}$), and real savings (b_t), satisfy

$$\frac{\partial z_{N,t}}{\partial \mathbb{E}_t\left[\pi_{w,t+1}\right]} > 0 \quad \& \quad \frac{\partial c_{N,t}}{\partial \mathbb{E}_t\left[\pi_{w,t+1}\right]} > 0 \quad \& \quad \frac{\partial b_t}{\partial \mathbb{E}_t\left[\pi_{w,t+1}\right]} < 0 \tag{9}$$

To illustrate the main mechanism at play when a consumption good features durability, Proposition 2 explores the consequences of a change in nominal wage expectations on HtM consumption behavior and contrasts these responses with those of the non-HtM household. When anticipating a higher future nominal wage, the resulting effect induces households to allocate future resources toward current spending and smooth consumption over time. However, a fundamental distinction emerges across agent types. While both households unambiguously raise non-durable expenditure, their reactions regarding durable investment may differ. Because of the financial friction, the HtM agent increases non-durable consumption entirely at the expense of durable expenditure, whereas the non-HtM agent relies on financial markets. In our partial equilibrium setting, the non-HtM household responds by reducing their net saving position to finance higher overall current expenditure: they increase their non-durable consumption while simultaneously accumulating a larger stock of the durable good. The non-HtM household transfers resources over time using the traditional saving channel, reducing future consumption in favor of higher current expenditure. The durable investment response of this household critically depends on the depreciation rate of the durable and the prevailing interest rate in the economy.

On the other hand, the HtM household smooths consumption across time by managing their within-period expenditure allocation. Because of the financial friction, this household engages in intertemporal substitution only by reallocating their fixed resources within the consumption bundle. When anticipating a higher future nominal wage, the resulting effect leads the household to postpone durable acquisition, resembling an agent borrowing from their future self and repaying this implicit debt once the expected higher income materializes. This behavior enables the household to reduce their current marginal utility in non-durable consumption in exchange for a higher future marginal utility, compensating this with the higher wage realization. This mechanism becomes more evident in the opposite scenario: when facing lower nominal wage inflation expectations, the HtM household saves through durable consumption. They accumulate a higher stock of the durable good at the expense of current non-durable consumption, buffering against future scarcity and smoothing consumption across periods. Various factors influence the extent of these responses, including the economy's depreciation rate. A higher depreciation reduces the value of durable accumulation, reducing the advantages of consumption substitution towards the durable good. As shown in Proposition 1, in the limiting case where the durable fully depreciates every period, the HtM household becomes completely unresponsive to changes in expectations.

This finding becomes more relevant in light of recent evidence concluding that the perceived pass-through from aggregate inflation expectation to nominal wage growth expectations is incomplete (Hajdini et al., 2022; Jain et al., 2024; Jiang et al., 2024). This dynamic resembles the one presented in Proposition 2, where a combination of higher inflation expectations on goods prices and rigid expectations on nominal wage inflation leads to lower expectations on real wages. In that case, our findings anticipate a possible different consumption reaction depending on the household type. To clarify our mechanism, we provide a simple example where the economy experiences a decrease in the expected nominal wage inflation, consistent with Proposition 2.

Example 1. Figure 2 illustrates the dynamic consumption responses to a decrease in the expected nominal wage inflation for non-HtM households and forward-looking and myopic HtM households. The top panel shows the evolution of households' non-durable consumption ($c_{i,t}$), the middle panel reports their acquisition of durable goods ($z_{i,t}$), and the bottom panel presents the corresponding durable stock ($d_{i,t}$) over time, with the x-axis of each subplot representing different periods. The figure distinguishes two main periods: during the "Anticipation" period, households adjust their behavior in response to a change in nominal wage inflation expectations, which is anticipated to materialize in the subsequent "Shock" period, as shown in the bottom panel. Since the non-HtM household obeys the permanent income hypothesis and fully smooths consumption over time, when inflation expectations on nominal wage change, they immediately adjust non-durable consumption and durable purchases to the new optimal level implied by the revised expected permanent income. The blue dashed line in Figure 2 illustrates this discrete non-HtM jump at the announcement period, where both $c_{\text{N},t}$ and $d_{\text{N},t}$ shift to lower optimal levels. The unrestricted access to credit markets enables the non-HtM household to reallocate resources freely across time. By contrast, the effect that pushes households to allocate current resources toward future spending leads the forward-looking HtM agent to accelerate durable accumulation, sacrificing current non-durable consumption. The red dashed-dotted line in the figure illustrates the optimal consumption response of this household, which is dictated by their specific policy function. Instead of reducing their stock of durables as the non-HtM household does while increasing their net saving position, the HtM household builds a larger durable stock as insurance against future scarcity. As highlighted in

Figure 2: Heterogeneous consumption responses to changes in expected wage inflation.

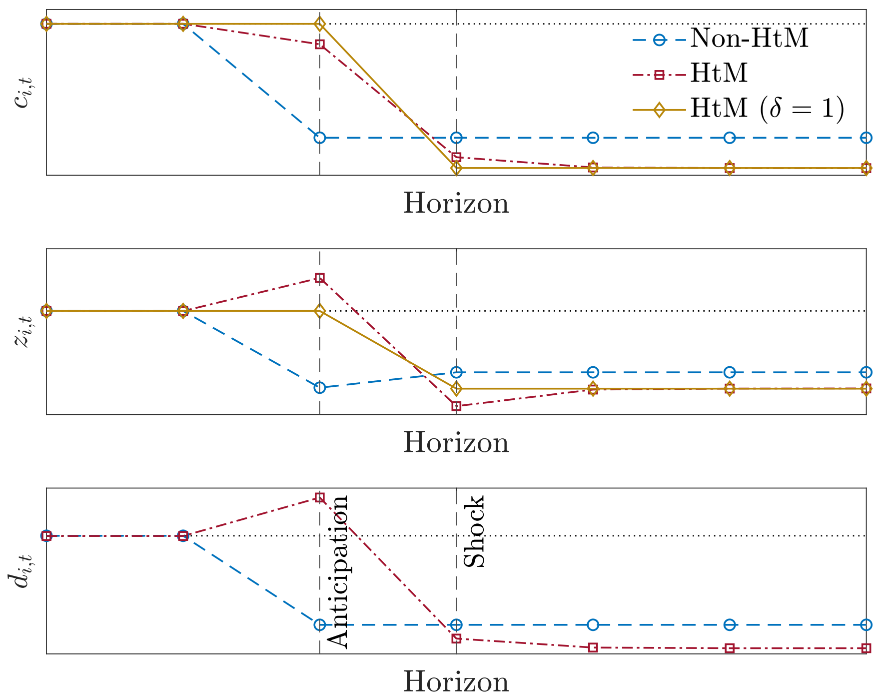

Note: The figure illustrates the optimal consumption choices of both households following a decrease in nominal wage expectations, which is anticipated one period in advance. When the durable good is available, one line represents the response of the non-HtM household (blue dashed line), and another represents the response of the HtM household (red dash-dotted line). The third case (yellow solid line) corresponds to the response of the HtM household in the special case where the durable good fully depreciates each period ($\delta = 1$).

Proposition 1, when durable goods fully depreciate within the period, the forward-looking behavior of the HtM household disappears, and they are unresponsive to expectations shifts. The yellow solid line in the figure represents the response of this myopic HtM agent that splits current income between both consumption goods according to a proportion defined by goods' relative prices and the household's preferences. In this scenario, at the moment of the announcement of the expectation shift, this household remains unresponsive regarding consumption decisions.

Proposition 3. HtM neutral response to uniform inflation expectations variation. *When the expected inflation of all the HtM-relevant nominal prices changes by the same proportion, i.e.,* $d\ln\left(\mathbb{E}_t\left[\pi_{j,t+1}\right]\right) = \bar{\pi}$ *for some constant* $\bar{\pi}$ *and* $j = \{d, c, w\}$*, the aggregated consumption responses of the HtM household satisfy*

$$\sum_{j=\{d,c,w\}} \frac{\partial c_{H,t}}{\partial \mathbb{E}_t\left[\pi_{j,t+1}\right]} \quad = \quad \sum_{j=\{d,c,w\}} \frac{\partial z_{H,t}}{\partial \mathbb{E}_t\left[\pi_{j,t+1}\right]} \quad = \quad 0 \tag{10}$$

In contrast, the aggregated responses of an initially net borrower household, i.e., $b_t \leq 0$, to an equally proportional change in the expected inflation across the same set of nominal prices satisfy

$$\sum_{j=\{d,c,w\}} \frac{\partial c_{N,t}}{\partial \mathbb{E}_t\left[\pi_{j,t+1}\right]} > 0 \quad \& \quad \sum_{j=\{d,c,w\}} \frac{\partial z_{N,t}}{\partial \mathbb{E}_t\left[\pi_{j,t+1}\right]} > 0 \quad \& \quad \sum_{j=\{d,c,w\}} \frac{\partial b_t}{\partial \mathbb{E}_t\left[\pi_{j,t+1}\right]} < 0 \tag{11}$$

This proposition shows that the HtM household remains unaffected by a uniform, proportional shift in the inflation expectations across all of their relevant nominal prices. When these prices are fully flexible and a purely nominal shock leaves expected relative prices unchanged, the HtM household has no incentives to modify their current consumption allocation. Proposition 3 illustrates a form of neutrality, as the shifts merely scale the expected budget restrictions of the forward-looking HtM household. By contrast, under nominal interest rate rigidity, the variation in inflation expectations in the set of flexible nominal prices decreases the real interest rate, reducing the savings return and the real value of future debt repayments. This intertemporal substitution effect leads the non-HtM household to increase current consumption of both durables and non-durables goods, financing this higher overall expenditure with a decline in their net savings position. The exercise highlights how the neutrality result for the HtM household contrasts with the non-HtM household's responsiveness, emphasizing how nominal rate rigidities can lead to heterogeneous consumption responses across financially constrained and unconstrained agents.

II.D Model's Predictions

Our model provides a useful framework for deriving predictions on agents' consumption responses to changes in inflation expectations. Assuming a rigid nominal interest rate and identical price stickiness for both consumption goods, this section summarizes predictions for two main scenarios, which we confront with our experimental evidence in the following section. We first derive households' predicted consumption responses to changes in aggregate inflation expectations, followed by responses to changes in durable-specific inflation expectations. Notice that the predicted behavior of the forward-looking HtM household differs substantially from that of the unresponsive and myopic HtM household of the canonical TANK framework.

Prediction 1. Responses to an increase in aggregate inflation expectations. In response to changes in overall inflation expectations, agents adjust their behavior according to their perceived pass-through from aggregate inflation to each specific nominal price in the economy. Recent evidence documents the existence of an incomplete pass-through from aggregate inflation expectations to nominal wage inflation expectations (Hajdini et al., 2022; Jain et al., 2024; Jiang et al., 2024). When distilling the model's predictions, we follow the literature and assume that higher aggregate inflation expectations lead households to anticipate lower future real wages due to the incomplete adjustment of nominal wages. Since the households' relevant flexible prices are those of both consumption goods and the economy's wage, our partial equilibrium framework offers the following consumption reactions as an answer:

- **The HtM household increases durable spending, while the non-HtM household increases total spending**. Changes in aggregate inflation expectations generate an effect associated with the incomplete pass-through to nominal wages, and an intertemporal substitution effect, arising from the rigidity of the nominal interest rate. In our partial equilibrium model, higher aggregate inflation expectations imply lower expected real wages. As in Proposition 2, and lacking access to traditional saving instruments, the HtM agent reacts to this expected decline in real wages by accumulating a larger durable stock. This behavior enables this household to smooth consumption over time, using durables as a saving vehicle. Additionally, as in Proposition 3, the fall in the economy's real interest rate reduces the cost of bringing future resources into the present. This intertemporal substitution effect leads the non-HtM household to reduce their net saving position and increase their current consumption level. When this last effect outweighs the one arising from the expected lower real wages, the non-HtM household expands both current non-durable and durable spending, at the expense of future spending.

- **The HtM household increases durable spending, while the non-HtM household decreases total spending**. When the wealth effect associated with declining real wages dominates the intertemporal substitution effect triggered by the fall in the economy's real interest rate, the non-HtM household responds by increasing their net saving position, thereby reducing both non-durable and durable spending in the present.

Prediction 2. Responses to an increase in durable inflation expectations. In contrast to Prediction 1, changes in the inflation expectations of one consumption good necessarily modify the expected relative price between consumption goods. Our partial equilibrium framework offers the following consumption reactions as an answer:

- **The HtM household increases durable spending, while the non-HtM household increases total spending**. Higher expected durable inflation triggers an intertemporal substitution effect, leading the HtM agent to expand its current durable stock by reallocating expenditure away from non-durable consumption. This strategy allows the HtM household to hedge against the anticipated price increase. The exact mechanism also operates for the non-HtM household, increasing current durable purchases, while the accompanying wealth effect pushes for a rise in non-durable spending. Together, these forces prompt a reduction in the non-HtM household's net saving position.

- **Both households increase only durable spending**. When the intertemporal substitution effect is particularly strong, both households favor durable acquisition.

Without considering the heterogeneous responses margin generated by the interaction between goods' durability and financial restrictions characteristic of our analysis, some recent literature provides causal evidence suggesting that households tend to increase overall spending following a rise in aggregate inflation expectations (Coibion et al., 2022b, 2024). As a result, we may anticipate that the HtM household will increase durable spending and the non-HtM household will increase total spending when experiencing higher aggregate or durable inflation expectations. Section III.C examines which mechanisms are reflected in the data, while we use our theoretical predictions to interpret those empirical results.

III Experimental Design and Empirical Results

III.A Survey Design and Sample

To empirically analyze consumption responses to changes in inflation expectations and evaluate the model's predictions, we introduced a novel module — which we specifically designed for our experiment — into one of the major U.S. consumer expectations surveys: the Federal Reserve

Bank of Cleveland's Survey of Consumer Expectations. This highly influential survey offers a clean and controlled environment to test our model's predictions. It provides an ideal opportunity to empirically evaluate the model's comparative statics by introducing a shock to survey respondents' inflation expectations through information treatments, and evaluate possible heterogeneous consumption responses across HtM and non-HtM respondents. This survey has a repeated cross-sectional format, and its sample is representative of the U.S. population aged 18 and above across key demographic factors, including age, gender, education, ethnicity, net income, and location.[2] To ensure high-quality responses, we implemented several safeguards against fraudulent or inattentive participation. After considering this set of restrictions and filters, our final sample comprises over 5,500 valid responses collected between November 21, 2023, and February 13, 2024. The remainder of this section provides descriptive details about our sample, offers an overview of the experimental setup, and describes the questions introduced in the new survey module.

The survey begins with a section collecting demographic, socioeconomic, and background information used to identify HtM consumers. It then elicits respondents' short-run inflation expectations — the beliefs we aim to influence — followed by the administration of the information treatments. Immediately afterward, we collect respondents' expectations on medium-run inflation. The survey concludes by measuring post-treatment information, including respondents' intended consumption expenditure across several categories and for two time horizons.

HtM Status. Before analyzing how consumption reacts to changes in inflation expectations across HtM and non-HtM respondents, we must first identify these groups. Identifying HtM consumers, however, is challenging due to the difficulty of collecting detailed information on households' financial positions. Thus, the empirical literature has developed different indirect identification strategies. While early studies mainly relied on measures of households' net worth value relative to labor earnings to identify HtM households (Zeldes, 1989), more recent works have shifted toward measures of households' liquid wealth and their proximity to borrowing limits (Kaplan et al., 2014). This paper adopts a hybrid identification approach using respondents' answers to four questions introduced at the beginning of the questionnaire, alongside standard demographic and socioeconomic items. The first question asks respondents to indicate how strongly they agree or disagree with the statement, "My household lives paycheck-to-paycheck." Responses are recorded on a five-point

[2] See Knotek et al. (2020) and Dietrich et al. (2022) for more details on the survey methodology, and Appendix D for further details and discussions on our new module.

scale ranging from "I strongly agree" to "I strongly disagree", allowing for varying intensities of agreement and covering all possible answers that respondents may provide. This qualitative, ordinal, closed-ended question explores respondents' self-perceived HtM status using a well-known U.S. colloquial expression — "living *paycheck-to-paycheck*." Although lacking a formal definition, the phrase commonly refers to individuals who are likely unable to cover basic living expenses if they suddenly lost their regular source of income. We interpret the answer to this question as a proxy of the strength of respondents' self-identification with being an HtM consumer. The feature of having minimal liquidity immediately before receiving a paycheck closely resembles the financially constrained agent in the theoretical literature, unable to carry assets from one period to the next.

The remaining three questions complement the efforts of the *paycheck-to-paycheck* question in identifying HtM respondents. The second question — also recording responses on a five-point scale — asks respondents how easy or difficult it would be for them to borrow an amount equivalent to one week of their monthly income.[3] This question captures respondents' perceived financial constraints by eliciting their beliefs about credit access using a hypothetical scenario. The third question provides information on actual financial market participation by asking respondents to indicate the approximate share of income they typically save. The five response options include a range of savings shares, as well as the possibility of reporting debt. Finally, the fourth question most closely mirrors the measures commonly used in the recent literature when classifying HtM households. This open-ended quantitative question asks respondents to report the number of weeks their financial assets could cover their usual expenses. Following Kaplan et al. (2014), the question provides a rough indication of liquid financial asset variables by inquiring about checking and savings accounts, stocks, government and corporate bonds, and mutual money market funds. Similar to the usual approach in the literature, this question implicitly evaluates households' liquid asset position relative to an earning or consumption measure.[4]

How can we use this multivariate descriptive information to identify HtM respondents? We address this by applying a Principal Component Analysis (PCA), which summarizes the common information across the four HtM indicators into a single measure that proxies each respondent's degree of HtM status. To ensure comparability across questions, we first grouped the responses to the open-ended final question into quintiles, matching the five-point response scale of the other

[3] The cutoff on the number of weeks referenced in the question is partially based on the taxonomy proposed by Kaplan et al. (2014). One of their definitions classifies households as constrained if their liquid wealth is no greater than one week's worth of their earnings.

[4] For transparency, the exact wording of these four survey questions is provided in Table A.1.

Figure 3: Relationship between the answers to the different HtM questions.

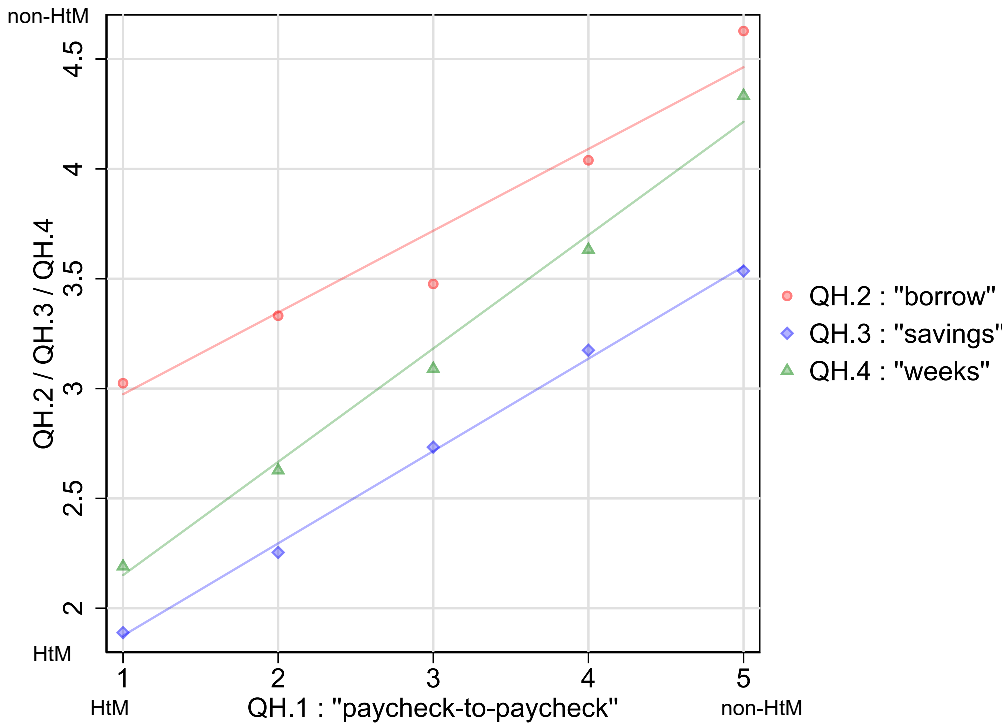

Note: The figure presents binned scatter plots showing pairwise correlations between the different HtM questions and the question labeled as "*paycheck-to-paycheck*" (QH.1), with the latter plotted on the x-axis and the remaining questions on the y-axis. The label "*borrow*" (QH.2) refers to the second question in the list, "*savings*" (QH.3) to the third, and "*weeks*" (QH.4) denotes the final question. Lower values in each question is associated with a higher probability of being classified as an HtM candidate. The number of bins used in each plot is 5.

three questions. Then, for all questions, response options were recoded into numerical values — the first response option was recoded as 1, the second as 2, and so on, up to 5 for the last option — and ordered such that lower values indicate a higher likelihood of being classified as an HtM respondent. Figure 3 reveals that responses to the HtM questions are positively correlated at the individual level, meaning that the probability of being identified as an HtM respondent is relatively consistent across questions. Given this strong correlation and the internal consistency of the five-point response scales, we use a PCA to synthesize the information provided by these questions. Since the first principal component captures a dominant common variation among the four HtM indicators,[5] we use the distribution of this principal component to define HtM status, classifying as

[5] As Figure B.1 illustrates, the first principal component has an eigenvalue well above the Kaiser criterion threshold of one and explains nearly 60% of the total variance in the data. These results indicate that the first principal

HtM those respondents who fall in the bottom 40% of this distribution.[6] Our results remain robust to alternative definitions of HtM status obtained by substantially varying the 40% baseline cutoff. Under our baseline taxonomy, HtM households report a notably lower average household income compared to non-HtM respondents. The median income of this financially constrained group falls within the \$35,000 and \$49,999 bracket, while the median income of the non-HtM group falls in the \$50,000 and \$99,999 range. The education gap between the two groups is relatively narrow, with the median response for both groups falling under the category "Some college, but no degree." While some other average demographic characteristics — such as age and number of children — are also quite similar across groups, respondents in the HtM group are more likely to be single or divorced and to identify as female, whereas respondents in the non-HtM group are more likely to be married and to identify as male.[7]

Inflation Expectations. Once HtM and non-HtM respondents have been identified, we turn to eliciting the beliefs that are fundamental to our analysis and that we aim to influence: respondents' inflation expectations. Following the section inquiring about demographic and socioeconomic information, the survey presents two blocks of questions asking about the expected evolution of inflation, both using the same frequency but differing in their forecast horizon: we ask respondents about their annual inflation expectations 12 months and 5 years ahead. The choice of different horizons is to help us mitigate concerns about survey fatigue and question repetition bias (Haaland et al., 2023). To collect these expectations, we implemented the two-step framework proposed by the Federal Reserve Bank of New York's Survey of Consumer Expectations (NYSCE) to gather inflation expectations data (Armantier et al., 2017). In each block, the survey participants first indicate whether they believe there will be inflation or deflation, then type in the corresponding inflation or deflation rate as a percentage. This method integrates qualitative and quantitative components, reducing potential measurement error associated with purely numeric forecasts.[8]

component captures a dominant and significant share of the common variation among the HtM variables.

[6] Our choice of the baseline cutoff in the first principal component distribution is informed by empirical findings in the literature. Using the Survey of Consumer Finances of the Board of Governors of the Federal Reserve Board (SCF), Kaplan et al. (2014) documents that the fraction of HtM households is around 30% in the U.S. and that this share has been relatively stable over recent decades. Similarly, Kaplan and Violante (2014) uses the same data source and estimates that up to 37% of U.S. households are HtM households. More recently, using the SCF and the Panel Study of Income Dynamics (PSID) of the University of Michigan, Aguiar et al. (2025) concluded that over 40% of U.S. households fall into the HtM category. In a much earlier study, using quarterly aggregate time series data from Data Resources, Inc., Campbell and Mankiw (1989) concluded that between 40% and 50% of the U.S. population consistently spends their whole current income without accumulating assets.

[7] Table A.5 provides more details on respondents' demographic and socioeconomic characteristics by HtM status.

[8] Table A.3 shows the exact wording used in each of these blocks of questions.

To evaluate the reliability of our collected inflation expectations data, we use the University of Michigan's Survey of Consumers (MSC) as a benchmark for comparison.[9] For the 12-month-ahead forecast, the Huber-robust mean in our full sample is 8.2%, with a median substantially lower of 5.0%. In the same period, the MSC reported a mean of 6.8% and a median of 6.7%. Although our sample exhibits somewhat higher expectations, the data align with several well-documented patterns in the literature measuring households' inflation expectations using survey data (Weber et al., 2022). First, medium-run inflation expectations are consistently lower than short-run expectations, reflecting that longer-term beliefs are more firmly anchored than their short-run counterparts. Second, our data exhibits a systematic upward bias relative to the actual inflation rate, which was around 3% during the sample period. This is a widespread pattern observed in prominent expectations surveys and often linked to socioeconomic and demographic elements such as cognitive abilities (D'Acunto et al., 2019) and grocery shopping roles (D'Acunto et al., 2021). Finally, inflation expectations in our sample tend to be higher among HtM candidates than non-HtM candidates. Overall, despite minor differences with the MSC, our collected inflation expectations data behave consistently with well-established survey-based evidence, providing a reliable foundation for further analysis.

Treatments Description. To obtain cleanly identified causal evidence on how inflation expectations shape the consumption decisions of HtM and non-HtM respondents, we conduct an information-provision experiment (Armantier et al., 2016; Cavallo et al., 2017; Coibion et al., 2018). After eliciting respondents' short-run inflation expectations and immediately before collecting their medium-run expectations, we administered information treatments to generate plausible exogenous variation in respondents' medium-run beliefs. To do this, the survey randomly assigned participants to groups, with all but one group receiving a distinct piece of information. The group that received no information serves as the benchmark against which we evaluate possible treatment effects. All information treatments were designed to be simple, concise, and to provide participants with truthful, official, and publicly available information. Therefore, we presume respondents paid

[9] Table A.6 reports summary statistics comparing our results with those of the MSC for the same period. Although the MSC also employs a two-step framework to elicit point forecasts, the wording of its inflation expectations questions differs slightly from ours. They emphasize the evolution of prices in general rather than the inflation rate *per se*. Because respondents' interpretation of inflation can vary substantially depending on the wording and framing of the questions (Armantier et al., 2013), comparisons between the two survey results should be made with caution, even though the forecast horizons match.

attention to the information received and considered it credible.[10]

In total, we assigned participants to four groups. The first group, referred to as Treatment 1 (T1), comprises 1,410 respondents who received quantitative information about aggregate inflation expectations. These participants were informed that — according to professional forecasters — the annual aggregate inflation rate at the end of 2024 was expected to be 2.5%, with our sample data being collected between November 2023 and February 2024. The treatment clarifies that the source of this projection is the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters. We denote the second group as Treatment 2 (T2). This group includes 1,354 respondents who received qualitative and quantitative information about past inflation for a semi-durable good. Specifically, the treatment informed participants that apparel prices remained stable throughout most of the 2010s but increased by 3.1% during the last 12 months preceding the survey implementation. The treatment mentions that the source of the information is the U.S. Bureau of Labor Statistics. We designed this treatment to induce variation in inflation expectations by highlighting a targeted change in the relative price of semi-durable goods. Therefore, we assume that any resulting variation in expectations generated by this treatment is anchored on the semi-durable price signal respondents received and that the prices of these goods mostly drive any resulting variation in inflation expectations associated with this treatment. The remaining two groups complete the experimental design. The third group, comprised of 1,370 respondents, serves as our Placebo group (P). They received quantitative information that was arguably neutral regarding the issue of interest since it is not relevant for the inflation expectation formation process. The treatment stated that the U.S. population grew 2.9% between 2018 and 2022, citing the U.S. Census Bureau as the source of information. The key role of this placebo group is to help us account for potential unconscious numerical anchoring issues, as well as experimenter demand effects (Haaland et al., 2023). The fourth and final group, comprising 1,395 respondents, received no information and serves as our passive or pure Control group (C). Finally, balance tests confirm the validity of the randomization process by showing that respondents' observable characteristics do not systematically predict their treatment assignment.[11]

Consumption Expenditure. After identifying HtM and non-HtM respondents and administering information treatments to exogenously change their inflation expectations, the survey concludes

[10] Table A.4 outlines the size of each treatment group and the specific information treatment they received.
[11] Table A.7 provides detailed results for the balance tests conducted on different pre-treatment covariates.

by collecting data on the main outcome of interest: respondents' intended consumption behavior. Following the treatment implementation section, we included a block of questions collecting information on respondents' intended spending in overall consumption and across four categories of goods and services,[12] designed to collectively cover the entire possible range of household consumption elements. These categories are services and non-durable, semi-durable, and durable goods.[13] To capture responses over different time frames, we considered two horizons for these questions: one month and one year ahead. As in the inflation expectations questions, we adopted the two-step framework used by influential household surveys to elicit information on beliefs (Armantier et al., 2017). Following this methodology, our survey respondents first indicated whether their consumption expenditure would increase or decrease, then specified the corresponding growth rate as a percentage, using last month's spending as the basis. Each question gives examples of the elements comprehended by the corresponding category. In particular, the semi-durable goods question mentions clothing and footwear, jewelry, watches, silverware, toys, tools and garden equipment, and household textiles and utensils as examples of this type of goods.[14] The information collected in this module comprises our main outcome variables. The results showed that the average respondent expects an increase of around 0.7% in the expenditure on durable and semi-durable goods one month ahead, while the expected rise in non-durable consumption is over 2%.

In addition to the post-treatment section eliciting respondents' intended consumption, the survey also includes a pre-treatment module collecting information on respondents' past consumption expenditure. In this section, participants reported how much they spent last month on the four aforementioned consumption categories. The respondents answered by providing dollar amounts for each category, which we then aggregated to construct our overall measure of past consumption expenditure. We use these pre-treatment measures to control for individual consumption patterns and response tendencies, which helps us to improve the estimation precision (Stantcheva, 2022).

[12] Although actual consumption data would be preferable, existing evidence shows that survey respondents' spending plans closely track their subsequent effective spending behavior (Coibion et al., 2022a; Schnorpfeil et al., 2023; Kotsogiannis and Sakellaris, 2025).

[13] We follow the Bank of Canada's Guide to the Income and Expenditure Accounts to define our taxonomy, which uses the international classification of individual consumption by purpose. According to this nomenclature, goods are tangible products that can be stored or inventoried, whereas services must be consumed at the point of delivery. Similarly, the distinction between non-durables and durables relies on whether the goods can be used only once or repeatedly over a period of more than one year. Finally, semi-durable goods, although they also last more than a year, have a significantly shorter lifespan than durable goods.

[14] Table A.2 presents the precise wording used by these questions to collect the consumption information.

III.B Empirical Strategy

Treatment Validation. To evaluate the effectiveness of our treatments, we use the inflation expectations over the next 12 months as the measure of the respondents' inflation beliefs before receiving one of the experiment's information treatments. We denote the respondent j's prior belief as $prior_j$. Similarly, we use the average annual inflation expectations over the next 5 years to measure updated beliefs after treatment exposure, and denote the respondent j's posterior as $post_j$. In our experimental setting, the only difference across groups is whether or not they received an information treatment, and the information contained in it. Following a Bayesian learning framework, when updating expectations, we expect treated respondents to effectively weigh the information signal, sacrificing some of the weight assigned to their prior. Figure 4 illustrates this idea by plotting the relationship between prior and posterior variables, stratified by treatment status. Since the control group answered the posterior question without receiving any external information, this group's relationship serves as our benchmark against which we evaluate and validate our treatments. In comparison, we observe a noticeably flatter relationship between prior and posterior for the first two treatment groups, suggesting that treated respondents effectively weigh the provided signal when updating beliefs. In other words, when answering the posterior question and compared to the respondents in the pure control group, the average respondent in the first two treated groups seems to adjust their expectations closer to the figures provided in the treatments. On the other hand, the group receiving the third treatment of the experiment reveals a slope indistinguishable from that of the control group — an outcome that aligns with our expectations, as these respondents serve as our placebo group. To formally validate these findings, we ran the following regression

$$post_j = \alpha + \eta \times prior_j + \mathbf{X}_j \boldsymbol{\beta}' + \sum_{\tau=\{1,2,\mathrm{P}\}} \left[\gamma_{\tau,1} \times treat_{j,\tau} + \gamma_{\tau,2} \left(treat_{j,\tau} \times prior_j\right)\right] + \upsilon_j \quad (12)$$

where the dichotomic variable $treat_{j,\tau}$ equals one when the respondent j received the treatment $\tau = \{1,2\}$ or the placebo treatment $\tau = P$, and zero otherwise. The omitted category in this regression is the control group. The vector $\mathbf{X}_j \in \mathbb{R}^{1\times n_{\mathrm{x}}}$ denotes a set of respondent-level covariates measured before treatment implementation — such as age, gender, education, household net income, among others — where $n_{\mathrm{x}} \geq 0$ is the number of covariates included in the regression. Since treatment

Figure 4: Relationship between the answers to the prior and posterior questions by treatment status.

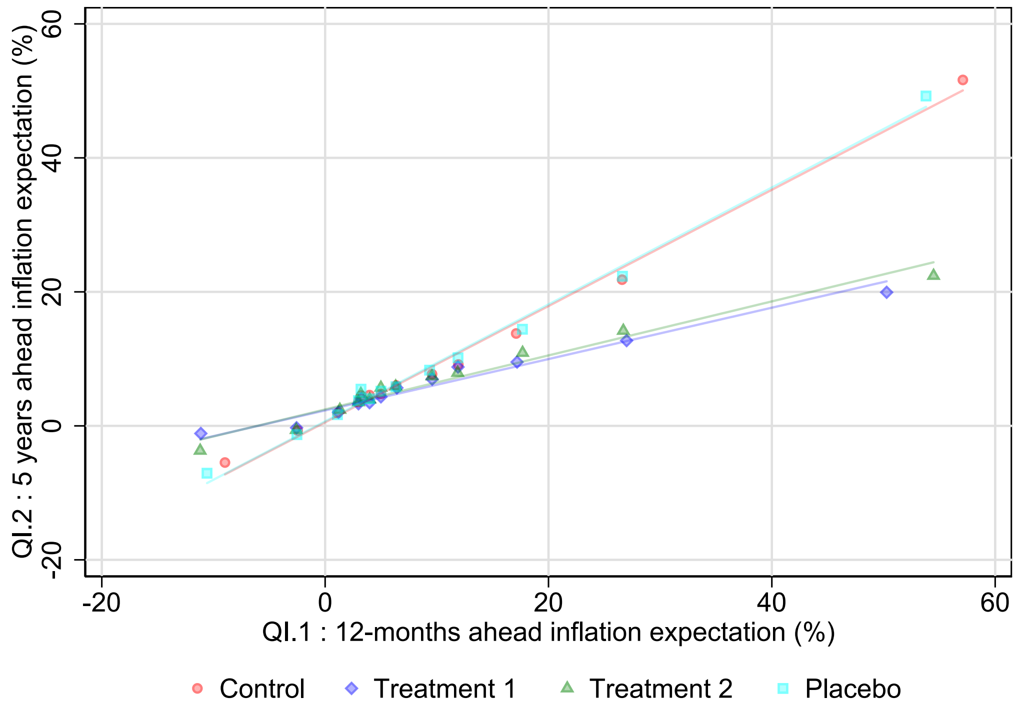

Note: The figure presents binned scatter plots of the answers to the prior (QI.1) and posterior (QI.2) questions by treatment group. The number of bins used in each plot is set to 15. The inflation expectations data have been trimmed at the weekly frequency by excluding the top and bottom 2.5% of the distribution.

was randomly assigned, including $\mathbf{X}_j$ mostly improves the precision of the estimated treatment effects. [15][16] As a result, we obtain that the slope linking the two inflation point forecasts is 0.87 for the passive control group. Since respondents' learning from treatment τ is best captured by the coefficient $\gamma_{\tau,2}$, which reflects the change in the slope of the relationship between prior and posterior for the corresponding treated group τ, the results indicate that both treatments significantly affect posterior beliefs. The estimates for these two coefficients are negative and statistically different from zero — -0.48 for the first treatment and -0.46 for the second treatment — indicating that both treatments significantly impact belief formation. These results reveal that the correlation

[15] Table A.9 presents the results of this regression.

[16] Due to the substantial noise in expectations and intended consumption survey data, we follow the literature (Coibion et al., 2023a,b, 2024; Georgarakos et al., 2024) and employ Huber-robust regressions to systematically control for outliers and minimize the adverse effects of highly influential observations (Huber, 1964).

between the prior and posterior is significantly lower for both treated groups than for the control group, meaning that these respondents reduce their prior relative weight when forming medium-run inflation expectations due to the treatment they received. As their correlations are roughly half those of the control group, the results are economically meaningful, indicating large average belief revisions in response to the treatments. This set of results closely aligns with the empirical literature studying the inflation expectations formation process using information treatment experiments (Cavallo et al., 2017; Coibion et al., 2022a; Hajdini et al., 2022). Additionally, since the coefficients for the third treatment are not statistically significant, we confirm that the experiment's placebo does not influence respondents' posterior beliefs, ruling out any numerical anchoring bias in our experiment. In other words, participants adjust their inflation expectations toward the information provided only when the figure conveyed in the treatment is meaningful for their inflation expectation formation process. Overall, these findings support the idea that the provided treatments constitute a valid source of exogenous variation in inflation expectations, as both treatments led respondents to significantly revise their beliefs toward the provided signal.

As a first approximation to analyze potential heterogeneous effects between HtM and non-HtM respondents, Figure 5 splits the relationship between the answers to the prior and posterior questions by treatment status and HtM condition. The slope of HtM respondents is indistinguishable from that of non-HtM respondents across all treatment groups, suggesting no significant differences in belief updating behavior between these two groups. To formally assess this observation, we introduce the dichotomic variable H_j in the regression specified in equation (12), where the value of this variable is one if we classify respondent j as an HtM individual and zero otherwise, and ran the following regression

$$\begin{aligned} post_j \quad = \quad & \alpha + \eta \times prior_j + \mathbf{X}_j \boldsymbol{\beta}' + \sum_{\tau=\{1,2,\mathrm{P}\}} \left[\gamma_{\tau,1} \times treat_{j,\tau} + \gamma^{\mathrm{H}}_{\tau,1} \left(treat_{j,\tau} \times H_j \right) \right. \\ & \left. + \gamma_{\tau,2} \left(treat_{j,\tau} \times prior_j \right) + \gamma^{\mathrm{H}}_{\tau,2} \left(treat_{j,\tau} \times prior_j \times H_j \right) \right] + \upsilon_j \end{aligned} \tag{13}$$

The last two columns of Table A.9 present the results from this modified regression. As anticipated, no significant differences exist in how respondents incorporate the treatments across HtM groups. The interaction coefficients of the treatment and HtM indicators are small and not statistically different from zero. Interpreting these results as in Weber et al. (2025), we conclude that there

Figure 5: Relationship between the answers to the prior and posterior questions by treatment and HtM status.

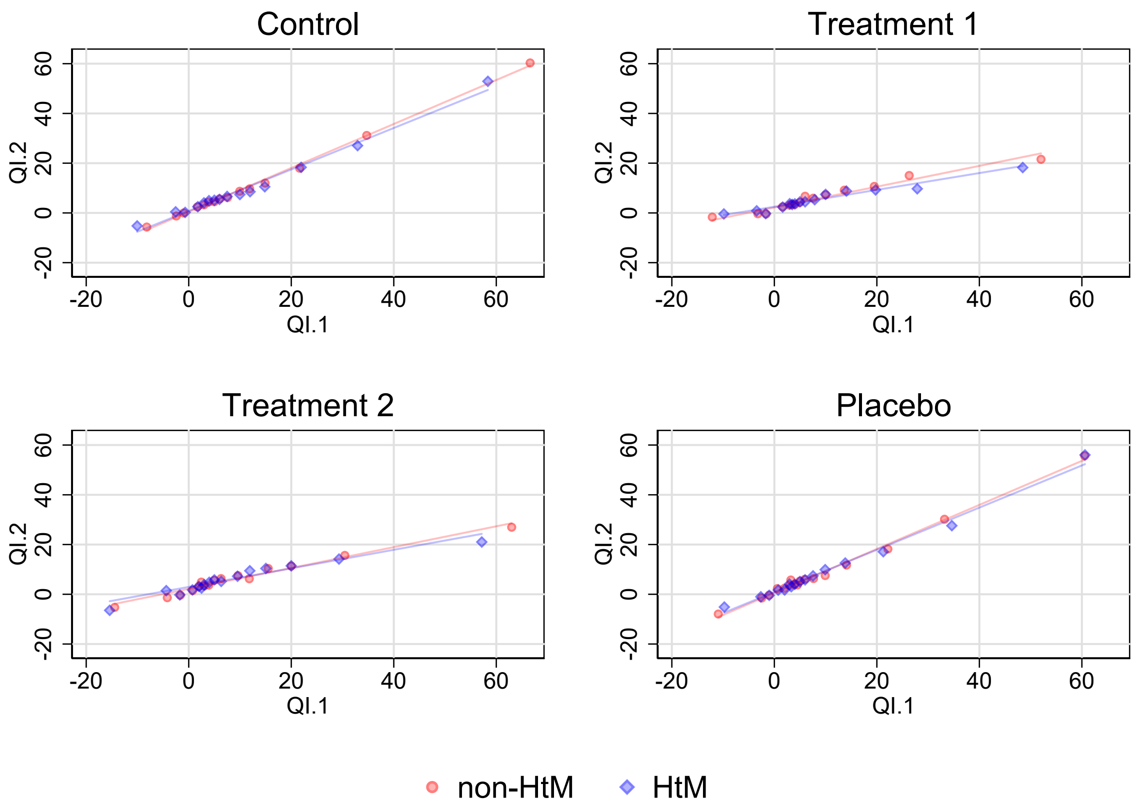

Note: The figure presents binned scatter plots of the answers to the prior (QI.1) and posterior (QI.2) questions by treatment group for HtM and non-HtM respondents. The number of bins used in each plot is set to 15. The inflation expectations data have been trimmed at the weekly frequency by excluding the top and bottom 2.5% of the distribution.

are similar levels of attention and updating reactions across HtM categories, suggesting a similar value and cost of the information for both groups (Mackowiak and Wiederholt, 2024). This lack of heterogeneity in beliefs updating is key to our analysis. Since there are no statistical differences in how treatments affect posterior belief formation between HtM and non-HtM respondents, and since this regression represents the first stage of our identification strategy, any heterogeneous effects observed in the second stage arise from the differential consumption responses to variations in inflation expectations, rather than to differences in treatment assimilation between the two types of respondents. For the remainder of the analysis, we adopt a more conservative approach and define the control group of our experiment as those respondents from either the initial passive control group or those who received the placebo treatment.

Inflation Expectations on Consumption. To quantify the causal effect of inflation expectations on households' intended consumption, we employ an instrumental variable strategy that relies on the exogenous variation in inflation expectations induced by the randomized treatment exposure. Under the exclusion restriction, this design allows us to rule out spurious co-movements between inflation expectations and consumption, enabling us to causally investigate how different consumption categories respond to changes in inflation expectations for both HtM and non-HtM respondents. Given our interest in heterogeneous responses between these two groups, we estimate separate first-stage regressions that instrument posterior beliefs separately for HtM and non-HtM respondents. To do so, we use the previously defined HtM indicator variable, H_j, and adapt the regression specified in equation (12) as follows

$$\begin{aligned} post_j \times \mathrm{H}_j \quad = \quad & \alpha^{\mathrm{H}} + \eta^{\mathrm{H}} \times prior_j + \mathbf{X}_j\,(\boldsymbol{\beta}^{\mathrm{H}})' + \sum_{\tau=\{1,2\}} \left[\gamma^{\mathrm{H}}_{\tau,1}\,(treat_{j,\tau} \times \mathrm{H}_j)\right. \\ & \left. + \gamma^{\mathrm{H}}_{\tau,2}\,(treat_{j,\tau} \times prior_j \times \mathrm{H}_j)\,\right] + \upsilon^{\mathrm{H}}_j \end{aligned} \tag{14}$$

where the dependent variable identifies observations of the posterior variable for the HtM group. In this case, the dichotomic variable $treat_{j,\tau}$ equals one when the respondent j received the treatment $\tau = \{1, 2\}$, and zero otherwise. Similarly, we obtain our second first-stage regression by adapting the same original regression but by defining the dependent variable to capture the treatment effect only among non-HtM respondents

$$\begin{aligned} post_j \times (1 - \mathrm{H}_j) \quad = \quad & \alpha^{\mathrm{N}} + \eta^{\mathrm{N}} \times prior_j + \mathbf{X}_j\,(\boldsymbol{\beta}^{\mathrm{N}})' + \sum_{\tau=\{1,2\}} \left[\gamma^{\mathrm{N}}_{\tau,1}\,(treat_{j,\tau} \times (1 - \mathrm{H}_j))\right. \\ & \left. + \gamma^{\mathrm{N}}_{\tau,2}\,(treat_{j,\tau} \times prior_j \times (1 - \mathrm{H}_j))\,\right] + \upsilon^{\mathrm{N}}_j \end{aligned} \tag{15}$$

where the dependent variable captures the observations of the posterior beliefs for non-HtM respondents. In the second-stage regression, we estimate how exogenous expectation shifts affect intended consumption expenditure. To do so, from the two separate first-stage regressions, we obtain the instrumentalized version of the vector grouping posterior beliefs, $\boldsymbol{\Sigma}_j = [\; post_j \times (1 - \mathrm{H}_j)\;,\; post_j \times \mathrm{H}_j \;]$, which we denoted as $\widehat{\mathbf{\Sigma}}_j$. Using this instrumentalized variable, we estimate the following regression, which corresponds to the reduced-form response of intended consumption expenditure to an exogenous variation in inflation expectations

$$\%\triangle c^e_{j,k,h} \quad = \quad \alpha_{k,h} + \eta_{k,h} \times prior_j + \mathbf{X}_j \boldsymbol{\beta}'_{k,h} + \widehat{\mathbf{\Sigma}}_j \mathbf{\Theta}'_{k,h} + \varepsilon_{j,k,h} \tag{16}$$

where the variable $\%\triangle c^e_{j,k,h}$ denotes our outcome of interest. This variable measures respondent j's intended percentage growth in expenditure on the consumption category k at horizon h. We consider two horizons that, when mapped onto the partial equilibrium model of Section II, correspond to $h = \{t, t+1\}$. Specifically, $h = t$ corresponds to respondents' one-month-ahead (current) consumption intentions, and $h = t + 1$ denotes their one-year-ahead (future) consumption intentions. In consequence, the parameters in the matrix $\mathbf{\Theta}_{k,h} = [\ \theta^{\mathrm{N}}_{k,h}\ ,\ \theta^{\mathrm{H}}_{k,h}\ ]$ capture the causal effect of inflation expectation variations on the corresponding consumption outcome. Specifically, $\theta^{\mathrm{N}}_{k,h}$ identifies the causal effect of shifts in inflation expectations on intended consumption in the consumption category k at horizon h in the case of non-HtM respondents, while $\theta^{\mathrm{H}}_{k,h}$ identifies the corresponding causal effect for HtM respondents. By estimating this regression across different horizons, we can trace the dynamic response of intended consumption to an exogenous variation in inflation expectations, effectively mapping out an impulse response function. The following section presents the empirical results obtained from implementing this identification strategy.

III.C Results

Responses to an Increase in Aggregate Inflation Expectations. Following the identification strategy outlined in the previous section, this section presents the empirical results on the causal effect of inflation expectation variations on households' consumption. The initial set of results focuses on estimates obtained by pooling both implemented treatments — Treatment 1 and Treatment 2 — thereby capturing the combined effect of these two interventions. Figure 6 displays the estimated impact of aggregate inflation expectations on intended consumption growth, disaggregated by consumption category and HtM status. As previously mentioned, we compute these estimates at two different time horizons — current and future consumption reactions — which are referenced along the x-axis of each subplot. The left panel of the figure reports results for non-HtM

Figure 6: Estimated effects of aggregate inflation expectations on consumption by HtM status

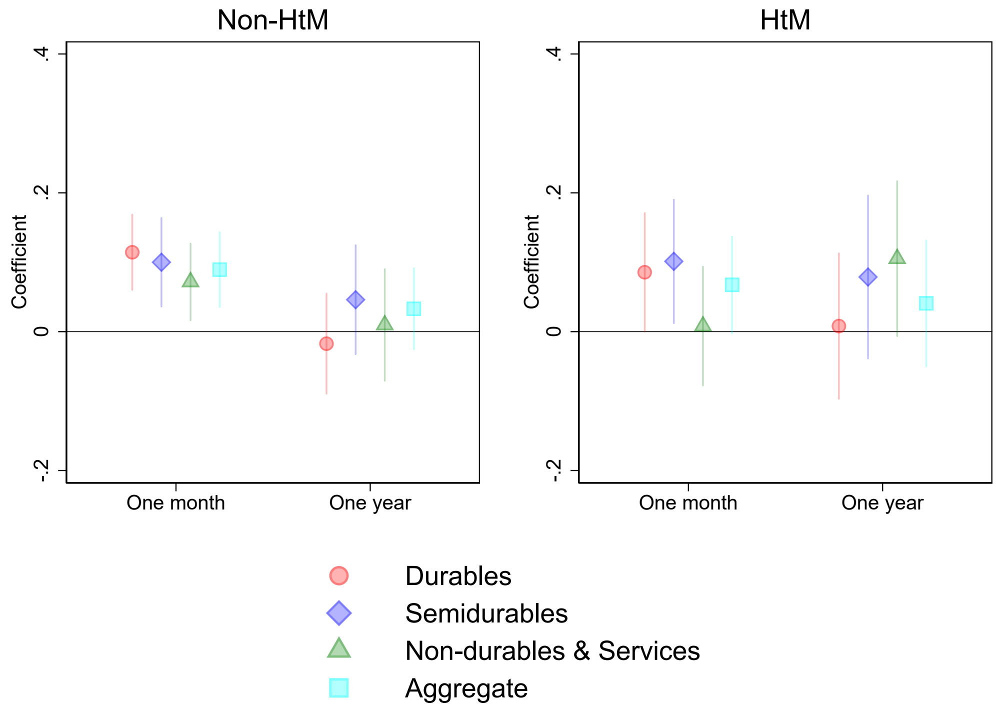

Note: The figure reports the estimated coefficients obtained from regressing intended consumption growth — in the category presented in the legend and at the horizon referenced in the x-axis — on aggregate inflation expectations, pooling observations across both treatments. All regressions include week and U.S. state fixed effects, as well as a set of control variables including age, gender, education level, household income, civil status, number of children, and political affiliation. The inflation expectations and consumption data have been trimmed at the weekly frequency by excluding the top and bottom 2.5% of the distribution. The regression employs robust standard errors, and the figure displays 90% confidence intervals. We present full regression results in Table A.10.

respondents, while the right panel reports our findings for HtM respondents.[17] The vertical lines associated with each coefficient indicate the corresponding 90% confidence intervals.[18]

By examining the first set of results in the left panel of the figure, we find evidence consistent with an intertemporal substitution effect operating through the bond Euler equation. For non-HtM respondents, we reject the null hypothesis that aggregate inflation expectations have no effect on current consumption choices — that is, on intended consumption one month ahead — across all

[17] Using the notation from the previous section, the coefficients plotted in Figure 6 correspond to the individual coefficients $\theta^i_{k,h}$ of the vector $\Theta_{k,h}$ in the equation (16). Coefficients for $i = \text{N}$ are plotted in the right panel, while those for $i = \text{H}$ are plotted in the left panel. The subindex k denotes the consumption category listed in the legend of the figure, and h refers to the time horizon on the x-axis of each subplot.

[18] We report full details of the regression results in Table A.10.

the consumption categories. An exogenous increase of 1 p.p. in non-HtM respondents' aggregate inflation expectations leads to a significant rise of around 0.1 p.p. in their current intended consumption for all categories.[19] When aggregating across categories, an increase of 1 p.p. in these households' aggregate inflation expectations leads to a rise of around 0.09 p.p. in their overall intended consumption one month ahead, equivalent to roughly 1.09 p.p. annually. Importantly, the first stage regressions yield sufficiently high F-statistics, confirming that the instruments employed are sufficiently strong for identification purposes. Turning to the longer horizon results, the same panel reports that for non-HtM respondents, we fail to reject the null hypothesis that aggregate inflation expectations change has no impact on non-HtM intended consumption decisions one year ahead. Across all consumption categories, the estimated coefficient is not statistically different from zero, suggesting that the observed short-run responses do not persist over time.

These non-HtM results are broadly consistent with Prediction 1. According to the model's prediction, higher aggregate inflation expectations generate two competing forces: an effect arising from rigid nominal wages and an intertemporal substitution effect arising from the rigidity of the nominal interest rate. While the first leads to an anticipated reduction in real wages, the second reduces the economy's real interest rate. When the latter effect dominates, the non-HtM agent expands both current non-durable and durable spending, at the expense of future spending by reducing their net saving position. In line with this mechanism, our findings show that higher aggregate inflation expectations significantly increase current consumption among non-HtM consumers across all consumption categories, likely reducing their net saving position, and with these effects dissipating over time. This pattern aligns with the one reported by Georgarakos et al. (2024), who find that for their full sample — comprised of HtM and non-HtM respondents — the positive durable consumption response peaks between two to three months after treatment implementation, vanishing thereafter.

Regarding HtM respondents, the right panel of Figure 6 presents the results for these households.[20] The current responses for durable and semi-durable goods are similar to those observed

[19] Although the differences across consumption categories are not statistically significant, the magnitude of the estimated short-run effects appears to increase with the durability of the good. We find statistically significant coefficients at the 1% level and magnitudes exceeding 0.1 p.p. for durable and semi-durable goods. The implied coefficient magnitude for non-durables and services reduces to around 0.07 p.p., but remains statistically significant at the 5%. This relatively weaker response among non-durable goods and services is consistent with their well-documented lower sensitivity to business cycle fluctuations. This observation is relevant for the economy's aggregate results discussed in the following theoretical section.

[20] A salient feature of this panel's estimates is the relatively large standard errors associated with most coefficients, partly reflecting the substantial noise inherent in survey data on expectations and intended consumption, especially for this group of financially constrained respondents.

for non-HtM respondents: the one-month ahead estimated coefficients are similar in magnitude, and despite having larger standard errors, both coefficients are significant at the 10% level. In contrast, the estimated current responses in non-durable goods and services are substantially smaller and not statistically different from zero. The effect size for these consumption categories is around one order of magnitude smaller than that observed for non-HtM consumers.

Once again, these results are mostly consistent with Prediction 1. Because of the financial friction, and different than the non-HtM agent, the HtM agent is insensitive to changes in the real interest rate. As a result, when changing aggregate inflation expectations, their consumption reactions are driven only by the effect associated with rigid nominal wages, rather than by the intertemporal substitution effect linked to the nominal interest rate rigidity. In consequence, Prediction 1 anticipates that, lacking access to traditional saving instruments, the HtM agent reacts to this real wage effect by reallocating current resources toward the accumulation of durable goods while reducing the expenditure on non-durables. In concordance, the one-month ahead results suggest that, when aggregate inflation expectations increase, a wealth effect pushes HtM candidates to increase their current consumption expenditure only in the durable and semi-durable categories, exhibiting no significant responses in non-durables and services. Several factors could explain the absence of a statistically significant negative response in non-durables and services, as anticipated by Prediction 1. First, unlike the partial equilibrium model, survey respondents may endogenously adjust their labor supply, partially offsetting their income constraints. Second, cognitive frictions may limit respondents' ability to internalize their budget constraints when forming spending plans. Third, respondents may incorporate higher inflation expectations into their nominal expenditure plans, keeping nominal spending on non-durables and services unchanged even if their intended real consumption in these categories declines. Still, despite showing a relatively large coefficient, the estimated effect on overall consumption is not statistically significant for HtM respondents. Once again, the F-statistics reported in the first-stage regressions are sufficiently high to conclude that our information treatments generate enough variation, supporting reliable inference. Moreover, Figure 7 demonstrates that the one-month-ahead estimated coefficients for HtM and non-HtM households are robust to alternative definitions of HtM status, as obtained by varying the 40% cutoff used in the first principal component distribution.

The broad patterns here reported are consistent with the evidence in Jiang et al. (2024), who find that respondents who maintain their overall spending but adjust their consumption bundle are mostly related to channels associated with fixed budget and liquidity constraints, whereas respon-

Figure 7: Stability of the effects of inflation expectations on consumption by HtM status.

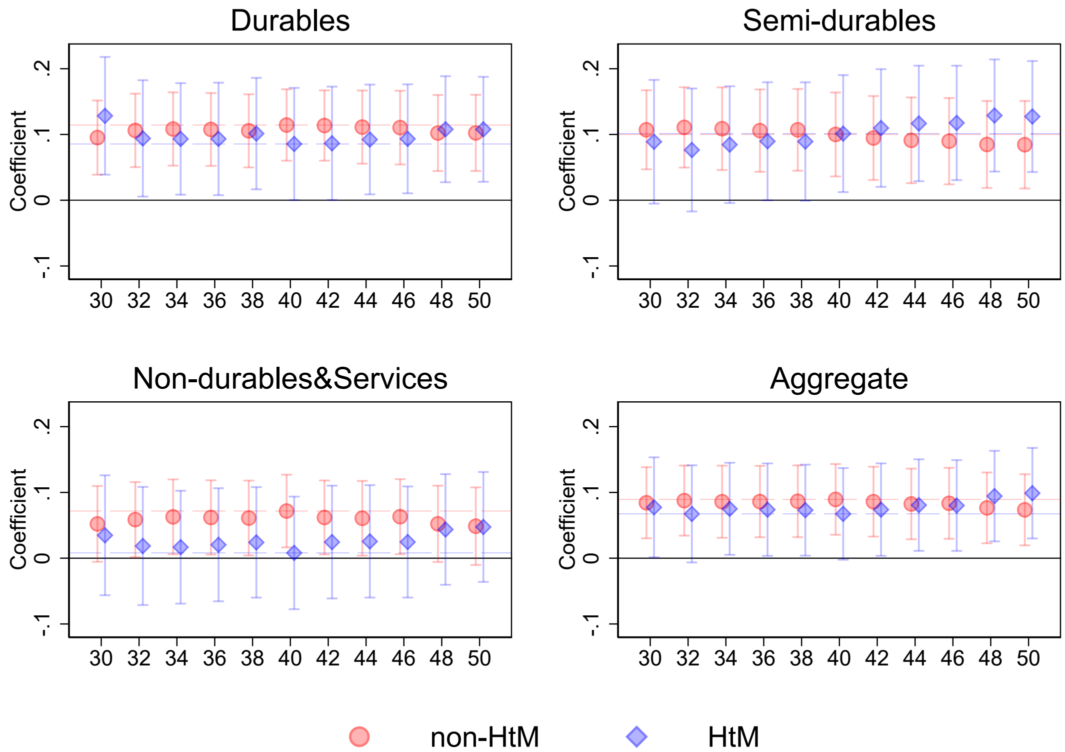

Note: The figure reports the regression coefficients of inflation expectations, where the dependent variable is the one-month ahead intended consumption growth in the corresponding consumption category indicated in each subplot title. To classify HtM households, each regression considers as the cutoff in the first principal component distribution the value presented in the x-axis. The regression includes week and U.S. state fixed effects, as well as a set of control variables including age, gender, education level, household income, civil status, number of children, and political affiliation. The inflation expectations and consumption data have been trimmed at the weekly frequency by excluding the top and bottom 2.5% of the distribution. The regression employs robust standard errors, and the figure displays 90% confidence intervals.

dents who increase their overall current consumption are mostly reacting through an intertemporal substitution channel. Although magnitudes differ substantially, Coibion et al. (2022a) shows that higher inflation expectations significantly increase overall and non-durable spending using either survey or scanner data and their full sample of HtM and non-HtM respondents. However, they find the opposite effect for durable goods at the extensive margin. Although our analysis abstracts from uncertainty, the short-run positive effect of inflation expectations on durable goods acquisition observed for HtM and non-HtM respondents is consistent with Georgarakos et al. (2024), who revise the opposite finding initially reported in Coibion et al. (2023a) once considering inflation

Figure 8: Estimated effects of semi-durable inflation expectations on consumption by HtM status.

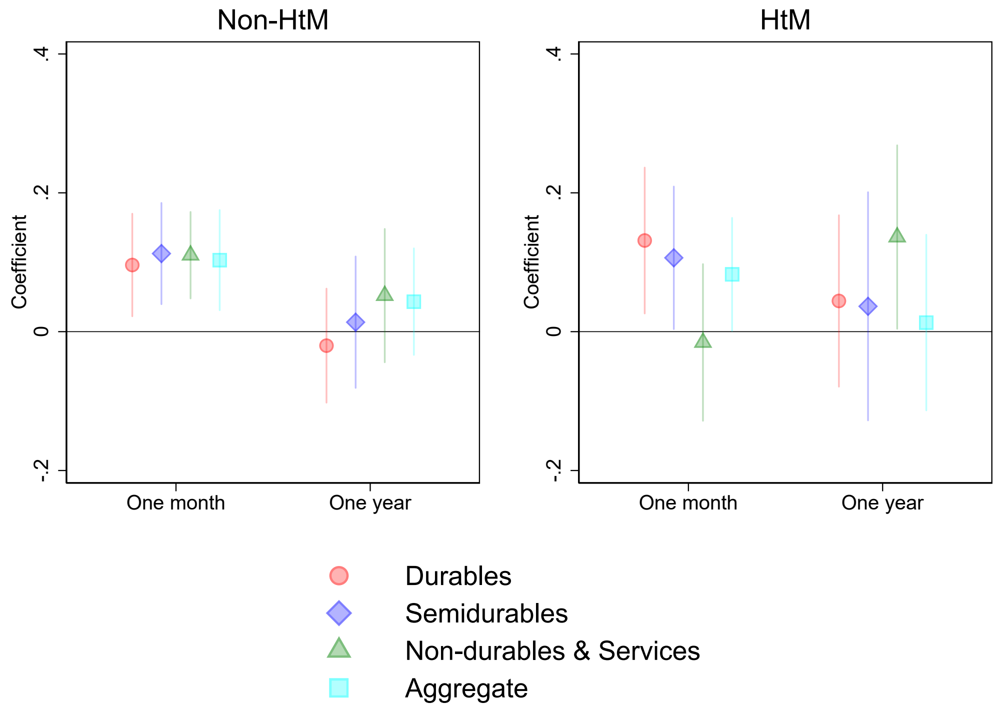

Note: The figure reports the estimated coefficients obtained from regressing intended consumption growth — in the category presented in the legend and at the horizon referenced in the x-axis — on aggregate inflation expectations, using only observations from Treatment 2. All regressions include week and U.S. state fixed effects, as well as a set of control variables including age, gender, education level, household income, civil status, number of children, and political affiliation. The inflation expectations and consumption data have been trimmed at the weekly frequency by excluding the top and bottom 2.5% of the distribution. The regression employs robust standard errors, and the figure displays 90% confidence intervals. We present full regression results in Table A.11.

uncertainty. By explicitly distinguishing between HtM and non-HtM respondents, our experiment uncovers heterogeneous mechanisms: when experiencing higher aggregate inflation expectations, non-HtM respondents react through a classic intertemporal substitution channel, while HtM respondents adjust within a fixed budget by reallocating expenditure toward durable goods.

Responses to an Increase in Semi-durable Inflation Expectations. We designed Treatment 2 to induce variation in inflation expectations through an intended targeted change in the relative price of semi-durable goods. Accordingly, we interpret the induced variation in inflation expectations arising from this specific treatment as being driven mainly through expected changes in the prices of this type of goods. Our second set of results focuses on the estimates obtained using Treatment 2 alone as an instrument. Figure 8 displays the estimated effect of higher semi-durable inflation expectations on respondents' intended consumption.[21] We observe that an exogenous increase of 1 p.p. in HtM households' semi-durable inflation expectations raises their current intended consumption of durables and semi-durables by more than 0.1 p.p., with these effects being statistically significant at the 5% and 10% levels, respectively. By contrast, the effects on current non-durable goods and services are not statistically different from zero. At longer horizons, these dynamics reverse, reflecting the intratemporal substitution effect associated with the anticipated eventual rise in the price of long-lived goods. Notice that the magnitude of the coefficients associated with the durable and semi-durable consumption decisions of HtM households estimated using only Treatment 2 is an amplification of those obtained when evaluating the effects of higher aggregate inflation expectations, pooling both treatments. In the case of non-durables and services, the estimated coefficient remains not significantly different from zero; in some cases, it even turns negative.

For comparison between these empirical results and those anticipated by our partial equilibrium model, we treat semi-durable and durable goods as equivalent categories, given that our model only considers one category of long-lived goods — durable goods. According to Prediction 2, higher durable inflation expectations increase the anticipated relative price of durables, generating an intertemporal substitution effect that leads HtM households to reallocate current expenditure toward durable accumulation while reducing non-durable spending. Consistent with this prediction, we empirically observe that a higher expected semi-durable inflation triggers an intertemporal substitution effect that leads HtM households to expand their current stock of durable and semi-durable goods. As previously discussed, several factors could explain the absence of a statistically significant negative response in non-durables and services, as predicted by Prediction 2.

Overall, the evidence is consistent with the predictions of the partial equilibrium model developed in Section II. The data validates our simple theoretical framework and the consideration that HtM consumers feature a forward-looking behavior, by using durable spending as a consumption smoothing mechanism. Having established this micro-level channel empirically, in the next section, we evaluate the possible aggregate implications of this forward-looking behavior and contrast these results with those of a benchmark economy where the HtM household lacks any saving-by-consuming mechanism, among other exercises.

[21] We report full details of the regression results in Table A.11.

IV General Equilibrium Framework

This section examines the general equilibrium implications of allowing households with heterogeneous financial market access to exhibit forward-looking behavior through durable consumption. We explore how monetary policy transmission operates when HtM households engage in intertemporal substitution by adjusting their durable holdings. To do this, we build on the partial equilibrium framework developed in Section II, and develop an infinite-horizon model with two types of households that consume a bundle composed of non-durable and durable goods. They differ in their access to financial markets. On the supply side, two industries are dedicated exclusively to producing either the non-durable or the durable good. In each sector, a perfectly competitive final-good producer aggregates intermediate varieties elaborated by monopolistically competitive firms that operate under nominal rigidities as in Calvo (1983). To focus on the interaction between household heterogeneity and durable consumption, we abstract from physical capital and relax the zero-net-supply constraint in the market of bonds by introducing an external agent who serves as the counterpart to households' savings.

Overall, introducing durable goods as a channel for intertemporal substitution fundamentally reshapes monetary policy transmission in heterogeneous-agent economies. Although the inclusion of durables amplifies the volatility of households' spending — particularly among financially constrained households that front-load durable acquisitions following an expansionary monetary policy shock — it stabilizes aggregate consumption. More importantly, allowing for durable consumption significantly attenuates the role of households' liquidity constraints in aggregate results. As durability increases— and the rate of depreciation of the durable good decreases — the aggregate dynamics of the heterogeneous-agent economy increasingly resemble those of a single representative-agent framework. The results show that durability introduces an intertemporal margin of consumption smoothing for financially constrained households, mitigating the traditional amplification result of monetary shocks typically attributed to these households.

IV.A Households

Time is discrete and indexed by t. The economy is populated by a continuum of infinitely lived households with identical preferences, with the total population having a constant unit measure over time. We divide these households into two groups distinguished only by heterogeneous access to financial markets. A share $\Lambda^{\text{H}} \in (0,1)$ of the population is exogenously excluded from these markets and is referred to as HtM households, while the remaining share $\Lambda^{\text{N}} = (1 - \Lambda^{\text{H}})$ has financial market participation and is referred to as non-HtM households. Within each group, households are homogeneous and face an identical optimization problem. Hence, we summarize the situation of each group using a single representative household indexed by $i = \{\text{H}, \text{N}\}$. We write $i = \text{H}$ for the representative HtM household, and $i = \text{N}$ for the representative non-HtM household.

The economy features a fully liquid financial instrument consisting of a one-period, non-contingent saving vehicle B_t offering a gross nominal return of R_t. In addition, non-HtM households can acquire equity shares in firms operating in the two production industries of the economy. We denote by $\Omega_{j,t-1}$ the shareholdings of households from firms operating in sector j at the beginning of period t, and by $v_{j,t}$ the corresponding average market value of these shares. When $j = c$, the subscript refers to the shares and market value of firms operating in the non-durable sector, and $j = d$ refers to the durable sector. Additionally, each representative household receives utility flows from consuming a bundle composed of the non-durables acquired during the period ($c_{i,t}$) and the stock of durable goods held at the end of the period.[22] Household i's stock of durable goods at the beginning of period t is denoted as $d_{i,t-1}$, while the stock it holds at the end of the same period is $d_{i,t}$. We assume that households' stock of durables depreciates at a constant rate $\delta \in (0,1)$ and denote household i's acquisitions of new durable goods during the period as $z_{i,t}$. In consequence, equation (1) continues to describe the law of motion of the stock of durables held by each household.

Given this setup, the endogenous state variables are specific to each household type. The only endogenous state variable for the representative HtM household is $d_{\text{H},t-1}$. In contrast, for the representative non-HtM household, the set of endogenous state variables is $\{d_{\text{N},t-1}, \{B_{t-1}, \{\Omega_{j,t}\}_{j=\{d,c\}}\}\}$.

Households experience disutility from labor. Household i supplies total labor $l_{i,t}$ in exchange for the nominal wage $P_{w,t}$. Since production occurs in two industries, the labor allocated by this household to the production sector j is $l_{i,j,t}$. Although we assume perfect labor mobility across industries, we allow for imperfect substitution of the labor supplied by a household across industries: a household's total labor supply satisfies $l_{i,t}^{1-\sigma_l} = l_{i,d,t}^{1-\sigma_l} + l_{i,c,t}^{1-\sigma_l}$, where $\sigma_l \in (0,1)$ denotes the inverse of the elasticity of substitution between sectoral labor inputs. Given this framework, the set of control variables of this household is $\{z_{i,t}, c_{i,t}, \{l_{i,j,t}\}_{j=\{d,c\}}, \{B_t, \{\Omega_{j,t}\}_{j=\{d,c\}}\} \cdot \mathbb{1}_{i=\text{N}}\}$, and

[22] As detailed in Section III.A, our empirical analysis considers four consumption categories designed so that they all combined cover the whole possible range of consumption elements. These categories are services, and non-durable, semi-durable, and durable goods. For tractability, the current theoretical section considers services and non-durables as interchangeable elements and groups semi-durable and durable goods into a single consumption category.

its objective problem is

$$\max_{\left\{z_{i,k},c_{i,k},\{l_{i,j,k}\}_{j=\{d,c\}},\left\{B_k,\{\Omega_{j,k}\}_{j=\{d,c\}}\right\}\cdot \mathbb{1}_{i=\text{N}}\right\}_{k\geq t}} \sum_{k\geq t}\beta^{k-t}\mathbb{E}_t\left[u\left(x_{i,k},l_{i,k}\right)\right] \tag{17}$$

where the parameter $\beta \in (0,1)$ is households' common and exogenous discount factor, and $\mathbb{E}_t[\bullet]$ denotes the full information rational expectations operator conditional on the information available on period t. As in Section II, we denote the consumption bundle as $x_{i,t} = x(c_{i,t}, d_{i,t}) : \mathbb{R}^2_+ \to \mathbb{R}_+$ and define the utility households receive using the function $u_{i,t} = u(x_{i,t}, l_{i,t}) : \mathbb{R}^2_+ \to \mathbb{R}$. Letting $P_{c,t}$ and $P_{d,t}$ denote the nominal market prices of non-durable and durable goods, the budget constraint of household i is

$$\begin{aligned} P_{c,t}c_{i,t} + P_{d,t}z_{i,t} + P_{d,t}\psi\left(d_{i,t};d_{i,t-1}\right) \quad = \quad & P_{w,t}\sum_{j=\{d,c\}} l_{i,j,t} + T_{i,t} + \left(R_{t-1}B_{t-1} - B_t\right)\cdot \mathbb{1}_{i=\text{N}} \\ & + \sum_{j=\{d,c\}}\left(\Omega_{j,t-1}\left(v_{j,t} + P_{j,t}\Pi_{j,t}\right) - \Omega_{j,t}v_{j,t}\right)\cdot \mathbb{1}_{i=\text{N}} \end{aligned} \tag{18}$$

where $P_{j,t}\Pi_{j,t}$ are the net nominal profits of firms operating in industry j, and $T_{i,t}$ stands for the household's received net transfers.[23] The indicator function $\mathbb{1}_{i=\text{N}}$ enforces the constraint that only non-HtM households can access financial markets. The term $\psi_{i,t} = \psi(d_{i,t}; d_{i,t-1}) : \mathbb{R}^2_+ \to \mathbb{R}_+$ is a twice continuously differentiable, strictly convex function capturing households' adjustment costs when changing the stock of durables they hold. It satisfies $\psi_{i,t}|_{d_{i,t}=d_{i,t-1}} = 0$, and we express this cost in terms of units of the durable good. Since non-HtM households have access to the risk-free saving instrument, their optimization problem considers the following transversality condition

$$\lim_{k\to\infty}\prod_{h=0}^{k}\left(R_{t+h-1}\right)^{-1}B_{t+k} \quad = \quad 0 \tag{19}$$

The price of the non-durable good serves as the numeraire. We denote the relative prices between

[23] As discussed in the policy setup, fiscal policy ensures that transfers, net of profit outcomes, remain constant across households. Combined with the zero-net-savings steady state condition, this implies a common steady-state budget constraint across all households in the economy.

the durable and non-durable good as $q_t = P_{d,t}/P_{c,t}$, the economy's real wage as $w_t = P_{w,t}/P_{c,t}$, and represent the gross inflation rate of each of the good prices as $\pi_{j,t} = P_{j,t}/P_{j,t-1}$, where $j = \{d, c\}$.

IV.B Firms

The economy's supply side follows a standard New Keynesian structure with two production industries dedicated exclusively to producing either the non-durable or the durable good.[24] Production in each industry takes place in two stages. At the first stage, there is a continuum of monopolistically competitive firms producing differentiated intermediate varieties. These firms use labor as the only input and operate using a linear technology subject to sector-specific productivity shocks, $\upsilon_{j,t}$, where the log of sectoral productivity evolves according to an exogenous first-order autoregressive process. We denote the aggregate labor demand in sector j as $n_{j,t}$. At the second stage, in each sector, a perfectly competitive final good producer aggregates the differentiated intermediate varieties to produce a homogeneous final good. For the industry j, we denote the total production of this final good as $y_{j,t}$.

Following Rotemberg and Woodford (1997), we assume that the government implements the standard New Keynesian optimal subsidy to eliminate markup distortions in the fully flexible price limit of the economy. More specifically, for each firm producing intermediate variety k, the government provides the production subsidy $\tau_k = (\varepsilon_j - 1)^{-1}$ where $\varepsilon_j > 1$ is the sector-specific elasticity of substitution across intermediate input varieties. The total sectoral cost of this subsidy, $\tau_{j,t}$, is financed using a sectoral profit tax $\nu_{j,t}$, which drives net profits to zero in the zero-inflation steady-state. As emphasized by Bilbiie (2020), this setup ensures an efficient, flexible price limit for the economy where the corresponding steady state is unaffected by profit distribution. Under the condition of a zero-net-savings steady state, this policy ensures uniform consumption allocations across households, such that $(\bar{c}_i, \bar{d}_i) = (\bar{c}, \bar{d})$ for every household i in the economy. Finally, we incorporate nominal rigidities *à la* Calvo (Calvo, 1983): in each period, a firm producing any intermediate variety in sector j has a probability $(1 - \theta_j)$ of being able to reset its price.

IV.C Monetary Policy

We assume that a single monetary authority sets the nominal interest rate associated with the risk-free saving instrument. The monetary policy follows a standard Taylor-rule framework, where

[24] See Appendix E.1 for further details on the supply-side block of the economy.

the monetary instrument adjusts in response to the economy's different sectoral inflation, while also being subject to a non-systematic disturbance element (Taylor, 1993). We assume that the log of the policy rate evolves according to a persistence coefficient of $\rho_R \in [0,1]$. Therefore, the following structure describes the instrument's rule

$$R_t \quad = \quad \bar{R}^{1-\rho_R} R_{t-1}^{\rho_R} e^{\iota_{m,t}} \prod_{j=\{d,c\}} (\pi_{j,t})^{\phi_{\pi,j}} \tag{20}$$

where the sector-specific reaction coefficient satisfies $\phi_j > 1$ with $j = \{d, c\}$ to avoid equilibrium indeterminacy. The variable $\iota_{m,t}$ captures a standard monetary policy shock, which evolves according to an exogenous first-order autoregressive process. The persistence coefficient of this shock is $\rho_m \in [0,1]$, and $\sigma_m > 0$ denotes the standard deviation of the exogenous innovations affecting the shock, which are represented using $\upsilon_{m,t} \overset{\text{iid}}{\sim} \mathcal{N}(0,1)$. Therefore, the law of motion of monetary policy shocks is

$$\iota_{m,t} \quad = \quad \rho_m \iota_{m,t-1} + \sigma_m \upsilon_{m,t} \tag{21}$$

Finally, the authorities manage transfers so that net transfers — after accounting for realized firms' profits — are equalized across households.

IV.D The Third Man

To relax the zero-net-supply constraint in the market of bonds and enable the non-HtM household to effectively save, we introduce a third agent who serves as the counterpart to households' savings. Since transfers are managed so that net transfers — after accounting for firms' profits — are equalized across households, the presence of this counterpart to non-HtM households' asset position is essential to obtain heterogeneous responses between HtM and non-HtM households. The two household groups would otherwise behave identically without this third agent, given that profits net of transfer are distributed uniformly. While there are alternative approaches to relax the zero-net-supply condition, we adopt this framework for its tractability and because it preserves the mechanisms of interest. We assume that, in each period, this agent receives a constant and exogenous real endowment ω_c of the non-durable good and derives utility exclusively from consuming this good. Denoting this agent's consumption as $c_{x,t}$, their preferences are represented by

$$u_x\left(c_{x,t}\right) = \frac{\left(c_{x,t}\right)^{1-\zeta_x}}{\left(1-\zeta_x\right)} \tag{22}$$

where $\zeta_x > 0$ denotes the inverse of the intertemporal elasticity of substitution. This agent has no labor market participation, but has access to the risk-free bond, which in equilibrium balances the non-HtM households' saving or borrowing position. We denote the amount this agent acquires of the one-period, non-contingent bond as $B_{x,t}$. Moreover, to address potential saving imbalances and induce stationarity, we assume that the return on borrowing for this agent includes a debt-elastic component (Schmitt-Grohé and Uribe, 2003). More precisely, the perceived interest rate of this agent satisfies

$$R_{x,t} = R_t e^{-\phi_B B_{x,t}} \tag{23}$$

where $\phi_B > 0$ governs the responsiveness of the return to the agent's level of debt.

IV.E Market Clearing

In this economy, several distinct markets operate: the riskless asset market, the equity market for the shares of intermediate firms, the sectoral labor markets, and the markets for both durable and non-durable goods. Given the homogeneity assumption among the non-HtM agents trading the riskless saving instrument — and the inclusion of the endowment-owning counterpart — the bond market clearing condition is

$$\Lambda^{\mathrm{N}} B_t + B_{x,t} = 0 \tag{24}$$

Similarly, equity market clearing implies that the non-HtM holdings of outstanding shares of firms satisfy $\Omega_{j,t} = (\Lambda^{\mathrm{N}})^{-1}$, either for firms operating in the durable or non-durable production industries. Moreover, labor market clearing requires that the total labor supply equals the total labor demand within each industry of the economy. Formally, this last condition corresponds to

$$\sum_{i=\{\mathrm{H,N}\}} \Lambda^i l_{i,j,t} \quad = \quad n_{j,t} \tag{25}$$

where $j = \{d, c\}$. Finally, both goods markets must also clear. For the non-durable and durable goods markets, respectively, the conditions are

$$y_{c,t} \quad = \quad \sum_{i=\{\mathrm{H,N}\}} \Lambda^i c_{i,t} + (c_{x,t} - \omega_c) \tag{26}$$

$$y_{d,t} \quad = \quad \sum_{i=\{\mathrm{H,N}\}} \Lambda^i \left[z_{i,t} + \psi\left(d_{i,t}; d_{i,t-1}\right) \right] \tag{27}$$

IV.F Competitive Equilibrium

Given a distribution of households in the economy $\{\Lambda^i\}_{i=\{\mathrm{H,N}\}}$, a stochastic process for the monetary policy innovation $\upsilon_{m,t}$ and sector-specific technology shocks $\{\upsilon_{j,t}\}_{j=\{d,c\}}$, and an initial value of the endogenous state variables $\{\{d_{i,t-1}, \{B_{t-1}, \{\Omega_{j,t}\}_{j=\{d,c\}}\} \cdot \mathbb{1}_{i=\mathrm{N}}\}_{i=\{\mathrm{H,N}\}}, B_{x,t-1}\}$, a competitive equilibrium in this economy is a stochastic sequence of households and firms decision variables $\{\{c_{i,t}, z_{i,t}, d_{i,t}, \{l_{i,j,t}\}_{j=\{d,c\}}\}_{i=\{\mathrm{H,N}\}}, B_t, c_{x,t}, B_{x,t}, \{\Omega_{j,t}, n_{j,t}\}_{j=\{d,c\}}\}$, of goods and assets prices $\{P_{c,t}, P_{d,t}, P_{w,t}, R_t, \{\upsilon_{j,t}\}_{j=\{d,c\}}\}$, and fiscal policy instruments $\{\{T_{i,t}\}_{i=\{\mathrm{H,N}\}}, \{\tau_{j,t}, \nu_{j,t}\}_{j=\{d,c\}}\}$, such that, at every period t: households and firms solve their respective optimization problem considering their different restrictions and taking as given the different prices and transfers; the monetary policy rule is satisfied; and all markets clear.

IV.G Solution

To characterize the equilibrium behavior of the economy, we first derive the first-order conditions of the demand side. We conclude that the optimality conditions for any household satisfy

$$w_t = -\left(\frac{mu^l_{i,t}}{mu^c_{i,t}}\right)\frac{\partial l_{i,t}}{\partial l_{i,j,t}} \tag{28}$$

$$mu^d_{i,t} = q_t mu^c_{i,t}\left(1+\frac{\partial \psi_{i,t}}{\partial d_{i,t}}\right) - \beta \mathbb{E}_t\left[q_{t+1} mu^c_{i,t+1}\left((1-\delta)\left(1+\frac{\partial \psi_{i,t+1}}{\partial d_{i,t+1}}\right) - \frac{\partial \psi_{i,t+1}}{\partial d_{i,t}}\right)\right] \tag{29}$$

where the term $mu^{\vartheta}_{i,t} = \partial u_{i,t}/\partial \vartheta_{i,t}$ denotes the marginal utility or disutility of the consumption or labor variable ϑ. Equation (28) characterizes the households' optimal labor supply condition, while equation (29) is the Euler equation for durable consumption. The latter relates the shadow value of additional durable consumption, expressed as the intratemporal optimality condition between the non-durable and durable good, and the discounted flow of utility from durable consumption. Since our partial equilibrium model developed in Section II assumes an exogenous labor supply, it lacks a condition analogous to equation (28). In contrast, equation (29) captures the possibility of forward-looking behavior among HtM households and is the counterpart of equation (4). In the special case when durables fully depreciate each period ($\delta = 1$), they no longer allow intertemporal wealth transfer. In consequence, the HtM household's optimization problem becomes purely static, and the durable consumption Euler equation collapses to a purely intratemporal tradeoff between both consumption goods, i.e., $mu^d_{i,t} = q_t mu^c_{i,t}(1 + \partial\psi_{i,t}/\partial d_{i,t})$.

Additionally, non-HtM households are the only agents with access to the risk-free financial instrument, and thus, they are the only ones that satisfy the standard bond's Euler equation. Moreover, they also have exclusive access to the firms' equity shares. As a result, for $i = \text{N}$, the following additional optimality conditions also apply for them

$$1 = \beta R_{t+1}\mathbb{E}_t\left[mu^c_{\text{N},t+1}\left(mu^c_{\text{N},t}\right)^{-1}(\pi_{c,t+1})^{-1}\right] \tag{30}$$

$$v_{j,t} = \beta \mathbb{E}_t\left[mu^c_{\text{N},t+1}\left(mu^c_{\text{N},t}\right)^{-1}(v_{j,t+1} + P_{j,t+1}\Pi_{j,t+1})(\pi_{c,t+1})^{-1}\right] \tag{31}$$

where $j = \{d, c\}$.[25] We allow for intratemporal substitution between durable and non-durable

[25] Conditional on the corresponding transversality condition, equation (31) represents the Euler equation for firms' equity and defines the fundamental firms' pricing equation $v_{j,t} = \sum_{k>t}\mathbb{E}_t\left[\Lambda_{k|t}P_{j,k}\Pi_{j,k}\right]$, where $\Lambda_{k|t}$ is the firms' relevant stochastic discount factor. Given our framework setup, this term satisfies $\Lambda_{k|t} = \beta^{k-t}(\lambda_{\text{N},k}/\lambda_{\text{N},t})$, where $\lambda_{i,t}$ denotes the shadow value of non-durable consumption for agent i in period t.

goods by adopting King-Plosser-Rebelo preferences (King et al., 1988). We assume the following regarding the structure of the consumption bundle ($x_{i,t}$) and the utility function ($u_{i,t}$)

$$x\left(c_{i,t}, d_{i,t}\right) = \left[\varrho^{\sigma}\left(c_{i,t}\right)^{1-\sigma} + (1-\varrho)^{\sigma}\left(d_{i,t}\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}} \tag{32}$$

$$u\left(x_{i,t}, l_{i,t}\right) = \frac{\left(x_{i,t}\right)^{1-\zeta}}{(1-\zeta)} - \chi\frac{\left(l_{i,t}\right)^{1+\eta}}{(1+\eta)} \tag{33}$$

where $\sigma \in (0,1)$ represents the inverse of the elasticity of substitution between goods, $\zeta > 0$ corresponds to the inverse of the intertemporal elasticity of substitution, and $\eta > 0$ is the inverse of the Frisch labor supply elasticity. The preference parameter $\varrho \in (0,1)$ and the utility cost parameter $\chi > 0$ determine households' consumption and labor relative preferences, shaping the steady-state values of consumption shares and labor supply. Regarding the adjustment cost function $\psi_{i,t}$, we restrict our attention to a quadratic cost specification such as

$$\psi(d_{i,t}; d_{i,t-1}) = \frac{\kappa}{2}\left(\frac{d_{i,t}}{d_{i,t-1}} - 1\right)^2 d_{i,t} \tag{34}$$

where the parameter $\kappa \geq 0$ governs the degree of convexity of this function. We use $\bar{\vartheta}$ to denote the steady-state value of the generic variable ϑ, and $\hat{\vartheta}$ to represent the log-deviation of this same variable from its corresponding steady-state value. Therefore, the set of log-linear equations that describes the optimality conditions of the demand side of the economy corresponds to

$$\hat{w}_t = \left(\eta + \sigma_l\right)\hat{l}_{i,t} - \sigma_l\hat{l}_{i,j,t} + \sigma\hat{c}_{i,t} - \left(\sigma - \zeta\right)\hat{x}_{i,t} \tag{35}$$

$$\begin{aligned}\sigma\left[1 - \beta\left(1-\delta\right)\right]\left(\hat{c}_{i,t} - \hat{d}_{i,t}\right) = {} & \hat{q}_t + \kappa\Big(\triangle\hat{d}_{i,t} - \beta\left(1 + (1-\delta)\right)\mathbb{E}_t\left[\triangle\hat{d}_{i,t+1}\right]\Big) \\ & -\beta\left(1-\delta\right)\Big(\mathbb{E}_t\left[\hat{q}_{t+1}\right] - \sigma\mathbb{E}_t\left[\triangle\hat{c}_{i,t+1}\right] + \left(\sigma-\zeta\right)\mathbb{E}_t\left[\triangle\hat{x}_{i,t+1}\right]\Big)\end{aligned} \tag{36}$$

$$\sigma\hat{c}_{\mathrm{N},t} = \sigma\mathbb{E}_t\left[\hat{c}_{\mathrm{N},t+1}\right] - \Big(\hat{R}_{t+1} - \mathbb{E}_t\left[\hat{\pi}_{c,t+1}\right]\Big) - \left(\sigma-\zeta\right)\mathbb{E}_t\left[\triangle\hat{x}_{\mathrm{N},t+1}\right] \tag{37}$$

where $i = \{\mathrm{H}, \mathrm{N}\}$, $j = \{d, c\}$, and $\triangle\vartheta_t$ denotes the time variation between periods t and $t-1$ of the generic variable ϑ. Regarding the supply side in the economy, we emphasize that a first-order log-linear approximation of the standard sectoral price equation around a zero-inflation steady state yields the following approximation to the sector-specific New Keynesian Phillips Curve

$$\hat{\pi}_{j,t} \quad = \quad \beta \mathbb{E}_t \left[\hat{\pi}_{j,t+1} \right] + \mu \widehat{mc}_{j,t} \tag{38}$$

where $j = \{d, c\}$ and $\mu = (1 - \theta_j)(1 - \beta\theta_j)/\theta_j$ measures the sensitivity of sectoral inflation $\pi_{j,t}$ to the evolution of firms' real marginal costs $mc_{j,t}$. Finally, the optimality condition of the third agent in the model — which serves as the counterpart to the saving and borrowing decisions of non-HtM households — corresponds to their bond Euler equation. This condition, when log-linearized, can be written as

$$\sigma_x \hat{c}_{x,t} \quad = \quad \sigma_x \mathbb{E}_t \left[\hat{c}_{x,t+1} \right] - \left(\hat{R}_{x,t} - \mathbb{E}_t \left[\hat{\pi}_{c,t+1} \right] \right) \tag{39}$$

where $\hat{R}_{x,t} = \hat{R}_t - \phi_B B_{x,t}$ denotes the log-deviation of the deb-elastic nominal interest rate.

IV.H Calibration

Table 1 summarizes the set of parameters considered in our external baseline calibration, which adopts standard annual-frequency values commonly employed in the macroeconomic literature. The households' discount factor is set to $\beta = 0.99$, implying an annual subjective discount rate of approximately 1% and yielding a steady-state gross nominal interest rate of $\bar{R} = 1.01$. In the baseline calibration, we assume costless adjustment in households' durable stock ($\kappa = 0$), and set both the intratemporal and intertemporal elasticity of substitution to unity, i.e., $\sigma^{-1} = \zeta^{-1} = 1$ (Beraja and Wolf, 2021). These choices imply a Cobb-Douglass aggregator structure for the consumption bundle $x_{i,t}$ and a logarithmic utility function $u_{i,t}$. The annual depreciation rate of durable goods is fixed to 5%, while the steady-state share of non-durable expenditure is calibrated to 70%, i.e. $\delta = 0.05$ and $\phi = 0.70$ (Barsky et al., 2016), implying households' relative preference for non-durable consumption of $\varrho = 0.63$. The Frisch elasticity of labor supply is $\eta^{-1} = 2$, and the disutility of labor is parametrized by $\chi = 12$ (Bachmann et al., 2024). Finally, in our baseline calibration, the share of HtM households in the population is set to 40%, in line with recent

Table 1: Parameter calibration for the annual-frequency general equilibrium model

Parameter	Value	Source	Description
		Households	
β	0.99	Beraja and Wolf (2021)	Households' discount factor
κ	0.00	Beraja and Wolf (2021)	Governs the convexity of durables' adjustment costs
σ	1.00	Beraja and Wolf (2021)	Inverse of the intratemporal elasticity of substitution
ζ	1.00	Beraja and Wolf (2021)	Inverse of the intertemporal elasticity of substitution
δ	0.05	Barsky et al. (2016)	Rate of depreciation of durables
ϕ	0.70	Barsky et al. (2016)	Steady-state share of non-durable consumption
η	0.50	Bachmann et al. (2024)	Inverse of the Frisch labor supply elasticity
χ	12.0	Bachmann et al. (2024)	Households' relative labor utility cost
Λ^{H}	0.40	Aguiar et al. (2025)	Share of HtM households in the population
		Firms	
θ_j	0.75	Bachmann et al. (2024)	Calvo's price stickiness in each industry
		The Third Man	
σ_x		σ	Inverse of the intratemporal elasticity of substitution
ζ_x		ζ	Inverse of the intertemporal elasticity of substitution
ϕ_B	0.10	Schmitt-Grohé and Uribe (2003)	Governs the responsiveness to the level of debt
		Monetary Policy	
$\phi_{\pi,j}$	1.50	Auclert et al. (2020)	Policy reaction to sectoral inflation
ρ_R	0.89	Auclert et al. (2020)	Persistence of nominal interest rate
ρ_m	0.83	Auclert et al. (2020)	Persistence of monetary policy shocks
σ_m	0.57	Auclert et al. (2020)	Standard deviation of monetary policy shocks

empirical estimates for the U.S. (Aguiar et al., 2025).

Regarding the supply side of the economy, the Calvo price stickiness parameter is set to $\theta_j = 0.75$ in both production industries, implying an average price duration of about one year (Bachmann et al., 2024). The third agent in the model is assigned the same parameter values as the HtM and non-HtM households, i.e. $\sigma_x = \sigma$ and $\zeta_x = \zeta$, while the parameter governing the responsiveness of the rate of return to this agent's level of debt is set to $\phi_B = 0.1$. We assume that the third agent is large, receiving an endowment equal to one thousand times the steady-state value of households' income in the economy during each period.

Finally, the parameters related to monetary policy are calibrated following Auclert et al. (2020). The Taylor rule coefficient on each sectoral inflation is set to $\phi_{\pi,j} = 1.5$, while the persistence parameter of the nominal interest rate is $\rho_R = 0.89$. Regarding the autoregressive monetary policy shocks, the persistence parameter of these series is $\rho_m = 0.89$, and the standard deviation of the exogenous innovation affecting these shocks is set to $\sigma_m = 0.83$.

IV.I Results

To study the aggregate implications of the forward-looking behavior of HtM households, this section analyzes the responses of the aggregate economy, as well as those of HtM and non-HtM households, to a one-standard-deviation expansionary monetary policy shock across different versions of the main model. To study the aggregate dynamics, we define aggregate inflation as $\pi_t = \phi\pi_{c,t} + (1 - \phi)\pi_{d,t}$ and aggregate output as $y_t = \phi y_{c,t} + (1 - \phi)y_{d,t}$, where ϕ denotes the steady-state share of non-durable expenditure. All variables shown in the figures are expressed as log-deviations from their respective steady-state, except for non-HtM households' bond holding, which is expressed as a share of households' steady-state nominal income.

HtM and non-HtM Responses. We begin by focusing on the heterogeneous responses of HtM and non-HtM households when both exhibit forward-looking behavior. Figure 9 presents the impulse response functions of relevant variables following an expansionary monetary policy shock. Consistent with the traditional New Keynesian narrative, non-HtM households respond directly to changes in the economy's interest rate through their Euler equation. Given the slow price adjustment, the decline in the real interest rate triggers an intertemporal substitution effect that raises the current consumption demand of these households. Due to price rigidities, intermediate firms operating under monopolistic competition are unable to immediately restore the economy's equilibrium by adjusting their prices. As a result, they expand their production level while also expanding their labor demand, putting upward pressure on real wages. These dynamics translate into counter-cyclical markups and lower firms' profits. The monetary expansion reaches HtM households mainly through higher labor income and reduced transfers associated with firms' profits.

Figure 9: HtM and non-HtM impulse responses to an expansionary monetary policy shock.

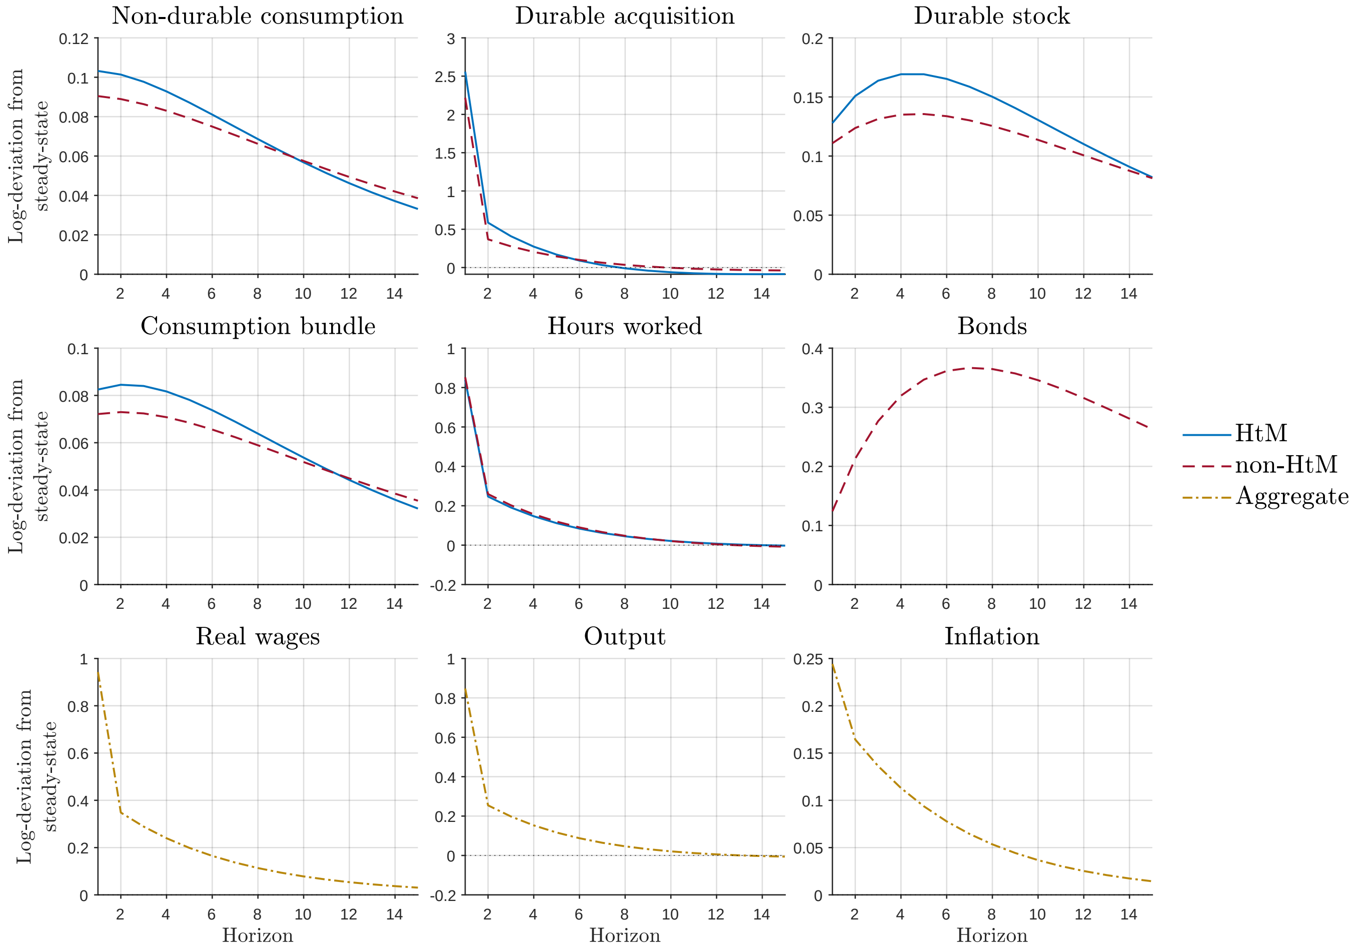

Note: The figure presents the estimated impulse responses to a one-standard-deviation expansionary monetary policy shock for HtM and non-HtM households. The blue solid line represents the responses of HtM households, while the red dashed line represents the responses of non-HtM households. The response of aggregate variables is represented with the yellow dash-dotted line.

Since these financially constrained households consume all of their higher income, they follow non-HtM households and expand their current consumption level. Notice that — consistent with the well-documented interest rate sensitivity of durable spending (Mankiw, 1985) — both groups of households sharply increase durable acquisitions. As a result, the stock of durables of both households builds gradually, peaking later due to the law of motion and depreciation dynamics. As prices gradually adjust, the effects of monetary policy dissipate, and aggregate inflation and output eventually return to their steady states as the economy re-equilibrates.

The estimated impulse responses in Figure 9 highlight the different intertemporal smoothing mechanisms of HtM and non-HtM households. One of the main visible differences across households' responses lies in their decisions on durable spending, which connects back to the partial equilibrium discussion in Section II. The larger durable acquisition by HtM households partly reflects their higher income level, but also their strategic use of durables as a vehicle for saving and intertemporal consumption smoothing. While non-HtM households increase current consumption and accumulate financial assets in response to the higher real wages, financially constrained HtM households rely on an alternative channel: they invest in durable goods to smooth consumption over time. The strong bond response of non-HtM households is consistent with their modest durable adjustment compared to HtM households. In contrast, the more pronounced and persistent response of HtM households in durable acquisition and durable stock, thus, reflects their reliance on durables as a substitute for financial savings. Examining the responses on the consumption bundle and aggregate output reveals that the possibility of durable consumption relaxes the consumption-income linkage typically characterizing the behavior of the myopic HtM household. Since the authorities manage transfers so that net transfers — after accounting for firms' realized profits — are equalized across households, the presence of the endowment-owning counterpart to the non-HtM households' saving position becomes essential for our results. Without this agent, there would be no source of heterogeneity in the responses of HtM and non-HtM households.

TANK Model with and without Durable Goods. To explore how durable consumption affects households' consumption behavior, we compare responses in two model variants: one allowing for durable consumption (TANK-D) and another where both available goods are non-durables (TANK).[26] Figure 10 shows the estimated impulse responses to an expansionary monetary policy shock for HtM and non-HtM households in both frameworks. First, consistent with the traditional TANK literature, the saving behavior among non-HtM households leads them to a modest consumption adjustment compared to HtM households in both models. Given their financial friction, HtM households exhibit larger on-impact consumption reactions across both goods in both frameworks. Second, as expected from the well-documented interest rate sensitivity of durables, expenditure on the second-good — which corresponds to the durable good in the TANK-D model — substantially amplifies households' expenditure responses, particularly among HtM households who front-load durable purchases to partially smooth consumption over time. As shown in the

[26] The TANK model follows the same structure as the TANK-D, except that it assumes the second good fully depreciates each period ($\delta = 1$), effectively behaving as a non-durable good.

Figure 10: HtM and non-HtM impulse responses to an expansionary monetary policy shock in the TANK-D and TANK models.

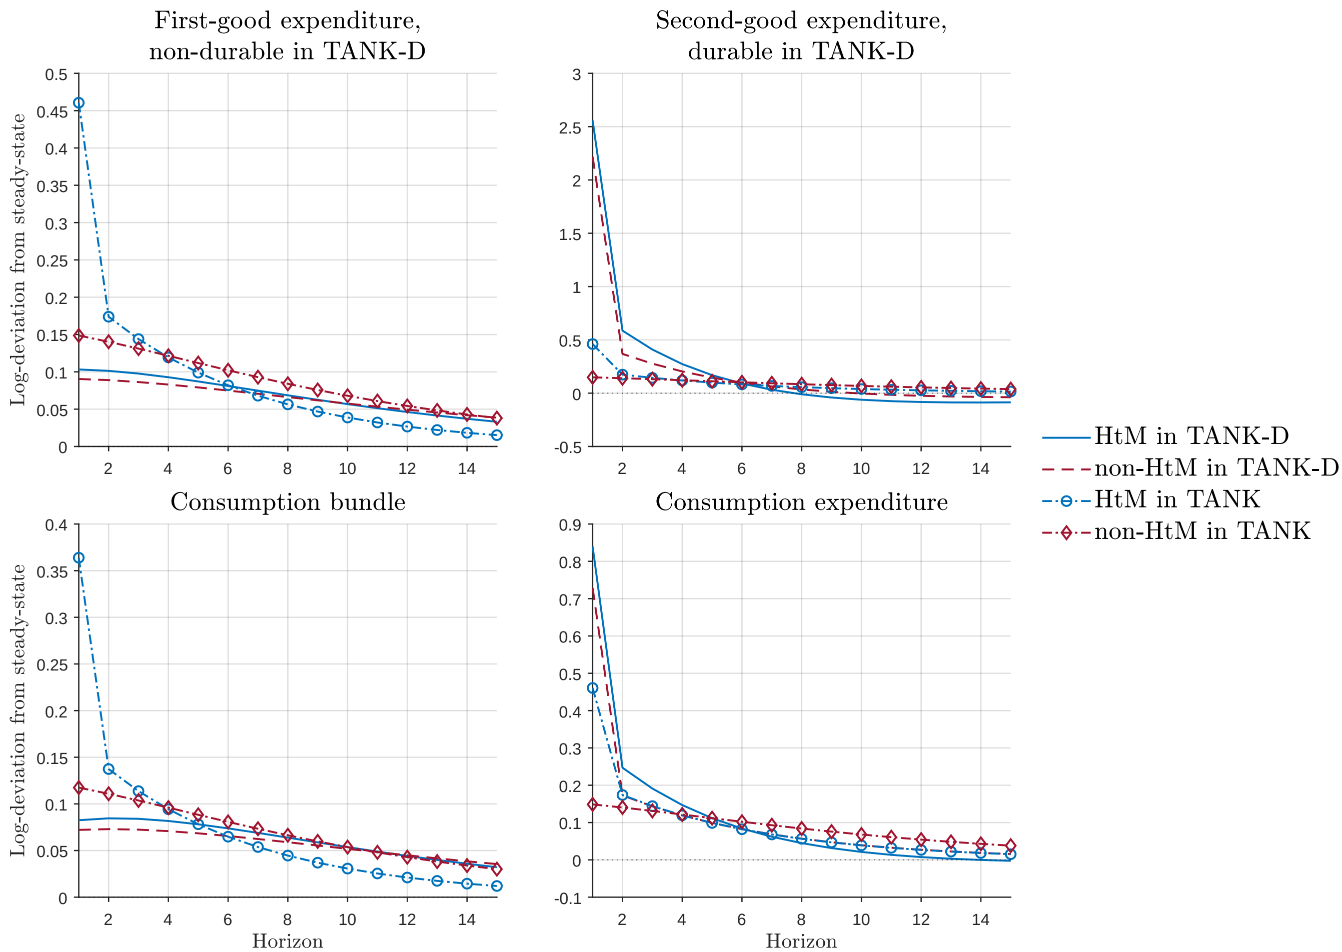

Note: The figure presents the estimated impulse responses to a one-standard-deviation expansionary monetary policy shock for HtM and non-HtM households in models with both durable and non-durable goods (TANK-D) and with two non-durable goods (TANK). The blue solid line represents the responses of HtM households in the TANK-D model, while the red dashed line represents the responses of non-HtM households in the same model. The blue dash-dotted line with circle markers represents the responses of HtM households in the TANK model, and the red dashed-dotted line with diamond markers represents the responses of non-HtM households in that model.

lower panels of the figure, and compared to the TANK model, the sharp reaction in durable acquisitions increases the volatility of households' overall consumption expenditure, but enables them to smooth their consumption bundle over time. These results are especially relevant for financially constrained HtM households. In conclusion, while introducing a durable good increases households' spending fluctuations, it stabilizes their overall consumption, especially among HtM households.

Figure 11: Aggregate impulse responses to an expansionary monetary policy shock in the TANK-D, TANK, and RANK-D models.

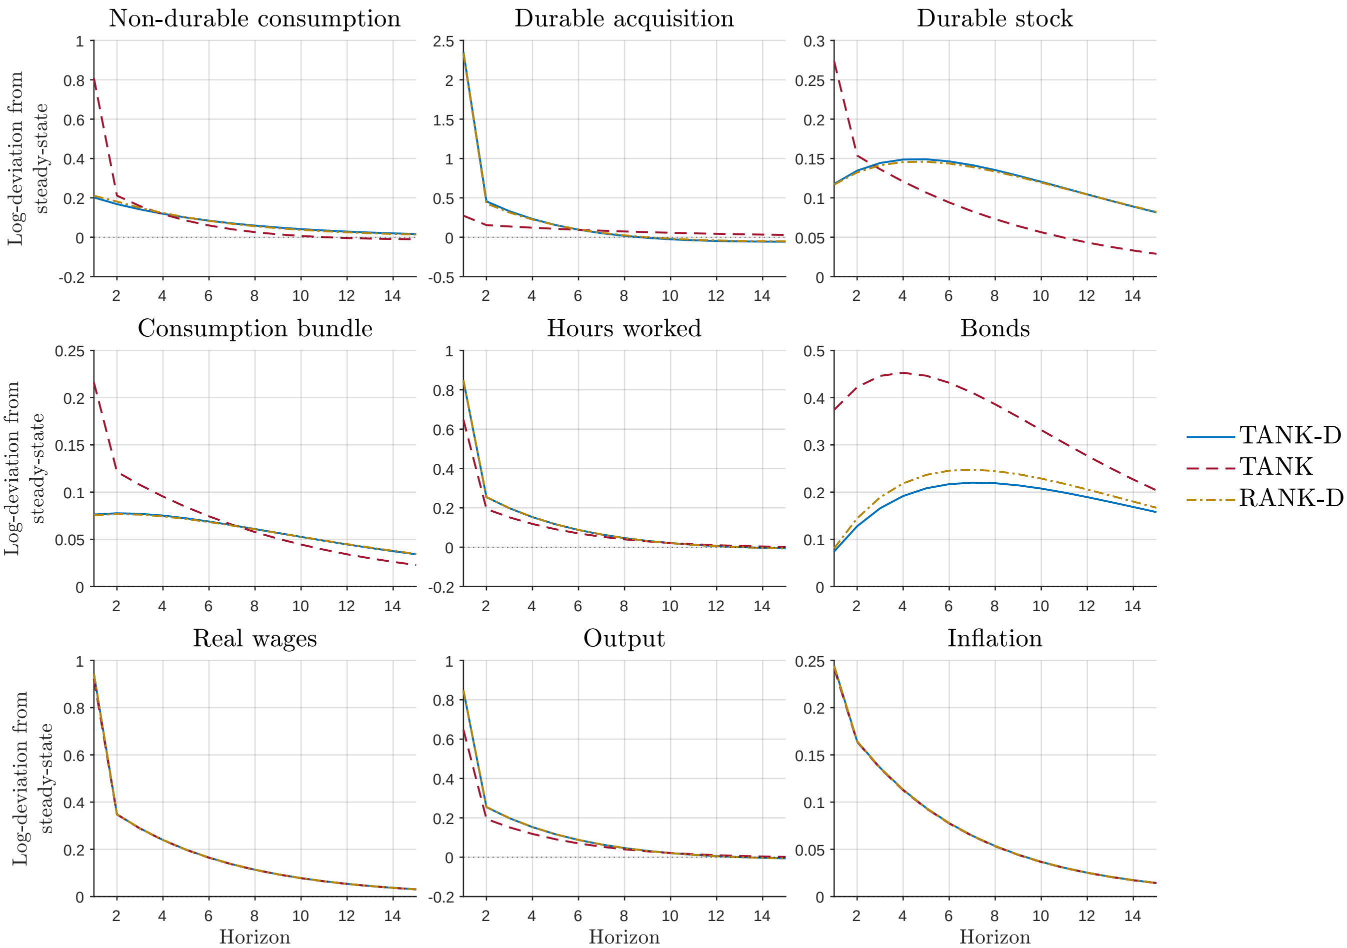

Note: The figure presents the estimated impulse responses of aggregate variables to a one-standard-deviation expansionary monetary policy shock. The blue solid line represents the aggregate responses in the TANK-D model, the red dashed line corresponds to the responses in the TANK model, and the yellow dash-dotted line represents the responses in the RANK-D model.

Once more, durables relax the tight contemporaneous link between income and consumption that typically characterizes HtM behavior.

Aggregate Effects of Household Heterogeneity. To examine the aggregate implications of household heterogeneity when a forward-looking dimension arises from consumption, we compare the aggregate impulse responses to an expansionary monetary policy shock across three models: TANK-D, TANK, and one where the economy is populated exclusively by durable-consuming

Figure 12: Aggregate output impulse responses to an expansionary monetary policy shock in the TANK-D and RANK-D models.

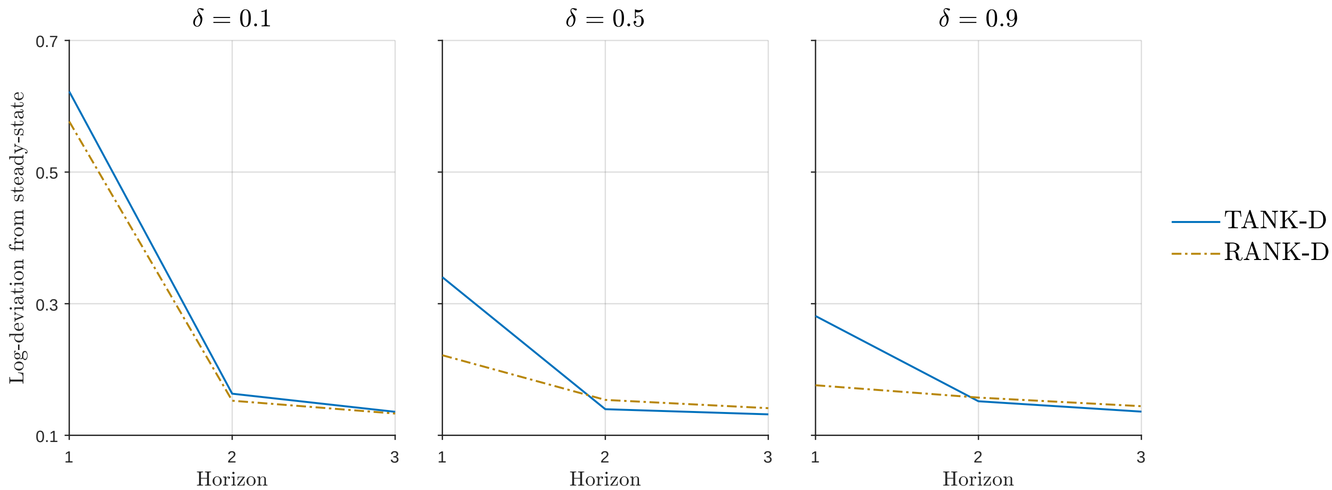

Note: The figure presents the estimated impulse responses of output — aggregated across HtM and non-HtM households — to a one-standard-deviation expansionary monetary policy shock under different levels of the durable goods depreciation rate, δ. The blue solid line represents the aggregate response in the TANK-D model, while the yellow dash-dotted line represents the responses in the RANK-D model.

non-HtM households. Traditionally, this latter case is identified as the representative-agent New Keynesian (RANK-D) model.[27] Figure 11 presents the impulse response functions of relevant aggregate variables in the TANK-D, TANK, and RANK-D models. A first observation is that the micro-level result — partially arising from the saving by consuming mechanism — that reflects the fact that introducing durables increases spending fluctuations while stabilizing overall consumption, also holds at the aggregate level. Comparing the impulse responses of aggregate output and the consumption bundle in the TANK-D and TANK models, we conclude that the sharp adjustments in durable acquisitions increase total production volatility but simultaneously allow the economy to smooth consumption over time. This pattern is consistent with the stronger bond accumulation observed in the TANK model relative to TANK-D.

Second, a notable finding from comparing the TANK-D and RANK-D models is that introducing forward-looking HtM households causes the aggregate responses in the TANK-D model to closely resemble those of the RANK-D model. In the literature, HtM households typically amplify aggregate responses to demand shocks, such as monetary policy shocks. However, once durable goods

[27] The RANK-D model follows the same structure as the TANK-D, except that it assumes that there are no HtM households in the economy ($\Lambda^{\text{H}} = 0$).

are introduced, HtM households engage in consumption smoothing through management of their durable stock, significantly dampening the role of household heterogeneity on aggregate outcomes. As a result, the TANK-D aggregate dynamics resemble those of the RANK-D model, where HtM households are entirely absent. The durable consumption and the consequent forward-looking behavior thus attenuate the traditional amplification associated with financially constrained households. Notice that, despite the similarities between the aggregate results of both frameworks, the model with heterogeneous agents continues to generate heterogeneous responses across household types.

As Figure 12 illustrates, this result critically depends on the depreciation rate of the durable good (δ). The figure shows the estimated impulse responses of output — aggregated across HtM and non-HtM households — to an expansionary monetary policy shock for a range of values of the depreciation parameter δ in the TANK-D and RANK-D models. First, across all levels of depreciation rates — and consistent with the traditional TANK literature — the results match the traditional amplification associated with the inclusion of financially constrained households. In each panel, the on-impact output response of the TANK-D model exceeds that of the RANK-D model. Second, the interest rate sensitivity of durable goods increases with their durability. In both models, on-impact responses to an expansionary monetary policy shock are larger when the depreciation rate is lower and the intertemporal benefits of durability increase. These larger output responses arise from households' larger durable acquisitions, particularly among HtM households, who front-load acquisitions to partially smooth consumption over time. One of the most important results from the figure is that as durability increases and the durables depreciation rate reduces, the responses of the TANK-D and RANK-D models converge. As anticipated in Figure 11, when durables persist over time — as in the panel where $\delta = 0.1$ — the heterogeneous-agent economy behaves more like the standard representative-agent framework, and household heterogeneity becomes less relevant for aggregate dynamics. On the other hand, when durables depreciate rapidly — as in the panel where $\delta = 0.9$ — heterogeneity across households amplifies the differences between the two models' responses to monetary policy shocks. Overall, the figure illustrates that introducing durable goods — and allowing HtM households to behave forward-looking through durable purchases — dampen the aggregate effects of heterogeneity. As δ decreases, the aggregate output responses in the TANK-D model increasingly resemble those of the RANK-D model, indicating that long-lived durables make the heterogeneous-agent economy behave more like a representative-agent environment.

V Conclusion

In this paper, we show that the forward-looking behavior of HtM households has a significant impact on both individual decisions and the aggregate economy. By allowing for durable consumption, our partial equilibrium model shows that even financially constrained HtM households can exhibit forward-looking behavior and smooth consumption over time. Since durable goods offer an intertemporal dimension by providing utility flows beyond the purchase period, HtM households use durables as an intertemporal margin of adjustment. Moreover, we show that there are significant behavioral differences between this forward-looking HtM household and the non-HtM household, and we provide experimental causal evidence consistent with the existence of this HtM forward-looking mechanism.

By incorporating durable goods into a TANK framework, we analyze the implications of the forward-looking HtM household for aggregate consumption dynamics and monetary transmission. We conclude that durables fundamentally reshape the transmission of monetary policy in heterogeneous-agent economies. Although durables amplify households' spending volatility — especially among HtM households who use durables as a saving vehicle — they stabilize aggregate consumption by relaxing the consumption-income linkage typically characterizing HtM households. More importantly, durables significantly dampen the role of households' liquidity constraints in aggregate results. As durability increases, the aggregate dynamics of the heterogeneous-agent economy increasingly resemble those of a single representative-agent framework.

From a policy perspective, the results suggest that the effectiveness of fiscal or monetary interventions depends critically on the composition of household consumption bundles and the durability of goods consumed. In particular, conventional monetary policy may be less significant in economies where durable goods provide alternative smoothing channels for financially constrained households. Our findings also suggest that economic authorities should remain attentive to the expectations of financially constrained households, despite limited access to financial markets.

References

M. Aguiar, M. Bils, and C. Boar. Who are the hand-to-mouth? *Review of Economic Studies*, 92 (3):1293–1340, 2025.

E. Aktug and O. E. Atesagaoglu. Inflation expectations and household spending in high inflation: Evidence from a randomized control trial. *Unpublished*, 2024.

F. Alvarez, R. E. Lucas, and W. E. Weber. Interest rates and inflation. *American Economic Review*, 91(2):219–225, 2001.

P. Andrade, E. Gautier, and E. Mengus. What matters in households' inflation expectations? *Journal of Monetary Economics*, 138:50–68, 2023.

O. Armantier, W. Bruine de Bruin, S. Potter, G. Topa, W. Van der Klaauw, and B. Zafar. Measuring inflation expectations. *Annual Review of Economics*, 5:273–301, 2013.

O. Armantier, S. Nelson, G. Topa, W. Van der Klaauw, and B. Zafar. The price is right: Updating inflation expectations in a randomized price information experiment. *The Review of Economics and Statistics*, 56(2):503–523, 2016.

O. Armantier, G. Topa, W. Van der Klaauw, and B. Zafar. An overview of the survey of consumer expectations. *Economic Policy Review*, 23(2):51–72, 2017.

A. Auclert, M. Rognlie, and L. Straub. Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model. *NBER Working Paper Series*, 26647, 2020.

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

R. Bachmann, B. Born, O. Goldfayn-Frank, G. Kocharkov, R. Luetticke, and M. Weber. A temporary vat cut as unconventional fiscal policy. *NBER Working Paper Series*, 2024.

A. Bar-Ilan and A. S. Blinder. Consumer durables: Evidence on the optimality of usually doing nothing. *Journal of Money, Credit and Banking*, 24(2):258–272, 1992.

R. Barsky, C. L. House, C. Boehm, and M. S. Kimball. Monetary policy and durable goods. *Unpublished*, 2016.

R. B. Barsky, C. L. House, and M. S. Kimball. Sticky-price models and durable goods. *American Economic Review*, 97(3):984–998, 2007.

M. Beraja and C. K. Wolf. Demand composition and the strength of recoveries. *NBER Working Paper Series*, Working Paper No. 29304, 2021.

D. Berger and J. Vavra. Consumption dynamics during recessions. *Econometrica*, 83(1):101–154, 2015.

B. Bernanke. Adjustment costs, durables, and aggregate consumption. *Journal of Monetary Economics*, 15(1):41–68, 1985.

F. O. Bilbiie. Limited asset market participation, monetary policy and (inverted) aggregate demand logic. *Journal of Economic Theory*, 140(1):162–196, 2008.

F. O. Bilbiie. The new keynesian cross. *Journal of Monetary Economics*, 114:90–108, 2020.

C. C. Binder and G. Brunet. Inflation expectations and consumption: Evidence from 1951. *Economic Inquiry*, 60(2):954–974, 2021.

W. Bruine de Bruin, W. Van der Klaauw, G. Topa, J. S. Downs, B. Fischhoff, and O. Armantier. The effects of question wording on consumers' reported inflation expectations. *Journal of Economic Psychology*, 33(4):749–757, 2012.

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

R. J. Caballero. Expenditure on durable goods: A case for slow adjustment. *The Quarterly Journal of Economics*, 105(3):727–743, 1990.

R. J. Caballero. Durable goods: An explanation for their slow adjustment. *Journal of Political Economy*, 101(2), 1993.

G. A. Calvo. Staggered prices in a utility-maximizing framework. *American Economic Review*, 12 (3):383–398, 1983.

J. Y. Campbell and N. G. Mankiw. Consumption, income, and interest rates: Reinterpreting the time series evidence. In *NBER Macroeconomics Annual 1989, Volume 4*, NBER Chapters, pages 185–246. National Bureau of Economic Research, Inc, 1989.

B. Candia. Inflation expectations and household spending: Different patterns in low and high-inflation settings. *Unpublished*, 2024.

A. Cavallo, G. Cruces, and R. Pérez-Truglia. Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics*, 9(3): 1–35, 2017.

M. P. Clements. Explanations of the inconsistencies in survey respondents' forecasts. *European Economic Review*, 54(4):536–549, 2010.

M. P. Clements. Probability distributions or point predictions? Survey forecasts of US output growth and inflation. *International Journal of Forecasting*, 30(1):99–117, 2014.

O. Coibion, Y. Gorodnichenko, and S. Kumar. How do firms form their expectations? New survey evidence. *American Economic Review*, 108(9):2671–2713, 2018.

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

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

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

O. Coibion, D. Georgarakos, Y. Gorodnichenko, and M. Weber. Forward guidance and household expectations. *Journal of European Economic Association*, 21(5):2131–2171, 2023b.

O. Coibion, Y. Gorodnichenko, E. S. Knotek, and R. Schoenle. Average inflation targeting and household expectations. *Journal of Political Economy Macroeconomics*, 1(2):403–446, 2023c.

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

F. D'Acunto, D. Hoang, M. Paloviita, and M. Weber. Cognitive abilities and inflation expectations. *AEA Papers and Proceedings*, 109:562–566, 2019.

F. D'Acunto, U. Malmendier, and M. Weber. Gender roles produce divergent economic expectations. *Proceedings of the National Academy of Sciense*, 118(21):1–10, 2021.

J. De Gregorio, P. E. Guidotti, and C. A. Végh. Inflation stabilisation and the consumption of durable goods. *The Economic Journal*, 108(446):105–131, 1998.

J. De Quidt, J. Haushofer, and C. Roth. Measuring and bounding experimenter demand effects. *American Economic Review*, 108(11):3266–3302, 2018.

D. Debortoli and J. Galí. Heterogeneity and aggregate fluctuations: Insights from TANK models. *NBER Working Paper Series*, Working Paper No. 32557, 2024.

A. M. Dietrich, K. Kuester, G. J. Müller, and R. S. Schoenle. News and uncertainty about covid-19: Survey evidence and short-run economic impact. *Journal of Monetary Economics*, 129:S35–S51, 2022.

L. Dräger and G. Nghiem. Are consumers' spending decisions in line with a euler equation? *The Review of Economics and Statistics*, 103(3):580–596, 2021.

J. C. Eberly. Adjustment of consumers' durable stocks: Evidence from automobile purchases. *Journal of Political Economy*, 102(3):403–436, 1994.

J. Galí, J. D. López-Salido, and J. Vallés. Rule-of-thomb consumers and the design of interest rate rules. *Journal of Money, Credit, and Banking*, 36(4):739–764, 2004.

J. Galí, J. D. López-Salido, and J. Vallés. Understanding the effects of government spending on consumption. *Journal of European Economic Association*, 5(1):227–270, 2007.

D. Georgarakos, Y. Gorodnichenko, O. Coibion, and G. Kenny. The causal effect of inflation uncertainty on households' beliefs and actions. *NBER Working Paper Series*, Working Paper No. 33014, 2024.

S. J. Grossman and G. Laroque. Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods. *Econometrica*, 58(1):25–51, 1990.

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

I. Hajdini, E. S. Knotek, J. Leer, M. Pedemonte, R. Rich, and R. Schoenle. Low passthrough from inflation expectations to income growth expectations: Why people dislike inflation. *Federal Reserve Bank of Cleveland*, Working Paper No. 22-21, 2022.

R. E. Hall. Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 86(6), 1978.

P. Huber. Robust estimation of a location parameter. *Annals of Mathematical Statistics*, 35(1): 73–101, 1964.

M. Jain, O. Kostyshyna, and X. Zhang. How do people view wage and price inflation? *Journal of Monetary Economics*, 145(103552), 2024.

J. H. Jiang, R. Kamdar, K. Lu, and D. Puzzello. How do households respond to expected inflation? an investigation on transmission mechanisms. *CAEPR Working Papers*, 004, 2024.

G. Kaplan and G. L. Violante. A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239, 2014.

G. Kaplan, G. L. Violante, and J. Weidner. The wealthy hand-to-mouth. *Brooking Papers on Economic Activity*, Spring, 2014.

G. Kaplan, B. Moll, and G. L. Violante. Monetary policy according to hank. *American Economic Review*, 108(3):697–743, 2018.

R. G. King, C. I. Plosser, and S. T. Rebelo. Production, growth and business cycles. *Journal of Monetary Economics*, 21:195–232, 1988.

E. S. Knotek, R. S. Schoenle, A. M. Dietrich, K. Kuester, G. J. Müller, K. O. R. Myrseth, and M. Weber. Consumers and covid-19: A real-time survey. *Federal Reserve Bank of Cleveland*, Economic Commentary(2020-08), 2020.

C. Kotsogiannis and P. Sakellaris. MPCs estimated from tax lottery, survey and administrative data. *Unpublished*, 2025.

J. V. Leahy and J. Zeira. The timing of purchases and aggregate fluctuations. *Review of Economic Studies*, 72(4):1127–1151, 2005.

B. Mackowiak and M. Wiederholt. Technical report, CEPR Discussion Papers, 2024.

N. G. Mankiw. Hall's consumption hypothesis and durable goods. *Journal of Monetary Economics*, 10(3):417–425, 1982.

N. G. Mankiw. Consumer durables and the real interest rate. *The Review of Economics and Statistics*, 67(3):353–362, 1985.

N. G. Mankiw, R. Reis, and J. Wolfers. Disagreement about inflation expectations. *NBER Macroeconomics Annual*, 18:209–270, 2003.

A. McKay and J. F. Wieland. Lumpy durable consumption demand and the limited ammunition of monetary policy. *Econometrica*, 89(6):2717–2749, 2021.

A. McKay, E. Nakamura, and J. Steinsson. The power of forward guidance revisited. *American Economic Review*, 106(10):3133–3158, 2016.

J. J. Rotemberg and M. Woodford. An optimization-based econometric framework for the evaluation of monetary policy. In B. S. Bernanke and J. J. Rotemberg, editors, *NBER Macroeconomics Annual 1997, Volume 12*, pages 297–346. MIT Press, Cambridge, MA, 1997.

C. Roth and J. Wohlfart. How do expectations about the macroeconomy affect personal expectations and behavior? *The Review of Economics and Statistics*, 102(4):731–748, 2020.

C. Roth, M. Wiederholt, and J. Wholfart. The effects of monetary policy: Theory with measured expectations. *Unpublished*, 2023.

S. Schmitt-Grohé and M. Uribe. Closing small open economy models. *Journal of International Economics*, 61:163–185, 2003.

P. Schnorpfeil, M. Weber, and A. Hackethal. Households' response to the wealth effects of inflation. *NBER Working Paper Series*, Working Paper No. 31672, 2023.

S. Stantcheva. How to run surveys: A guide to creating your own identifying variation and revealing the invisible. *Annual Review of Economics*, 15(1):205–234, 2022.

S. Stantcheva. Why do we dislike inflation? *Brookings Papers on Economic Activity*, Spring, 2024.

J. B. Taylor. Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, 39:195–214, 1993.

N. Vellekoop and M. Wiederholt. Inflation expectations and choices of households. *Leibniz Institute for Financial Research SAFE*, SAFE Working Paper Series 250, 2019.

M. Weber, F. D'Acunto, Y. Gorodnichenko, and O. Coibion. The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives*, 7(3):157–184, 2022.

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

S. P. Zeldes. Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy*, 97(2):305–346, 1989.

Appendix for

Saving by Consuming: The Intertemporal Behavior of Hand-to-Mouth Households

Mathieu Pedemonte

Inter-American Development Bank

Esteban Verdugo

University of Michigan

December, 2025

Contents

Appendix A Tables

Table A.1: Questions on respondents' HtM condition.

	Question		Answer
QH.1	Thinking about your current household's situation, how strongly do you agree or disagree with the statement "My household lives *paycheck-to-paycheck*"?	◯ ◯ ◯ ◯ ◯	I strongly agree. I somewhat agree. I neither agree nor disagree. I somewhat disagree. I strongly disagree.
QH.2	If your household needed to borrow one week's worth of your monthly household income, how easy or difficult would it be for your household to come up with the money?	◯ ◯ ◯ ◯ ◯	It would be extremely difficult. It would be relatively difficult. It would be neither easy nor difficult. It would be relatively easy. It would be extremely easy.
QH.3	What fraction of your income do you save?	◯ ◯ ◯ ◯ ◯	I spend all of my income each month. I spend more money than I earn. I often use credit cards or other loans to supplement my monthly income. I save around 10% of my monthly income. I save around 25% of my monthly income. I save at least 50% of my monthly income.
QH.4	Suppose your household loses its primary source of income today. How long can your household cover its weekly expenses with your household's total money in checking and savings accounts or other liquid financial assets such as cash, stocks, government and corporate bonds, and mutual and money market funds? Please give your best guess.	•	If my household loses its primary source of income today, we can cover our expenses for ______ week(s) using the money in our checking and saving accounts plus other liquid financial assets.

Table A.2: Questions on respondents' intended consumption expenditure.

	Question	Answer
QC.1	Compared with your spending last month, how do you expect your total spending to change in the next...[(1) one month/(2) one year]?	◯ Go up by ___ %. ◯ No change. ◯ Go down by ___ %.
QC.2	Compared with your spending last month, how do you expect your total spending on services - such as medical and dental care, haircuts, and restaurant meals - to change in the next...[(1) one month/(2) one year]?	◯ Go up by ___ %. ◯ No change. ◯ Go down by ___ %.
QC.3	Compared with your spending last month, how do you expect your total spending on non-durable goods - such as food and beverage products from grocery stores, and gasoline and other energy goods - to change in the next...[(1) one month/(2) one year]?	◯ Go up by ___ %. ◯ No change. ◯ Go down by ___ %.
QC.4	Compared with your spending last month, how do you expect your total spending on semi-durable goods - such as clothing and footwear, jewelry, watches, silverware, toys, tools and garden equipment, and household textiles and utensils - to change in the next...[(1) one month/(2) one year]?	◯ Go up by ___ %. ◯ No change. ◯ Go down by ___ %.
QC.5	Compared with your spending last month, how do you expect your total spending on durable goods - such as household appliances, radio and television sets, and sporting and wheeled goods - to change in the next...[(1) one month/(2) one year]?	◯ Go up by ___ %. ◯ No change. ◯ Go down by ___ %.

Table A.3: Questions on respondents' prior and posterior inflation expectations.

	Question		Answer
QI.1	Over the next 12 months, do you think that there will be inflation or deflation?	○ ○	Inflation. Deflation (the opposite of inflation).
↪	What do you expect the rate of [inflation/deflation] to be over the next 12 months? Please give your best guess.	•	I expect the rate of [inflation/deflation] to be ______ percent over the next 12 months.
QI.2	Over the next 5 years, do you think there will be inflation or deflation on average?	○ ○	Inflation. Deflation (the opposite of inflation).
↪	What do you expect the average annual rate of [inflation/deflation] to be over the next 5 years? Please give your best guess.	•	I expect the average annual rate of [inflation/deflation] to be ______ percent per year over the next 5 years.

Table A.4: Trial groups and information treatments.

Group	Treatment	Observations
Treatment 1	"*According to the Survey of Professional Forecasters conducted by the Federal Reserve Bank of Philadelphia, the annual aggregate inflation rate at the end of the year 2024 will be 2.5%.*"	1,410
Treatment 2	"*According to the Bureau of Labor Statistics, between 2012 and 2019, average apparel prices stayed relatively constant. On average, these prices decreased by 0.24% annually. In the last 12 months, average apparel prices have increased by 3.1%.*"	1,354
Placebo	"*According to the Census Bureau, the U.S. population grew 2.9% between 2018 and 2022.*"	1,370
Control	No treatment	1,395

Table A.5: Respondents' demographic and socioeconomic characteristics by HtM status.

	Full sample	non-HtM	HtM
Age (*years, sample mean*)	49.1	49.8	48.1
Gender (*% of sample*)			
◯ Female	49.9	43.8	58.6
◯ Male	49.8	55.9	41.2
◯ Other	0.3	0.3	0.2
Education level (*% of sample*)			
◯ Less than high school	1.9	1.0	3.2
◯ High school diploma or equivalent	22.0	18.5	26.8
◯ Some college, but not degree	26.8	23.0	32.3
◯ Bachelor's degree	30.6	32.8	27.4
◯ Master's degree	15.0	19.4	8.8
◯ Doctorate or professional degree	3.7	5.3	1.3
Civil status (*% of sample*)			
◯ Single	25.6	21.6	31.2
◯ Cohabiting partner	7.5	5.6	10.2
◯ Not cohabiting partner	1.6	1.0	2.4
◯ Married	48.8	58.7	34.7
◯ Divorced	11.5	8.9	15.2
◯ Widowed	5.1	4.3	6.3
Household net income (*% of sample*)			
◯ Less than $10,000	8.2	5.1	12.6
◯ Between $10,000 and $19,999	7.4	4.2	12.0
◯ Between $20,000 and $34,999	16.8	11.2	24.6
◯ Between $35,000 and $49,999	15.4	14.2	17.0
◯ Between $50,000 and $99,999	29.2	32.6	24.5
◯ Between $100,000 and $199,999	16.3	22.6	7.4
◯ More than $200,000	6.6	9.9	1.2
Number of children (*children, sample mean*)	1.50	1.55	1.44

continues in next page

Table A.5 – continued from previous page

	Full sample	non-HtM	HtM
Political affiliation (% *of sample*)			
◯ Democrat	40.2	44.3	34.4
◯ Independent	25.4	23.2	28.5
◯ Republican	30.3	29.8	31.0
◯ Other	4.1	2.7	6.1
English fluency (% *of sample*)			
◯ Native	72.0	71.0	73.2
◯ Fluent	27.4	28.4	25.8
◯ Less than fluent	0.6	0.6	1.0
Last month expenditure (% *of total expenditure*, *mean*)			
Durable goods	14.0	22.6	7.7
Semi-durable goods	5.7	9.0	3.3
Non-durable goods	16.6	22.2	12.5
Services	63.7	46.3	76.4
Aggregate	$5,124	$3,708	$7,129
Intended consumption growth, one month ahead (%, *mean*)			
Durable goods	0.70	1.43	-0.34
Semi-durable goods	0.69	1.42	-0.35
Non-durable goods	2.38	2.61	2.04
Services	1.63	2.19	0.84
Intended consumption growth, one year ahead (%, *mean*)			
Durable goods	1.38	1.65	1.01
Semi-durable goods	1.51	1.66	1.31
Non-durable goods	3.54	2.89	4.47
Services	2.50	2.62	2.33
Number of observations	5,529	3,231	2,298

Table A.6: Summary statistics on respondents' inflation expectations.

Panel A: Summary statistics on respondents' inflation expectations across surveys.

Inflation expectations	Michigan SCE		Full sample		HtM		non-HtM	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
12 months ahead	6.8	6.7	8.2	5.0	8.7	5.0	7.8	4.0
5 years ahead	8.3	7.5	6.8	4.0	6.9	5.0	6.6	4.0

Panel B: Summary statistics on respondents' inflation expectations across treatment status.

Inflation expectations	Control		Treatment 1		Treatment 2		Placebo	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
12 months ahead	9.0	5.0	7.5	5.0	8.1	4.0	8.2	4.9
5 years ahead	8.4	5.0	5.2	4.0	5.6	4.0	7.8	5.0

Table A.7: Balance tests on pre-treatment covariates.*

	T-test or χ^2-test p-value					
	C & T1	C & T2	C & P	T1 & T2	T1 & P	T2 & P
Age	0.92	0.82	0.23	0.92	0.34	0.40
Gender	0.54	0.37	0.99	0.81	0.59	0.44
Education level	0.55	0.14	0.82	0.13	0.39	0.75
Civil status	0.22	0.20	0.24	0.70	0.34	0.64
Household net income	0.09	0.32	0.07	0.05	0.10	0.05
Number of children	0.58	0.91	0.67	0.70	0.94	0.77
Political affiliation	0.51	0.95	0.74	0.27	0.77	0.45
English fluency	0.46	0.08	0.39	0.56	0.21	0.05

* **Note**: The columns in the table report p-values from balance tests conducted across treatment groups for different demographic and socioeconomic variables measured before treatment implementation. We conducted a t-test or a χ^2-test depending on the variable's nature. C denotes the passive control group, P corresponds to the placebo group, and T1 and T2 are the first and second treatment groups. The second row in the table provides information about the comparison groups in the analysis performed in each column.

Table A.8: Respondents' demographic and socioeconomic characteristics by treatment status.

Variable	Control	Treatment 1	Treatment 2	Placebo	Full sample
Age (*years, sample mean*)	48.7	48.9	49.4	49.6	49.1
Gender (% *of sample*)					
Female	48.4	51.0	51.9	48.4	49.9
Male	51.3	48.9	47.7	51.4	49.8
Other	0.4	0.2	0.4	0.2	0.3
Education level (% *of sample*)					
Less than high school	2.4	2.0	1.2	2.2	1.9
High school diploma or equivalent	21.7	23.3	21.2	21.7	22.0
Some college, but not degree	26.0	28.1	27.0	26.3	26.8
Bachelor's degree	30.6	29.7	32.0	30.0	30.6
Master's degree	15.3	13.4	15.2	16.4	15.0
Doctorate or professional degree	4.2	3.5	3.6	3.5	3.7
Civil status (% *of sample*)					
Single	23.7	26.6	25.3	26.6	25.6
Cohabiting partner	8.9	8.0	7.1	6.1	7.5
Not cohabiting partner	1.5	1.3	1.8	1.6	1.6
Married	50.3	47.1	48.8	48.9	48.8
Divorced	11.8	11.1	11.2	12.0	11.5
Widowed	3.8	5.8	5.8	4.9	5.1
Number of children (*children, sample mean*)	1.6	1.5	1.5	1.4	1.5
Political affiliation (% *of sample*)					
Democrat	41.3	38.3	40.0	41.3	40.2

continues in next page

Table A.8 – continued from previous page

Variable	Control	Treatment 1	Treatment 2	Placebo	Full sample
Independent	25.2	26.2	25.5	24.7	25.4
Republican	30.1	30.5	30.9	29.8	30.3
Other	3.4	5.0	3.6	4.2	4.1
Household's net income (% *of sample*)					
Less than \$10,000	7.2	9.7	8.1	7.8	8.2
Between \$10,000 and \$19,999	8.2	7.8	6.1	7.6	7.4
Between \$20,000 and \$34,999	16.8	16.6	18.4	15.3	16.8
Between \$35,000 and \$49,999	14.9	16.3	15.8	14.6	15.4
Between \$50,000 and \$99,999	27.6	29.7	29.7	31.0	29.2
Between \$100,000 and \$199,999	19.1	13.9	15.9	16.4	16.3
More than \$200,000.	6.1	5.9	6.1	8.3	6.6
English fluency (% *of sample*)					
Native	72.4	72.0	73.0	70.5	72.0
Fluent	26.5	27.4	26.7	28.7	27.4
Less than fluent	1.1	0.6	0.3	0.8	0.6
Last month expenditure (% *of expenditure*, *mean*)					
Durable goods	18.3	6.7	23.6	20.7	14.0
Semi-durable goods	8.9	2.3	8.3	9.3	5.7
Non-durable goods	23.3	9.2	24.9	22.6	16.6
Services	49.4	81.7	43.2	47.3	63.7
Consumption growth, one month ahead (%, *mean*)					
Durable goods	0.98	0.87	0.44	0.50	0.70
Semi-durable goods	0.83	0.22	0.86	0.86	0.69

continues in next page

Table A.8 – continued from previous page

Variable	Control	Treatment 1	Treatment 2	Placebo	Full sample
Non-durable goods	2.99	1.69	2.21	2.64	2.38
Services	1.98	1.17	1.78	1.61	1.63
Consumption growth, one year ahead (%, *mean*)					
Durable goods	1.55	0.85	1.89	1.29	1.38
Semi-durable goods	2.00	1.16	1.21	1.66	1.51
Non-durable goods	3.95	3.19	3.24	3.79	3.54
Services	3.04	2.21	2.69	2.08	2.50
Number of observations	1,395	1,410	1,354	1,370	5,529

Table A.9: Treatment effects on medium-run inflation expectations.*

	Average annual inflation expectations over the next 5 years				
	(1)	(2)	(3)	(4)	
				Regressor	Regressor$\times$H$_j$
$prior_j$	0.867***	0.863***	0.856***	0.868***	-0.033
	(0.011)	(0.012)	(0.012)	(0.013)	(0.023)
$treat_{j,1}$	1.771***	1.754***	1.798***	1.665***	0.284
	(0.191)	(0.193)	(0.197)	(0.224)	(0.289)
$prior_j \times treat_{j,1}$	-0.484***	-0.484***	-0.486***	-0.465***	-0.041
	(0.018)	(0.018)	(0.018)	(0.023)	(0.037)
$treat_{j,2}$	1.919***	1.916***	1.894***	1.553***	0.926***
	(0.196)	(0.199)	(0.204)	(0.226)	(0.324)
$prior_j \times treat_{j,2}$	-0.464***	-0.463***	-0.459***	-0.455***	-0.010
	(0.017)	(0.017)	(0.017)	(0.021)	(0.035)
$treat_{j,\mathrm{P}}$	0.150	0.149	0.149	-0.036	0.424
	(0.187)	(0.189)	(0.194)	(0.221)	(0.278)
$prior_j \times treat_{j,\mathrm{P}}$	0.004	0.003	0.001	0.000	0.005
	(0.016)	(0.016)	(0.016)	(0.019)	(0.035)
Observations	5,529	5,529	5,529	5,529	
Fixed-effects	No	Yes	Yes	Yes	
Controls	No	No	Yes	Yes	
R^2	0.784	0.786	0.792	0.793	

* **Note**: In all regressions, the dependent variable is the posterior belief measure $post_j$. The omitted category is the passive control group. Robust standard errors are reported in parentheses. The set of fixed effects includes week and U.S. state fixed effects, while the set of control variables includes age, gender, education level, household income, civil status, number of children, and political affiliation. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. The inflation expectations data have been trimmed at the weekly frequency by excluding the top and bottom 2.5% of the distribution.

Table A.10: Effects of aggregate inflation expectations on consumption, using pooled treatments.*

	$\%\triangle c^e_{j,k,h}$: Intended consumption growth on category k at horizon h							
	Durables		Semi-durables		Non-durables and Services		Aggregate	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$post_j \times (1 - H_j)$	0.114***	-0.017	0.100***	0.046	0.071**	0.009	0.089***	0.033
	(0.033)	(0.043)	(0.038)	(0.047)	(0.033)	(0.048)	(0.032)	(0.035)
$post_j \times H_j$	0.085*	0.008	0.101*	0.078	0.008	0.105	0.067	0.040
	(0.051)	(0.063)	(0.054)	(0.071)	(0.052)	(0.067)	(0.042)	(0.055)
Horizon	One month	One year	One month	One year	One month	One year	One month	One year
Observations	5,099	5,087	5,095	5,106	4,946	4,953	5,147	5,160
Fixed-effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F-stat, non-HtM	179.9	176.9	174.0	178.5	171.6	173.7	183.5	184.2
F-stat, HtM	154.8	157.4	158.5	159.0	145.1	151.3	154.0	155.9
R^2	0.134	0.069	0.139	0.079	0.150	0.077	0.223	0.205

* **Note**: In each regression, the dependent variable is the intended consumption growth for the consumption category shown in the column and the time horizon indicated in the row. Robust standard errors are reported in parentheses. The set of fixed effects includes week and U.S. state fixed effects, while the set of control variables includes age, gender, education level, household income, civil status, number of children, and political affiliation. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. The inflation expectations data have been trimmed at the weekly frequency by excluding the top and bottom 2.5% of the distribution.

Table A.11: Effects of durable inflation expectations on consumption, using Treatment 2.*

	$\%\triangle c^e_{j,k,h}$: Intended consumption growth on category k at horizon h							
	Durables		Semi-durables		Non-durables and Services		Aggregate	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$post_j \times (1 - H_j)$	0.096**	-0.020	0.112**	0.013	0.110**	0.052	0.103**	0.043
	(0.044)	(0.049)	(0.044)	(0.057)	(0.037)	(0.058)	(0.043)	(0.046)
$post_j \times H_j$	0.131**	0.044	0.106*	0.036	-0.015	0.136*	0.082*	0.013
	(0.063)	(0.075)	(0.062)	(0.099)	(0.037)	(0.058)	(0.049)	(0.076)
Horizon	One month	One year	One month	One year	One month	One year	One month	One year
Observations	3,788	3,776	3,796	3,810	3,689	3,690	3,837	3,843
Fixed-effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F-stat, non-HtM	190.0	186.5	188.3	201.1	198.2	197.1	208.3	203.7
F-stat, HtM	164.4	164.4	167.7	169.7	155.9	163.5	163.1	166.5
R^2	0.146	0.074	0.153	0.093	0.161	0.094	0.224	0.204

* **Note**: In each regression, the dependent variable is the intended consumption growth for the consumption category shown in the column and the time horizon indicated in the row. We use only Treatment 2 as an instrument. Robust standard errors are reported in parentheses. The set of fixed effects includes week and U.S. state fixed effects, while the set of control variables includes age, gender, education level, household income, civil status, number of children, and political affiliation. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. The inflation expectations data have been trimmed at the weekly frequency by excluding the top and bottom 2.5% of the distribution.

Appendix B Figures

Figure B.1: Eigenvalue and variance explained each principal component.

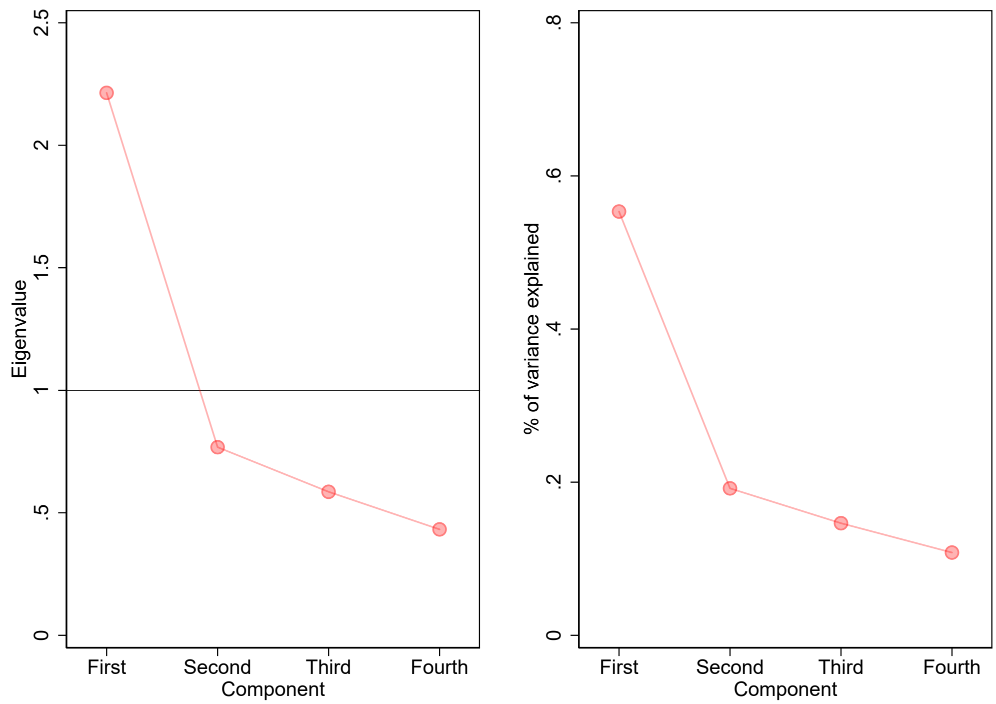

Note: The figure provides two scree plots summarizing the results of the PCA. The left panel illustrates the proportion of total variance explained by each principal component, while the right panel plots the eigenvalue associated with each principal component.

Figure B.2: Inflation expectations and actual inflation rate.

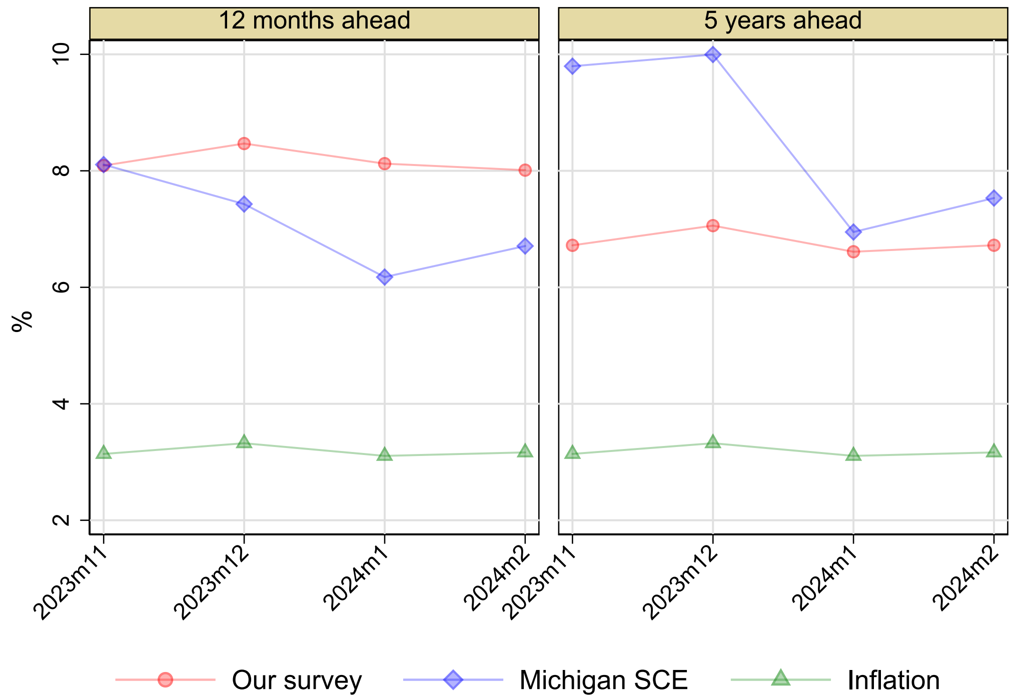

Note: The figure presents the time series of the mean of our collected data and the raw data from the University of Michigan's Survey of Consumers. The series plotted are the annual inflation expectations 12 months and 5 years ahead, alongside the actual inflation rate. Our collected inflation expectations data have been trimmed at the weekly frequency by excluding the top and bottom 2.5% of the distribution. The inflation expectations data from the University of Michigan's Survey of Consumers have been trimmed at the monthly frequency using the same percentile-based approach. The actual inflation data corresponds to the series "Consumer Price Index for All Urban Consumers: All Items in U.S. City Average" of the U.S. Bureau of Labor Statistics, and retrieved from FRED, Federal Reserve Bank of St. Louis.

Figure B.3: Kernel density of inflation expectations by treatment status.

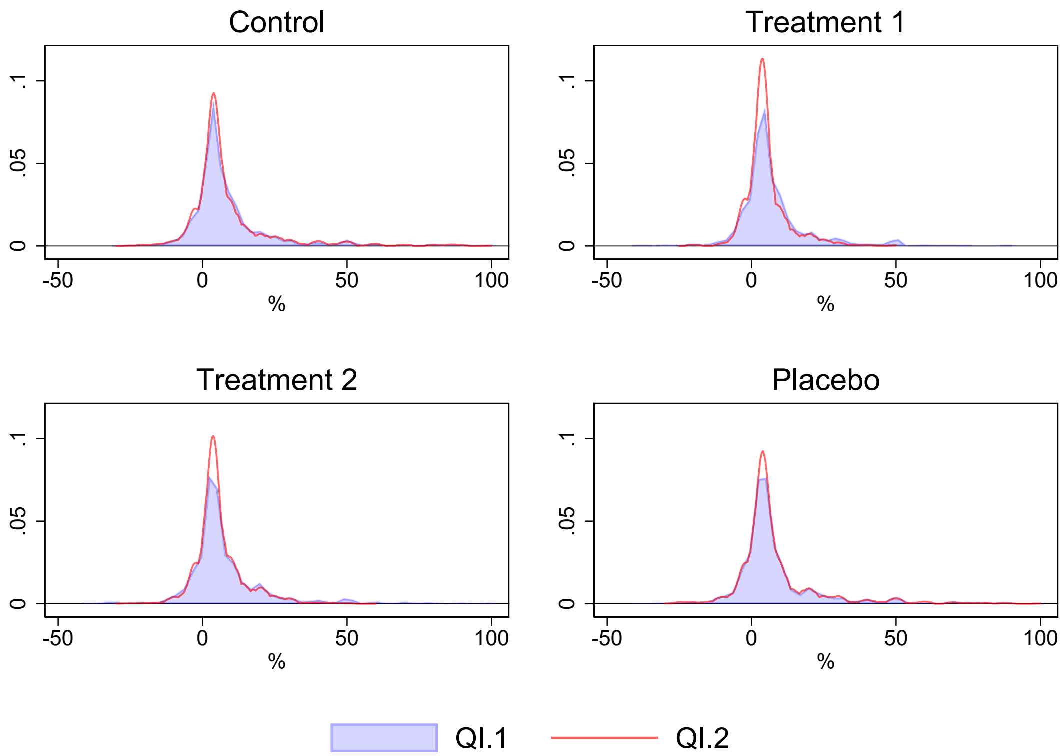

Note: The figure presents the kernel density of the answers to the prior (QI.1) and posterior (QI.2) questions by treatment group, where the density was estimated using a bandwidth of 1.5. The inflation expectations data have been trimmed at the weekly frequency by excluding the top and bottom 2.5% of the distribution.

Figure B.4: Treatment effects on long-run inflation expectations.

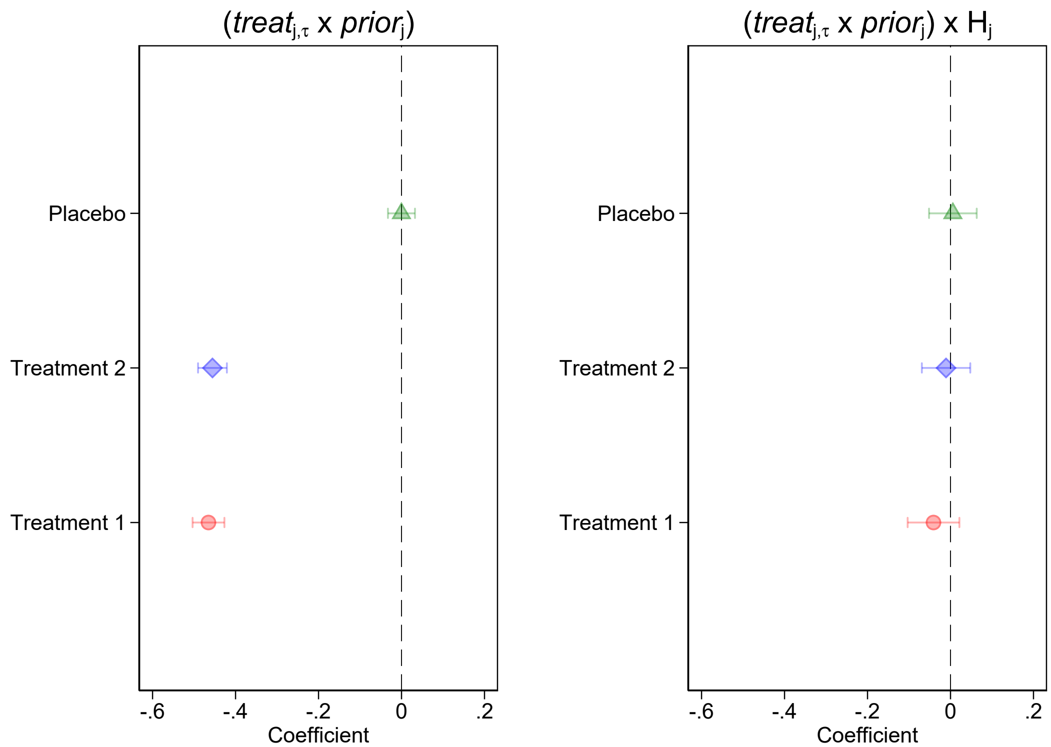

Note: The figure reports the regression coefficients of the experiments' information treatments, where the dependent variable is the posterior belief measure $post_i$. The omitted category is the passive control group. The regression includes week and U.S. state fixed effects, as well as a set of control variables including age, gender, education level, household income, civil status, number of children, and political affiliation. The inflation expectations data have been trimmed at the weekly frequency by excluding the top and bottom 2.5% of the distribution. The left panel reports the coefficients for non-HtM respondents, while the right panel illustrates the differential response of HtM respondents relative to non-HtM respondents. The regression uses robust standard errors, and the figure depicts 90% confidence intervals. Full regression results are presented in Table A.9.

Appendix C Proofs

C.1 A Two-Agents, Two-Goods Economy

This appendix section outlines the proofs behind the Proposition and corollaries drawn from the two-period model presented in Section II. First, we describe the procedure followed to derive the linear functions of log-deviation consumption for both representative agents. Then, we characterize the consumption responses of each agent when facing a change in the expected inflation rate of each of the different types of goods. Lastly, we examine the relationships between the magnitudes of these various responses.

Proof. **Lemma 1**

From the households' optimization problem, we derive the following key condition defining agents' optimal consumption allocations for this two-period model

$$mu^d_{i,t} \quad = \quad q_t mu^c_{i,t} - \beta_\delta \mathbb{E}_t \left[q_{t+1} mu^c_{i,t+1} \right] \mathbb{1}\left(t = 1\right) \tag{40}$$

where $i = \{\text{H}, \text{S}\}$ and $t = \{1, 2\}$. As in the main text, the function $\mathbb{1}(\vartheta)$ serves as a dummy indicator of the event ϑ, and we define the parameter $\beta_\delta = \beta(1-\delta) \in (0,1)$. Notice that the indicator function restricts the role of expectations in the last period of analysis, i.e., $t = 2$. Additionally, $\pi_{j,t} = p_{j,t}/p_{j,t-1}$ denotes the inflation rate of each of the prices $j = \{d, c, w\}$, and $q_t = p_{d,t}/p_{c,t}$ corresponds to the relative price between durable and non-durable goods. Similarly, we denote the economy's real wage as $w_t = p_{w,t}/p_{c,t}$. Furthermore, the variable $mu^\vartheta_{i,t} = \partial u(x_{i,t})/\partial \vartheta_{i,t}$ denotes the marginal utility related to the consumption good $\vartheta = \{d, c\}$.

For the household with access to the financial instrument, we obtain the following additional optimality condition corresponding to the household's traditional Euler equation

$$mu^c_{\text{S},t} \quad = \quad \beta R_{t+1} \mathbb{E}_t \left[mu^c_{\text{S},t+1} \left(\pi_{c,t+1}\right)^{-1} \right] \tag{41}$$

where $t=1$, in accordance with the transversality condition of the problem. We complement this set of optimality conditions using the corresponding optimization problem restrictions. For simplicity, and using the notation introduced earlier in the main text, we define the following set of elements

$$
\begin{aligned}
\psi_{0,t}^{i} &= \left[\psi_{5,t}^{i}\left(R_{\pi,t}\right)^{-1}\frac{\partial \grave{mu}_{i,t-1}^{c}}{\partial d_{i,t-1}}-\psi_{3,t-1}^{i}-\psi_{4,t}^{i}\psi_{6,t}^{i}\beta\left(\grave{\pi}_{q,t}\right)^{-1}\left(\grave{R}_{\pi,t}-\grave{R}_{\ell,t}\left(\delta\right)\right)\right] && (42)\\
\psi_{1,t}^{i} &= -\beta\left(\psi_{0,t}^{i}\grave{q}_{t}\right)^{-1} && (43)\\
\psi_{2,t}^{i} &= \left(\grave{q}_{t}\frac{\partial \grave{mu}_{i,t}^{c}}{\partial c_{i,t}}-\frac{\partial \grave{mu}_{i,t}^{d}}{\partial c_{i,t}}\right) && (44)\\
\psi_{3,t}^{i} &= \left(\grave{q}_{t}\frac{\partial \grave{mu}_{i,t}^{c}}{\partial d_{i,t}}-\frac{\partial \grave{mu}_{i,t}^{d}}{\partial d_{i,t}}\right) && (45)\\
\psi_{4,t}^{i} &= \grave{q}_{t}\left(\frac{\partial \grave{mu}_{i,t}^{c}}{\partial d_{i,t}}\frac{\partial \grave{mu}_{i,t}^{d}}{\partial c_{i,t}}-\frac{\partial \grave{mu}_{i,t}^{c}}{\partial c_{i,t}}\frac{\partial \grave{mu}_{i,t}^{d}}{\partial d_{i,t}}\right)\left(\psi_{2,t}^{i}\grave{q}_{t}-\psi_{3,t}^{i}\right)^{-1} && (46)\\
\psi_{5,t}^{i} &= \grave{R}_{\pi,t}\left(\psi_{2,t-1}^{i}-\psi_{4,t}^{i}\beta_{\delta}\grave{R}_{\pi,t}\right) && (47)\\
&\quad \left[\frac{\partial \grave{mu}_{t-1}^{c}}{\partial c_{t-1}}\mathbb{1}\left(i=\text{S}\right)+\frac{\partial \grave{mu}_{t-1}^{c}}{\partial d_{t-1}}\left(\grave{q}_{t-1}\right)^{-1}\mathbb{1}\left(i=\text{H}\right)-\psi_{4,t}^{i}\beta\left(\grave{R}_{\pi,t}\right)^{2}\left(\grave{q}_{t}\right)^{-1}\right]^{-1}\\
\psi_{6,t}^{i} &= \left(\psi_{5,t}^{i}-\grave{q}_{t-1}\grave{R}_{\ell,t}\left(\delta\right)\right) && (48)\\
\psi_{7,t}^{i} &= -\left(\psi_{2,t-1}^{i}-\psi_{4,t}^{i}\beta_{\delta}\grave{R}_{\pi,t}\right)^{-1} && (49)\\
\psi_{8,t}^{i} &= \left[\psi_{3,t-1}^{i}-\psi_{4,t}^{i}\beta_{\delta}\left(\grave{q}_{t-1}\right)^{-1}\left(\grave{R}_{\pi,t}-\grave{R}_{\ell,t}\left(\delta\right)\right)\right] && (50)\\
\psi_{9,t}^{i} &= \left[\left(\frac{\grave{mu}_{i,t}^{c}}{\psi_{2,t}^{i}}\right)\frac{\partial \grave{mu}_{i,t}^{d}}{\partial c_{i,t}}-\psi_{4,t}^{i}\left(\grave{\imath}_{i,t}-\frac{\grave{mu}_{i,t}^{c}}{\psi_{2,t}^{i}}\right)\right] && (51)
\end{aligned}
$$

where $t=2$ and $R_{\ell,t}\left(\delta\right)=\left(1-\delta\right)\grave{\pi}_{q,t}$. As mentioned in the main text, $\grave{\vartheta}$ denotes the value of the variable ϑ in the absence of disturbances to inflation expectations. As a next step, and considering exogenous changes in the expected inflation of each of the economy's prices, we compute the full derivatives of the set of equations that characterize the solution of the model. This procedure yields a linear system of equations, which define the relationship between changes in each of the model's control variables and changes in the expected values of the model's exogenous state variables. Due to the heterogeneity in household participation in financial markets, the number of equations and control variables of the system differ across household types. Solving this system, which captures the agent's optimal consumption and saving responses, we conclude that, given a change in the expected inflation rate $\pi_{j,2}$ of each of the prices $j=\{d,c,w\}$ the solution for the households' current durable expenditure satisfies

$$\frac{\partial i_{i,1}}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]} = \mathcal{D}_{1,1}^{i}\left(\left.\frac{\partial \mathbb{E}_1\left[q_2\right]}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]}\right|\grave{\Theta}\right)+\mathcal{D}_{2,1}^{i}\left(\left.\frac{\partial \mathbb{E}_1\left[w_2\right]}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]}\right|\grave{\Theta}\right)+\mathcal{D}_{3,1}^{i}\left(\left.\frac{\partial \mathbb{E}_1\left[R_{\pi,2}\right]}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]}\mathbb{1}\left(i=\mathrm{s}\right)\right|\grave{\Theta}\right) \tag{52}$$

where the functions $\mathcal{D}_{1,1}^{i}\left(\bullet\right), \mathcal{D}_{2,1}^{i}\left(\bullet\right), \mathcal{D}_{3,1}^{i}\left(\bullet\right) : \mathbb{R} \rightarrow \mathbb{R}$ are linear in their respective argument. Additionally, we can demonstrate that the function associated with the marginal variation in the expected value of the relative price between durable and non-durable goods corresponds to

$$\mathcal{D}_{1,1}^{i}\left(\vartheta_{j,t}\right|\grave{\Theta}\left.\right) = \psi_{1,t}^{i}\left[\psi_{5,t}^{i}\grave{mu}_{i,t}^{c}+\psi_{6,t}^{i}\psi_{9,t}^{i}\right]\vartheta_{j,t} \tag{53}$$

where $t = 2$ and $\vartheta_{j,t} = \partial \mathbb{E}_t\left[q_{t+1}\right]/\partial \mathbb{E}_t\left[\pi_{j,t+1}\right]$ in this case. Similarly, we can demonstrate that the function $\mathcal{D}_{2,1}^{i}\left(\bullet\right)$ exhibits a linear dependency on the response of the expected real wages in the economy when the inflation expectations on price $p_{j,t}$ are altered. We can formally express this relationship as

$$\mathcal{D}_{2,1}^{i}\left(\vartheta_{j,t}\right|\grave{\Theta}\left.\right) = \psi_{1,t}^{i}\psi_{4,t}^{i}\psi_{6,t}^{i}\vartheta_{j,t} \tag{54}$$

where $t = 2$ and $\vartheta_{j,t} = \partial \mathbb{E}_t\left[w_{t+1}\right]/\partial \mathbb{E}_t\left[\pi_{j,t+1}\right]$. Finally, notice that the remaining linear function applies exclusively to non-HtM households. This function uses as input the response of the real interest rate associated with the liquid bond when inflation expectations on price $p_{j,t}$ change. This function satisfies

$$\mathcal{D}_{3,1}^{i}\left(\vartheta_{j,t}\right|\grave{\Theta}\left.\right) = \psi_{1,t}^{i}\left[\psi_{4,t}^{i}\psi_{6,t}^{i}\left(\frac{\grave{b}_{i,t-1}}{\grave{p}_{c,t-1}}\right)-\psi_{5,t}^{i}\left(\frac{\grave{q}_t\grave{mu}_{i,t}^{c}}{\grave{R}_{\pi,t}}\right)\right]\vartheta_{j,t} \tag{55}$$

where $t = 2$ and $\vartheta_{j,t} = \partial \mathbb{E}_t\left[R_{\pi,t+1}\right]/\partial \mathbb{E}_t\left[\pi_{j,t+1}\right]\mathbb{1}\left(i=\mathrm{s}\right)$. Regarding non-durable consumption, the solution of the system of equations reveals that, in response to a change in the expected inflation rate $\pi_{j,2}$ of each price $j = \{d, c, w\}$, households' current reactions in non-durable expenditure satisfies

$$\frac{\partial c_{i,1}}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]} = \mathcal{C}^i_{0,1}\left(\frac{\partial i_{i,1}}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]}\right) + \mathcal{C}^i_{1,1}\left(\frac{\partial \mathbb{E}_1\left[q_2\right]}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]}\middle|\,\grave{\Theta}\right) + \mathcal{C}^i_{2,1}\left(\frac{\partial \mathbb{E}_1\left[w_2\right]}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]}\middle|\,\grave{\Theta}\right) + \mathcal{C}^i_{3,1}\left(\frac{\partial \mathbb{E}_1\left[R_{\pi,2}\right]}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]}\mathbb{1}\left(i = \mathrm{s}\right)\middle|\,\grave{\Theta}\right) \tag{56}$$

where the functions $\mathcal{C}^i_{0,1}(\bullet), \mathcal{C}^i_{1,1}(\bullet), \mathcal{C}^i_{2,1}(\bullet), \mathcal{C}^i_{3,1}(\bullet) : \mathbb{R} \to \mathbb{R}$ are also linear in their respective argument. We can demonstrate that, given the corresponding input reaction $\vartheta_{j,t}$, each of these functions satisfies the following definition

$$\mathcal{C}^i_{0,1}\left(\vartheta_{j,t}\middle|\,\grave{\Theta}\right) = \psi^i_{7,2}\psi^i_{8,2}\vartheta_{j,t} \tag{57}$$

$$\mathcal{C}^i_{1,1}\left(\vartheta_{j,t}\middle|\,\grave{\Theta}\right) = \psi^i_{7,t}\psi^i_{9,t}\beta_\delta\vartheta_{j,t} \tag{58}$$

$$\mathcal{C}^i_{2,1}\left(\vartheta_{j,t}\middle|\,\grave{\Theta}\right) = \psi^i_{4,t}\psi^i_{7,t}\beta_\delta\vartheta_{j,t} \tag{59}$$

$$\mathcal{C}^i_{3,1}\left(\vartheta_{j,t}\middle|\,\grave{\Theta}\right) = \psi^i_{4,t}\psi^i_{7,t}\beta_\delta\left(\frac{\grave{b}_{i,t-1}}{\grave{p}_{c,t-1}}\right)\vartheta_{j,t} \tag{60}$$

where $t = 2$. From the first-period budget restriction associated with the optimization problem of the non-HtM household, we can derive the optimal saving response of this household based on its optimal consumption reaction. As a result, we conclude

$$\frac{\partial b_{\pi,1}}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]} = -\frac{\partial c_{\mathrm{s},1}}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]} - \grave{q}_1\frac{\partial d_{\mathrm{s},1}}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]} \tag{61}$$

Therefore, the set of equations described by equations (52)-(61), along with the corresponding definition of elements involved in each of these functions, determines the optimal responses of households' durable and non-durable consumption, as well as non-HtM real savings, as outlined in the main text.

□

Proof. **Proposition 1**

We eliminate the durability of the goods in the economy by assuming that the stock of durable products fully depreciates each period. Specifically, we set $\delta = 1$, where the parameter δ is the depreciation rate of the stock of durables owned by households. This assumption leads to a calibration where $(\beta_\delta)|_{\delta=1} = (\psi^{\text{H}}_{5,2})|_{\delta=1} = (\psi^{\text{H}}_{6,2})|_{\delta=1} = 0$. With this assumption in place, and given the corresponding consumption reactions of the HtM household outlined in equations (52)-(60), we arrive at the HtM unresponsiveness conclusion presented in the main text.

To study the consumption reactions of the non-HtM household in the scenario of no durable goods in the economy, we repeat the procedure employed when analyzing the HtM case. Setting $\delta = 1$ results in a calibration where $(\beta_\delta)|_{\delta=1} = 0$, $(\psi^{\text{S}}_{7,2})|_{\delta=1} = -(\psi^{\text{S}}_{2,1})^{-1}|_{\delta=1}$, and $(\psi^{\text{S}}_{8,2})|_{\delta=1} = (\psi^{\text{S}}_{3,1})|_{\delta=1}$. Replacing these conditions into the set of equations governing the households' current reactions in non-durable expenditure, we conclude

$$\left.\frac{\partial c_{\text{S},1}}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]}\right|_{\delta=1} = -\left(\grave{q}_1 \frac{\partial \grave{mu}^c_{\text{S},1}}{\partial d_{\text{S},1}} - \frac{\partial \grave{mu}^d_{\text{S},1}}{\partial d_{\text{S},1}}\right)\left(\grave{q}_1 \frac{\partial \grave{mu}^c_{\text{S},1}}{\partial c_{\text{S},1}} - \frac{\partial \grave{mu}^d_{\text{S},t}}{\partial c_{\text{S},1}}\right)^{-1} \left.\frac{\partial i_{\text{S},1}}{\partial \mathbb{E}_1\left[\pi_{j,2}\right]}\right|_{\delta=1} \tag{62}$$

This equation implies that, in response to a change in inflation expectations, the non-HtM household adjusts its consumption proportionally across both categories of goods, consistent with the discussion stated in the main text.

□

Proof. **Proposition 2**

In this partial equilibrium framework, changes in the expected inflation rate of the economy's nominal wage affect only one expected relative price: real wages. As established in Lemma (1), the linear function $\mathcal{D}^i_{2,1}(\bullet)$, defined by equation (54), governs households' optimal durable consumption reaction given variations in the expected value of the economy's real wages. For HtM households, we can demonstrate that the signs of the relevant elements involved in this function are $\psi^{\text{H}}_{1,2}, \psi^{\text{H}}_{4,2} > 0$ and $\psi^{\text{H}}_{6,2} < 0$. These signs directly imply a negative relationship between the consumption response of the HtM household in terms of durable expenditure and their expectations of the inflation rate of the economy's nominal wage, as stated in the main text.

Using the set of equations (52)-(55), we can demonstrate that the solution for the optimal durable consumption response from the HtM household satisfies

$$\frac{\partial i_{\text{H},1}}{\partial \mathbb{E}_1 [\pi_{j,2}]} = \beta_\delta \left(\psi^{\text{H}}_{0,2}\right)^{-1} \left(\psi^{\text{H}}_{9,2} \frac{\partial \mathbb{E}_1 [q_2]}{\partial \mathbb{E}_1 [\pi_{j,2}]} + \psi^{\text{H}}_{4,2} \frac{\partial \mathbb{E}_1 [w_2]}{\partial \mathbb{E}_1 [\pi_{j,2}]}\right) \tag{63}$$

Substituting this optimal durable consumption reaction into the linear functions $\mathcal{C}^{\text{H}}_{1,1}(\bullet)$ and $\mathcal{C}^{\text{H}}_{2,1}(\bullet)$, and incorporating the conclusion $(\psi^{\text{H}}_{0,2} + \psi^{\text{H}}_{8,2}) = -q_1(\psi^{\text{H}}_{7,2})^{-1}$ into the equation (56), which characterizes households' optimal non-durable consumption response given variations in the expected value of the different relative prices of the economy, we conclude that for the HtM household the relationship between their consumption responses in durable and non-durable goods satisfy

$$\frac{\partial c_{\text{H},1}}{\partial \mathbb{E}_1 [\pi_{j,2}]} = -\grave{q}_1 \frac{\partial i_{\text{H},1}}{\partial \mathbb{E}_1 [\pi_{j,2}]} \tag{64}$$

This expression simply reflects this agent's first-period budget constraint. As a result, and given the negative consumption response in terms of durable expenditure, we conclude that the HtM response in non-durable goods must be positive.

To study non-HtM consumption reactions, we repeat the procedure used in the HtM case. Lemma (1) established that households' optimal durable consumption reaction to a variation in the expected value of the economy's real wages follows the linear equation (54). In the case of non-HtM households, and using two simplifying assumptions that are not strictly necessary to reach our conclusions, we can determine the sign of each component directly involved in this function. First, assuming separable preferences over durable and non-durable goods, we obtain $\psi^{\text{S}}_{4,2} > 0$. Additionally, since the real interest rate is sufficiently high such that $\grave{R}_{\pi,2} > \grave{R}_{\ell,t}(\delta)$, we conclude $\psi^{\text{S}}_{6,2} > 0$, and these two previous assumptions lead to $\psi^{\text{S}}_{1,2} > 0$. These signs directly imply a positive relationship between the response of the durable consumption of non-HtM households and their expectations about the inflation rate of future nominal wages, as stated in the main text.

Regarding the non-durable consumption response of the saving household, separable preferences between durable and non-durable goods allow us to conclude $\psi^{\text{S}}_{3,1}, \psi^{\text{S}}_{4,2}, \psi^{\text{S}}_{7,2} > 0$. Given the additional assumption that the real interest rate is sufficiently high such that $\grave{R}_{\pi,2} > \grave{R}_{\ell,t}(\delta)$, it follows that $(\grave{q}_1)^{-1}(\grave{R}_{\pi,2} - \grave{R}_{\ell,t}(\delta))\psi^{\text{S}}_{1,2}\psi^{\text{S}}_{4,2}\psi^{\text{S}}_{6,2} \in (0,1)$. Replacing these results into the equations (56) and (59), which together fully describe the linear dependency of the non-HtM non-durable consumption response when changing the expected inflation rate on the economy's nominal wages, we conclude that these variables also exhibit a positive relationship, as stated in the main text. Finally, the result on the optimal real savings response of non-HtM households follows directly from both previous findings and considering equation (61).

□

Proof. **Proposition 3**

Given the linear structure of households' consumption responses to changes in inflation expectations, as described by the set of equations (52)-(60), HtM will remain unresponsive if there is no expected variation in the relative prices q_2 and w_2 following a shift in inflation expectations. Therefore, a sufficient condition for the HtM insensitivity result is

$$\sum_{j=\{d,c,w\}} \frac{\partial \mathbb{E}_1 [\vartheta_2]}{\partial \mathbb{E}_1 [\pi_{j,2}]} = 0 \tag{65}$$

where $\vartheta = \{q, w\}$. Notice that a sufficient condition for this equality to hold is that the variation rate in inflation expectations must be the same across all the relevant nominal prices. In other words, for the HtM unresponsiveness result to hold, we require $d \ln (\mathbb{E}_1 [\pi_{j,2}]) = \bar{\pi}$ with $j = \{d, c, w\}$ and $\bar{\pi}$ being a constant.

Notice that the previous result, where the condition $d \ln (\mathbb{E}_1 [\pi_{j,2}]) = \bar{\pi} \neq 0$ for $j = \{d, c, w\}$ leads to no expected variation in the economy's relative prices q_2 and w_2, still applies. However, unlike HtM households, when facing an equal percentage change in the expected inflation across all nominal prices $p_{j,t}$ for $j = \{c, d, w\}$, the non-HtM households still adjust their consumption in response to the consequent change in the expected real return of the economy's financial instrument. Lemma (1) established that households' optimal durable consumption response to a change in the expected real interest rate follows the linear equation (55). By imposing two simplifying but non-essential assumptions, we can easily determine the sign of each relevant component directly involved in this function. First, assuming separable preferences over durables and non-durables, we obtain $\psi^{\mathrm{s}}_{4,2}, \psi^{\mathrm{s}}_{5,2} > 0$. Additionally, since the initial real interest rate is sufficiently high such that $\grave{R}_{\pi,2} > \grave{R}_{\ell,t}(\delta)$, it follows that $\psi^{\mathrm{s}}_{6,2} > 0$, and these two previous assumptions together imply $\psi^{\mathrm{s}}_{1,2} > 0$. As a result, these sign conditions directly imply a positive aggregate response in the durable consumption of a net borrower non-HtM household, i.e. $\grave{b}_{\text{s},1} \leq 0$, as stated in the main text.

Since the proposed condition on the variation of inflation expectations leaves relative prices q_2 and w_2 unchanged while changing the expected real return of the economy's traditional saving instrument $R_{\pi,2}$, the non-durable consumption reaction of the non-HtM household is fully characterized by equations (56) and (60). Since the assumption of separable preferences over durable and non-durable goods leads to $\psi^{\text{s}}_{4,2}, \psi^{\text{s}}_{7,2} > 0$, we simply require $\psi^{\text{s}}_{8,2} > 0$ so the net borrower non-HtM household reacts by reducing non-durable consumption as an aggregate response to the change in inflation expectations. This element $\psi^{\text{s}}_{8,2}$ exhibits a positive sign only if

$$\grave{R}_{\pi,2} \quad < \quad (1-\delta)\,\grave{\pi}_{q,2} + \grave{q}_1 \frac{\partial m\grave{u}^d_{\text{s},1}}{\partial d_{\text{s},1}} \left(\grave{q}_2 \frac{\partial m\grave{u}^c_{\text{s},2}}{\partial c_{\text{s},2}} + \frac{1}{\grave{q}_2} \frac{\partial m\grave{u}^d_{\text{s},2}}{\partial d_{\text{s},2}} \right) \left[\beta\,(1-\delta) \frac{\partial m\grave{u}^c_{\text{s},2}}{\partial c_{\text{s},2}} \frac{\partial m\grave{u}^d_{\text{s},2}}{\partial d_{\text{s},2}} \right]^{-1} \tag{66}$$

as stated in the main text. Finally, the result on the optimal real savings response of this net borrower household follows directly from both previous consumption findings and equation (61).

□

$$\phi^i_{c,t} \quad = \quad \left[\frac{\grave{c}_{i,t}}{\grave{c}_{i,t} + \grave{q}_t \grave{i}_{i,t}} \right] \tag{67}$$

$$\phi^i_{w,t} \quad = \quad \left[\frac{\grave{w}_t}{\grave{w}_t + (1-\delta)\,\grave{q}_t d_{i,t-1} + \left[\left(\grave{R}_t \grave{b}_{i,t-1} - \grave{b}_{i,t} \right) / \grave{p}_{c,t} \right] \mathbb{1}\,(i = \text{s})} \right] \tag{68}$$

$$\phi^i_{d,t} \quad = \quad \left[\frac{(1-\delta)\,\grave{q}_t \grave{d}_{t-1}}{\grave{w}_t + (1-\delta)\,\grave{q}_t \grave{d}_{i,t-1} + \left[\left(\grave{R}_t \grave{b}_{i,t-1} - \grave{b}_{i,t} \right) / \grave{p}_{c,t} \right] \mathbb{1}\,(i = \text{s})} \right] \tag{69}$$

where $t = \{1, 2\}$. Notice that these ratios satisfy $\phi^i_{w,t} = 1$ when $i = \text{H}$ and $t = 1$, and $\phi^i_{d,t} = 0$ when $t = 1$.

Appendix D Survey Design and Treatment Validation

We use data from a novel module added to the Federal Reserve Bank of Cleveland's Survey of Consumer Expectations, one of the leading U.S. consumer expectations surveys. As mentioned in the main text, this module was specifically designed for our experiment, following several rounds of feedback and pretesting multiple versions of the questionnaire with expert and non-expert groups. Before the official launch, we conducted and evaluated a small-scale pilot questionnaire using the Qualtrics platform. The main survey has a daily frequency and a repeated cross-sectional format, and it draws respondents from different, actively managed, double-opt-in market research panels complemented by social media recruitment. The sample is restricted to respondents aged 18 and above who live in the U.S. Beyond these restrictions, it is representative of the U.S. population in terms of various key demographic factors, including age, gender, education, ethnicity, net income, and location. The survey implements several safeguards against fraudulent or inattentive participation. It reduces the probability of bots engaging in abusive activities on the experiment by using reCAPTCHA scores to filter out responses from likely illegitimate users. Moreover, to complement these measures, we included ex-ante methods to check respondents' attention, such as tracking their time answering each question and their click patterns. We also applied ex-post filters to exclude individuals with careless responses and inconsistent answers in the survey, and included a standard question to evaluate their numerical abilities. Following recommendations suggested by Haaland et al. (2023), we excluded respondents who completed the survey in an unreasonably short or long time. To mitigate biases such as experimenter demand effects or social desirability, we ensured the complete anonymity of the respondents and provided a monetary reward for completing the survey to induce truthful revelation.[28] Considering all these restrictions and filters, our final sample comprises around 5,500 responses collected between November 21, 2023, and February 13, 2024, with an average of around 500 responses per week.

In the initial section of the survey, each respondent completed a questionnaire collecting demographic and socioeconomic information. Table A.5 in Appendix A provides some descriptive statistics on the main sample results for these variables. On average, respondents are approximately 49 years old, and the median respondent has an annual disposable income that falls within the $35,000 to $49,999 range. Half of the sample identifies as female, and over two-fifths express

[28] Moreover, demand effects tend to be weaker in online survey settings, which is the mode through which we collected our data (De Quidt et al., 2018).

more affinity with the Democratic Party. More than half of the respondents are married or cohabiting with a partner, and the majority report having at least one child, with a mean value of 1.5 children per respondent. Most respondents have at least some college education, and when asked about the highest education level completed, most selected the bachelor's degree option. The median survey completion time was around 12 minutes.

D.1 HtM Status

As explained in the main text, identifying HtM respondents is challenging. Because obtaining detailed household financial information is costly, the literature has used different strategies to identify potential HtM candidates. In this paper, we adopted a strategy mixing those used in the literature, and based our identification using respondents' answers to the four questions presented in Table A.1, located in Appendix A. Figure D.5 highlights the first of these questions, where we ask "*Thinking about your current household's situation, how strongly do you agree or disagree with the statement "My household lives paycheck-to-paycheck"?*". We designed this qualitative, ordinal, closed-ended question to explore respondents' self-perception of their HtM status using a well-known U.S. colloquial expression - "living *paycheck-to-paycheck*". The term is widely understood to describe individuals who cannot cover basic living expenses if they suddenly lose access to their regular source of income. These individuals, likely financially constrained, often lack liquidity between paychecks, and when faced with a change in expectations, liquidity constraints heavily determine their consumption responses. The feature of having minimal liquidity immediately before receiving a paycheck closely resembles the agent in the theoretical literature, who is unable to carry assets from one period to the next (Kaplan and Violante, 2014). Although we use only five answer options, they are exhaustive, covering all possible answers that respondents may provide. Moreover, we directly map the answers' scale to the question in a balanced, bipolar, ordinal format, partially addressing a possible acquiescence bias. Following the recommendations of Stantcheva (2022), we included a middle option that reflects a possible indifference of the survey respondent.

The remaining three questions in Table A.1 were developed following the same design principles as the first one. In particular, the second question in the list asks respondents how easy or difficult it would be for them to borrow an amount equivalent to one week's worth of their monthly income. As explained in the main text, this threshold is motivated by the taxonomy introduced by Kaplan et al. (2014), where one of their definitions classifies households as financially constrained if their

Figure D.5: First of the questions included to study respondents' HtM condition.

QH.1 Thinking about your current household's situation, how strongly do you agree or disagree with the statement "My household lives *paycheck-to-paycheck*"?

- ◯ I strongly agree.
- ◯ I somewhat agree.
- ◯ I neither agree nor disagree.
- ◯ I somewhat disagree.
- ◯ I strongly disagree.

liquid wealth is less than one week's worth of their earnings. This qualitative, ordinal, closed-ended question provides a direct self-assessment of the respondent's perceived access to formal or informal credit. This multiple-choice scaled question elicits respondents' opinions on how easy it would be for them to borrow an amount proportional to their household income. Access restrictions from formal or informal financial markets are a key feature of the theoretical literature agent, restricted from borrowing options. We get information on effective current financial market participation through the third question on the list, which asks respondents to report the fraction of their income they typically save. Finally, the fourth question most closely mirrors the standard approach in the literature for identifying HtM households. This final quantitative question gets information on the number of weeks respondents' financial assets could cover their usual expenses. Drawing on Kaplan et al. (2014), the question provides a rough intuition of liquid financial asset variables by inquiring about checking and savings accounts, stocks, government and corporate bonds, and mutual money market funds. Similar to the usual literature approach, this question implicitly evaluates households' liquid assets position relative to an earning or consumption measure.

We assign a numerical value ranging from one to five for each of the qualitative, ordinal, closed-ended questions included in this section. In other words, the first response option is recoded as 1, the second as 2, and so on, up to 5 for the last option. According to this scale, lower values correspond to a higher probability of being classified as HtM. Given the relatively strong correlation between the HtM respondents' classifications obtained using each question by itself and the internal consistency of their response scales, we employ a principal component analysis (PCA) to synthesize the multivariate descriptive information provided by this set of questions. More precisely, we observe the distribution of the first PCA component and classify as HtM the respondents who are at the bottom 40% of this distribution. This first component represents the linear combination of the responses to the original four HtM questions that explains the largest share of their common variation. In our case, it captures almost 60% of the total variance of these variables. Table A.5 in Appendix A details respondents' demographic and socioeconomic characteristics by HtM status.

D.2 Consumption Expenditure

In the information provision literature, second-stage variables are the dependent variables that we causally influence by affecting respondents' beliefs through the implementation of treatments (Stantcheva, 2022). In our experiment, these second-stage variables are the measures of intended consumption expenditure that we collect. As explained in the main text, we gather information on past consumption before treatment implementation, and we collect post-treatment information on future consumption behavior. First, we ask respondents how much they spent last month on four categories of goods and services that together cover the entire possible range of consumption elements. These categories are services and non-durable, semi-durable, and durable goods. The respondents provided dollar amounts for each consumption category, which we aggregated across these four groups to obtain our overall measure of past consumption expenditure. We use these pre-treatment consumption measures to enhance our estimation accuracy, as they enable us to control for individual consumption tendencies (Stantcheva, 2022). Table A.5 in Appendix A provides details on the past consumption behavior for the total sample and each household type. On average, respondents dedicate around 80% of their consumption to non-durable goods and services. This figure reduces to about 70% for non-HtM respondents and increases to nearly 90% for HtM respondents.

We then ask about expected consumption behavior. After treatment implementation, we included a block of questions collecting information on respondents' intended spending in overall consumption and the previously described four categories.[29] As the main text explains, we have considered two horizons for these questions to capture expectations over different time frames and read our results as a possible impulse response function. We asked respondents about their planned consumption expenditure one month and one year ahead. We adopt the two-step framework suggested by some

[29] Although actual consumption data would be preferable given that intended consumption may be influenced by survey noise, measurement error, respondents' overconfidence in their ability to overcome liquidity constraints, and other distortions, analysis conducted by Schnorpfeil et al. (2023) concludes that survey respondents' spending plans closely track their actual spending behavior.

Figure D.6: Question about respondents' planned consumption of semi-durable goods.

QC.4 Compared with your spending last month, how do you expect your total spending on semi-durable goods - such as clothing and footwear, jewelry, watches, silverware, toys, tools and garden equipment, and household textiles and utensils - to change in the next…

(1) One month.
- ◯ Go up by ______ %.
- ◯ No change.
- ◯ Go down by ______ %.

(2) One year.
- ◯ Go up by ______ %.
- ◯ No change.
- ◯ Go down by ______ %.

influential household surveys to elicit information on beliefs (Armantier et al., 2017), which requires survey respondents to answer by selecting whether they believe their consumption expenditure would increase or decrease, and then specify the corresponding growth rate as a percentage. Table A.2 in Appendix A shows the exact wording of these questions. The question instructed them to use last month's consumption spending as a reference point to compute the growth rate. Each question gives a rough intuition of the elements comprehended by the corresponding category. In particular, Figure D.6 illustrates the approach we followed, focusing on the semi-durable goods question. It mentions clothing and footwear, jewelry, watches, silverware, toys, tools and garden equipment, and household textiles and utensils as examples of semi-durable goods. Importantly, the information collected in this module corresponds to our set of main outcome variables.

D.3 Inflation Expectations

One of the key components of information provision experiments is the respondents' belief that researchers intend to influence through the information treatment implementation. These beliefs serve as the experiment's first-stage variables. In our setup, they correspond to two measures of inflation expectations. We follow one of the literature's standard approaches to gather information about respondents' inflation beliefs. This methodology is widely adopted by highly influential household surveys such as the University of Michigan's Survey of Consumers (MSC) and the Federal Reserve Bank of New York's Survey of Consumer Expectations (NYSCE). Although some major surveys include probability-distribution questions on expectations, we opt for a simpler, more tractable alternative and rely on quantitative point forecast-type answers. The literature argues that directly eliciting point estimates, rather than using probability distribution predictions, provides more accurate information regarding first-moment predictions (Clements, 2010, 2014). After the section inquiring about respondents' demographic and socioeconomic information, the questionnaire presents two sets of questions asking about the expected evolution of aggregate prices, both using the same frequency but differing in the forecast horizon. We administer the information treatments between these two questions. In line with the MSC, we ask respondents about their annual inflation expectations 12 months and 5 years ahead. The reasons behind this horizon difference are the usual concerns about survey fatigue and the repetition of questions in a survey (Haaland et al., 2023). Using the same question twice or more could lead to methodological problems such as consistency pressure, experimenter demand, and social desirability bias (Stantcheva, 2022). Following one of the standard literature solutions, we elicited beliefs on a highly related but different outcome. As in Coibion et al. (2023c), we used similar but not identical questions and elicited annual inflation beliefs for two different horizons. To gather these inflation expectations data, we implemented the two-step framework suggested by the NYSCE to elicit point forecast (Armantier et al., 2017). Table A.3 in Appendix A shows the exact wording employed in these questions. This method integrates qualitative and quantitative components to diminish a possible classical measurement error. Specifically, the survey participants answer each inflation expectation question by selecting whether they believe there will be inflation or deflation, then typing the corresponding inflation or deflation rate as a percentage. This procedure allows for straightforward and meaningful interpersonal comparisons.

In our setup, and using the language of the information provision literature, the variable associated with short-term inflation expectations corresponds to the respondents' prior. These expectations are the respondents' beliefs elicited before receiving or not receiving one of the information treatments included in the experiment. Conversely, the medium-run inflation expectations variable serves as respondents' updated beliefs or posterior. We collect these expectations immediately after respondents receive or do not receive an information treatment. By eliciting priors for the treated and the control groups before treatment implementation, we mitigate concerns related to priming, increase our ability to capture genuine belief updating, and enable us to examine heterogeneous treatment effects based on the alignment between this first-stage variable and the provided information (Haaland et al., 2023).

Table A.6 presents summary statistics on respondents' inflation expectations. As the first columns reveal, we use the median and mean across all MSC respondents as benchmarks to analyze and validate our collected data.[30] The fact that our forecast horizons match those used in the University of Michigan's survey facilitates comparisons between our inflation expectations results and those reported by the MSC. The table reports statistics for both inflation expectations questions, using the full sample as well as samples restricted by HtM status. Between November 2023 and February 2024, the mean 12-month-ahead forecast across all MSC's respondents was 6.8%, with a median of 6.7%. Despite the 2.5% trimming procedure applied to our version of this series to mitigate the influence of outliers, more than 20% of our survey's respondents still reported inflation expectations of at least 20% for the following 12 months. To address such extreme responses, and following recent contributions in this literature (**?**Coibion et al., 2023b, 2024; Georgarakos et al., 2024), we report statistics robust to these extreme perceptions. For the 12-month-ahead forecast, the Huber-robust mean of our full sample is 8.2%, while the corresponding median is substantially lower at 5.0%. As the remaining columns of Table A.6 document, we observe some heterogeneity in expectations when splitting our sample into the two household categories. Both the mean and median inflation expectations tend to be higher for HtM candidates compared to non-HtM candidates. Regarding 5-year-ahead expectations, the MSC reports a raw average of approximately 8.3% during our relevant period, while the corresponding median is 7.5%. The full-sample mean of our collected data for this post-treatment variable is slightly lower, at nearly 7.0%, while the median is considerably lower, at 4.0%. HtM households also report higher values for these longer-term expectations than their non-HtM counterparts. However, the statistics for this post-treatment variable are not directly comparable, as they combine information from groups of respondents who received different information treatments.

[30] Although this influential University of Michigan survey also adopts a two-step framework to elicit point forecasts (Armantier et al., 2017), the wording of its inflation expectations questions differs slightly from the one used in our survey. Specifically, the MSC asks: "*During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?*". This question is followed by: "*By about what percent do you expect prices to go <up><down>, on average, during the next 12 months?*". The emphasis in these questions is on the evolution of prices in general rather than on the inflation rate per se. As shown in Armantier et al. (2013), respondents' interpretation of inflation can vary substantially depending on the wording and framing of the questions. Therefore, comparisons between the results of these two surveys should be interpreted with caution. In particular, Bruine de Bruin et al. (2012) concludes that respondents' average expectations are lower and less dispersed when questions are focused on inflation rather than prices in general.

Finally, Figure B.2 presents the time series of the average inflation expectations at the two different time horizons, using data from both previously described sources, alongside the actual inflation rate over the same period. Our collected data aligns with several of the most robust patterns documented in the literature on eliciting household inflation expectations from survey data (Weber et al., 2022). In particular, our data reveal that medium-run inflation expectations are consistently lower than short-run expectations. A pattern that some researchers interpret as evidence that the population's medium-run expectations are more firmly anchored than their short-run counterparts. Moreover, our data also exhibit a systematic upward bias in numerical expectations relative to the actual inflation rate. This bias is a widespread pattern observed in different prominent expectations surveys, and the literature relates it to various specific socioeconomic and demographic elements, including cognitive abilities (D'Acunto et al., 2019) and gender (D'Acunto et al., 2021), among others. Finally, there seems to be a positive correlation between updating short-run and medium-run inflation expectations. This degree of co-movement is frequently used in the literature to indicate the degree of anchoring of expectations. When shocks in the economy influence the inflation beliefs of respondents over the short run and lead these individuals to revise their expectations on a longer horizon, it suggests the perception that these shocks are not short-lived.

D.4 Treatments Description

As the main text explains, we opted for an information provision experiment to obtain cleanly identified evidence of a causal relationship from inflation expectations into consumption choices (Armantier et al., 2016; Cavallo et al., 2017; Coibion et al., 2018). We modified the information set of a well-defined group of respondents to correct misperceptions and evaluate their learning and belief-updating process. To achieve this, we designed the survey by creating four groups of respondents and randomly assigning each participant to one of these groups. Table A.7 confirms that respondents' observable characteristics do not systematically predict their treatment assignment, confirming the success of the randomization process. Table A.4 outlines the size of each group and the information treatment they received. As previously mentioned, we administered treatments after eliciting priors and before collecting respondents' posterior beliefs. Notice that all treatments were deliberately designed to be concise and straightforward in content, and they provide respondents with truthful, official, and publicly available information. Therefore, we presume respondents pay attention to the information received and consider it credible. Moreover, since we elicit beliefs on an objective, quantifiable benchmark, such as the annual aggregate inflation rate measure, we can closely link each treatment information to the elicited belief.

The treatment of the first group provided quantitative information on the aggregate inflation expectations of professional forecasters. As the main text explains, the 1,416 respondents in this group were informed that, according to professional forecasters, the annual aggregate inflation rate at the end of 2024 would be 2.5%. The treatment clarifies that the source of this projection is the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters. The figure considered in this treatment was the most recent average expectation of professional forecasters at the time we initiated our survey implementation. The second group received qualitative and quantitative information about past inflation, specifically related to a semi-durable good. For these 1,329 respondents, the treatment informed that apparel prices remained stable throughout most of the decade of the 2010s; however, during the last 12 months preceding the survey implementation, these prices increased by 3.1%. The treatment mentions that the source of the information is the U.S. Bureau of Labor Statistics. This treatment was designed to induce variation in inflation expectations by highlighting a targeted change in relative prices. As a result, we interpret that changes in the prices of semi-durable and durable goods mostly drive any resulting variation in inflation expectations from this treatment. The third group comprises 1,386 respondents and serves as our placebo group. They received quantitative information that is arguably neutral regarding the issue of interest. The treatment stated that the U.S. population grew 2.9% between 2018 and 2022, citing the U.S. Census Bureau as the source of information. One of the key roles of this placebo group is to help us account for potential unconscious numerical anchoring issues. By providing a figure arguably irrelevant to the inflation expectation formation process but comparable in magnitude to that given in both actual treatments and also framing the information using as a reference an official government source, it allows us to assess whether respondents' inflation expectations were influenced by numerical priming rather than the content of the information (Haaland et al., 2023). When testing treatment effects on belief updating, the placebo results will help us to evaluate the possible presence of numerical anchoring bias in our experiment. Lastly, the fourth and final group, comprising 1,391 respondents, received no information treatment and serves as our passive or pure control group. Table A.8 provides descriptive statistics on key demographic and socioeconomic variables by treatment status. Similarly, Table A.7 reports the p-value results of conducting balance tests across the different treatment groups and for several demographic and socioeconomic variables, which are measured before treatment implementation. As expected from the simple random assignment procedure, the sample is mostly balanced across the different treatment arms. Only the variable related to household net income reveals significance at the 10% confidence level in some comparisons. Our main regression analysis includes this slightly unbalanced variable as a control.

As previously discussed, a key element of information provision experiments is the respondents' beliefs that researchers intend to influence through the implementation of the treatment. Figure B.3 presents the estimated kernel density of the inflation expectations data for each treatment group. From this figure emerges one of the most prominent features associated with survey data eliciting household inflation expectations, which corresponds to the substantial disagreement existing among respondents (Mankiw et al., 2003). However, notice that despite the wide dispersion of responses, the distributions of prior beliefs appear broadly similar across treatment groups, suggesting a relatively homogeneous belief structure before treatment implementation. Importantly, this profound disagreement is not unique to our data. Since it holds across surveys conducted in different regions and time periods, specific survey design characteristics are unlikely to be the reason behind this feature. In particular, Weber et al. (2022) argues that the main driving force behind this pattern is likely the heterogeneous sources agents rely on to form expectations and the differences in interpretation of the economic consequences of common shocks. Moreover, the figure reveals that prior and posterior beliefs have nearly identical densities in the case of the pure control group. Although the two densities appear broadly similar for the placebo group, the posterior beliefs exhibit some slight bunching around the figure provided in this treatment. However, this visual pattern is still insufficient to conclude whether the placebo treatment had a significant influence on respondents' beliefs. While the evidence is suggestive, it does not allow us to confirm or rule out the presence of numerical anchoring bias in our experiment. Finally, the densities associated with posterior beliefs for both treatment groups exhibit a notably larger kurtosis compared to prior beliefs, suggesting that both information treatments are successful in affecting respondents' belief updating.

D.5 Treatments Validation

In a Bayesian learning framework, individuals update their beliefs when new information becomes available. This approach allows us to empirically characterize how arguably previously unknown information influences respondents' expectations and validate the effectiveness of our treatments. Let $prior_j$ denote the respondent j's point forecast for the aggregate inflation rate over the next 12 months. Remember that this variable captures respondents' beliefs before they received or did not receive one of the information treatments included in the experiment. Similarly, let $post_j$ denote the respondent j's point forecast on the average annual inflation rate over the next 5 years. Again, this variable reflects respondents' updated inflation beliefs and is collected immediately after they are exposed to or not exposed to an information treatment. The intentional difference in the time horizons between these two questions is a methodological choice to mitigate concerns about survey fatigue and question repetition. Using the exact same question twice or more could introduce methodological issues such as consistency pressure, experimenter demand, and social desirability bias (Stantcheva, 2022). Following one of the solutions in the literature, we elicited beliefs on a highly related but different outcome.

Under a Bayesian learning framework, we expect survey respondents to revise their beliefs by weighing their prior expectations and the signal received through the information treatment. Figure 4 illustrates this idea by depicting the relationship between prior and posterior variables, stratified by treatment status. In particular, the figure shows that for the passive control group in our experiment, the slope is less than one, likely reflecting that the population's medium-run inflation expectations are more anchored than their short-run expectations. Since this group answered the posterior question without receiving any external information, this relationship serves as our benchmark against which we evaluate and validate the experiment's treatments. In contrast, we observe a significantly flatter relationship between prior and posterior beliefs for the first two treatment groups. The slopes for these groups are much smaller, suggesting that treated respondents effectively weigh the provided signal in the treatment and sacrifice some of the relative weight assigned to their prior when updating beliefs. In other words, when answering the posterior question and compared to the respondents in the pure control group, the average respondent in the first two treated groups adjusts their expectations closer to the figures provided in the treatments. Respondents located at the left side of the belief distribution tend to increase their expectations on average, while those at the right side tend to revise them downward on average. Notably, the figure suggests that respondents in these first two treated groups similarly updated their beliefs when receiving the treatment, even though the information provided was conceptually different. Finally, the group receiving the third treatment of the experiment reveals a slope indistinguishable from that of the control group. This outcome aligns with our expectations, as this group served as our placebo group. As discussed earlier, one of the roles of this group is to help us circumvent unconscious numerical anchoring issues, given the fact that its associated treatment provides a figure arguably irrelevant to the inflation expectation process but comparable in magnitude to that given in both previous treatments, while also mentioning an official government entity as the source of information. Given the lack of significant visual differences between the results of the control and placebo groups, we anticipate that there is no evidence of numerical anchoring bias in our experiment. Formally, we run the following regression presented in equation (12) where the dichotomic variable $treat_{j,\tau} = \{0,1\}$ equals one if respondent j received the treatment $\tau = \{1,2\}$ or the placebo treatment $\tau = P$, and a zero value otherwise. In this regression, the omitted category is the passive control group. The element $\mathbf{X}_j \in \mathbb{R}^{1\times n_{\mathrm{x}}}$ denotes a set of relevant respondent-level covariates measured before treatment implementation, where $n_{\mathrm{x}} \geq 0$ is the number of covariates included in the regression. Since treatment was randomly assigned, the inclusion of $\mathbf{X}_j$ serves solely to improve the precision of the estimated treatment effects. Table A.9 presents the results of this regression.[31] As expected, the results for the coefficients $\gamma_{\tau,1}$ and $\gamma_{\tau,2}$ with $\tau = \{1,2\}$ indicate that both treatments significantly affect posterior beliefs, with heterogeneous effects depending on the respondent's prior. As Weber et al. (2025) explains, respondents' learning from treatment τ is best captured by the coefficient $\gamma_{\tau,2}$, which reflects the change in the slope of the relationship between prior and posterior beliefs for the corresponding treated group. The first two columns of this table indicate that treated respondents exhibit a significantly weaker relationship between their posterior and prior beliefs compared to the control group. As anticipated, respondents in either of these treatment groups incorporate the newly received information when forming posteriors, effectively adjusting their expectations to align with the values provided in the treatments. This set of results aligns with the empirical literature studying the formation process of inflation expectations using information treatment experiments (Cavallo et al., 2017; Coibion et al., 2022a; Hajdini et al., 2022). Additionally, given that the coefficients for the placebo treatment $\gamma_{P,1}$ and $\gamma_{P,2}$ are not statistically significant, we confirm that the experiment's placebo does not influence posterior beliefs, ruling out numerical anchoring bias in our experiment. Overall, these findings suggest that the provided treatments constitute a valid source of exogenous variation in inflation expectations, as both treatments led respondents to significantly revise their beliefs toward the provided signal.

Since our main objective is to analyze potential heterogeneous effects between HtM and non-HtM respondents, Figure 5 splits the relationship between the answers to the prior and posterior

[31] Due to the substantial noise in expectations and intended consumption survey data, we follow the literature (?Coibion et al., 2023b, 2024; Georgarakos et al., 2024) and employ Huber-robust regressions to systematically control for outliers and minimize the adverse effects of highly influential observations (Huber, 1964).

questions by treatment status and HtM condition. A visual inspection of this figure suggests that the slope of HtM respondents is indistinguishable from that of non-HtM respondents across all treatment groups, indicating no significant differences in belief updating behavior between these two groups. To formally assess this observation, we introduce the dichotomic variable $\mathrm{H}_j = \{0, 1\}$ in the regression specified in equation (12), where the value of this variable is one if we classify respondent j as HtM and zero otherwise. Therefore, we run the following regression

$$\begin{aligned} post_j \quad = \quad & \alpha + \eta \cdot prior_j + \eta^{\mathrm{H}} \left(prior_j \times H_j \right) + \mathbf{X}_j \boldsymbol{\beta}' + \sum_{\tau = \{1,2,\mathrm{P}\}} \Big[\gamma_{\tau,1} \cdot treat_{j,\tau} \\ & + \gamma^{\mathrm{H}}_{\tau,1} \left(treat_{j,\tau} \times H_j \right) + \gamma_{\tau,2} \left(treat_{j,\tau} \times prior_j \right) + \gamma^{\mathrm{H}}_{\tau,2} \left(treat_{j,\tau} \times prior_j \times H_j \right) \Big] + \upsilon_j \end{aligned} \tag{70}$$

The last four columns of Table A.9 present the results from this modified regression. As anticipated, there are no significant differences in how respondents incorporate the treatments across HtM status since the coefficients $\gamma^{\mathrm{H}}_{\tau,2}$ are small and not statistically different from zero across for every treatments τ. In particular, Figure B.4 illustrates these findings by plotting the estimated regression coefficients $\gamma_{\tau,2}$ and $\gamma^{\mathrm{H}}_{\tau,2}$. These results correspond to those presented in the fourth and fifth columns of Table A.9. The left panel of this figure focuses on the fourth column results and confirms that the placebo treatment did not influence non-HtM posterior beliefs, while both first treatments led non-HtM respondents to revise their beliefs significantly. Similarly, the right panel confirms that there are no significant differences in how HtM respondents incorporate each of the treatments compared to non-HtM respondents. HtM households tend to report a significantly lower average household income than non-HtM respondents. Although differences in educational attainment were less pronounced, some disparities still existed between the two groups. The literature suggests that lower-income and less educated respondents are systematically less informed than their more affluent and educated counterparts. Based on this, we anticipated a possible differential impact of information treatments on belief updating across HtM status, reflecting a possible differential impact of the provided information on respondents' idiosyncratic initial information set employed to form expectations. However, our findings do not support this hypothesis. This lack of heterogeneity in beliefs updating is key for our experiment. Since there are no statistical differences in how treatments affect posterior belief formation between HtM and non-HtM, and since this regression represents the first stage of our identification strategy, any heterogeneous effects observed in the second stage can be attributed to differential consumption responses to variations in inflation expectations, rather than to differences in treatment assimilation between the two types of respondents. For the remainder of the analysis, we adopt a more conservative approach and define the control group of our experiment as those respondents from either the initial passive control group or those who received the placebo treatment.

Appendix E General Equilibrium Framework

E.1 Firms

The supply side of the economy operates using a standard macroeconomic structure. We assume the existence of two distinct industries indexed by $j = \{d, c\}$, where $j = c$ and $j = d$ represent industries exclusively dedicated to the non-durable and durable production, respectively. We split the production process into two levels. At one level, a perfectly competitive final good producer aggregates intermediate varieties elaborated by monopolistically competitive firms operating in each sector to produce a homogeneous good. At the other level, there is a continuum of monopolistically competitive firms producing these intermediate good varieties. Both monopolistically competitive firms and varieties are indexed by k. We denote the sector j-specific set of intermediate goods-producing firms as $\mathcal{F}_j$ and assume it has a unit mass. Moreover, we assume that sector j's final good producer creates output $y_{j,t}$ by combining quantities $m_{k,t}$ of each one of the corresponding intermediate input varieties $k \in \mathcal{F}_j$ according to a standard constant elasticity of substitution aggregator function, such as

$$y_{j,t} = \left(\int_{k \in \mathcal{F}_j} (m_{k,t})^{\frac{\varepsilon_j - 1}{\varepsilon_j}} \, dk \right)^{\frac{\varepsilon_j}{\varepsilon_j - 1}} \tag{71}$$

where $\varepsilon_j > 1$ denotes the sector-specific elasticity of substitution across intermediate input varieties. The profit maximization problem of the perfectly competitive final good producer of sector j leads to the following isoelastic demand function for each intermediate input variety $k \in \mathcal{F}_j$

$$m_{k,t} = \left(\frac{P_{k,t}}{P_{j,t}} \right)^{-\varepsilon_j} y_{j,t} \tag{72}$$

where $P_{k,t}$ is the price of the intermediate good k, and $P_{j,t}$ is the standard sector j's Dixit-Stiglitz price aggregator. We assume that monopolistically competitive intermediate goods producers use a linear production function. They require labor amounts $n_{k,t}$ to produce the corresponding variety k according to the following technology

$$y_{k,t} \quad = \quad A_{j,t} n_{k,t} \tag{73}$$

where $A_{j,t}$ denotes the sector-specific total factor productivity, common to all the different intermediate producers operating in the industry j. We assume that the logarithm of this variable, denoted as $a_{j,t} = \log(A_{j,t})$, evolves according to an exogenous first-order autoregressive process, such as

$$a_{j,t} \quad = \quad \rho_{j,a} a_{j,t-1} + \sigma_{j,a} \upsilon_{j,t} \tag{74}$$

where $\rho_{j,a} \in [0,1]$ is the persistence coefficient of this process, and $\upsilon_{j,t}$ denotes a standard independent and identically normally distributed shock, i.e., $\upsilon_{j,t} \overset{\text{iid}}{\sim} \mathcal{N}(0,1)$. In consequence, the coefficient $\sigma_{j,a} \geq 0$ captures the standard deviation value of the innovation affecting the productivity process $a_{j,t}$. Since firms demand labor in exchange for the nominal wage $P_{w,t}$, in a given period t, the net profit function of the firm producing intermediate variety k corresponds to

$$P_{j,t} \Pi_{k,t} \quad = \quad (1 + \tau_k) P_{k,t} y_{k,t} - P_{w,t} n_{k,t} - \nu_{k,t} \tag{75}$$

where τ_k is a production subsidy and $\nu_{k,t}$ denotes a lump-sum profit tax. Following the literature (Debortoli and Galí, 2024), we assume that the government implements the standard New Keynesian optimal subsidy, given by $\tau_k = (\varepsilon_j - 1)^{-1} > 0$, to eliminate markup distortions in the fully flexible price limit of the economy. The total cost of this subsidy is fully financed through the total sectoral profit tax $\nu_{j,t} = \int_{k \in \mathcal{F}_j} \nu_{k,t} dk$, resulting in zero sectoral profits in the zero-inflation steady-state of the economy. As discussed in Bilbiie (2020), this assumption ensures an efficient fully flexible price limit, providing households with full insurance for both non-durable and durable consumption. As a consequence, the economy's steady state is not affected by profits' allocation and satisfies $(\bar{c}_i, \bar{d}_i) = (\bar{c}, \bar{d})$ for all household $i \in \mathcal{I}$. Although this assumption is not necessary to obtain our main results, it considerably simplifies the algebra. Moreover, under these conditions, sector j's aggregate real profits correspond to

$$\Pi_{j,t} = \int_{k\in\mathcal{F}_j} \Pi_{k,t} dk = y_{j,t} - \frac{P_{w,t} n_{j,t}}{P_{j,t}} \tag{76}$$

where $n_{j,t} = \int_{k\in\mathcal{F}_j} n_{k,t} dk$ corresponds to the aggregate sectoral labor demand. We incorporate nominal price rigidities following the pricing mechanism introduced by Calvo (1983). We assume that in every period, a firm producing any sector j's intermediate good has a probability $(1-\theta_j) \in (0,1)$ of being able to reset its price. Firms that receive this opportunity set their optimal price $P^*_{k,t}$ to maximize the expected stream of future net profits discounted over time. Formally, when setting prices, these firms solve the following optimization problem

$$P^*_{k,t} = \underset{\{P_{k,t}\}}{\arg\max} \sum_{h\geq t} \mathbb{E}_t \left[(\theta_j)^{h-t} \Lambda_{h|t} \left\{ (1+\tau_k) P_{k,t} y_{k,h|t} - P_{j,h} mc_{k,h} y_{k,h|t} - \nu_{k,h} \right\} \right] \tag{77}$$

where $mc_{k,t} = mc_{j,t} = P_{w,t} / (A_{j,t} P_{j,t})$ denotes the real marginal cost of firm producing variety k, and $\Lambda_{h|t}$ corresponds to the firm's relevant stochastic discount factor.[32] Since intermediate good producers operate in a monopolistically competitive environment, they maximize their future net profits, accounting for the isoelastic demand of final good producers. According to equation (72), for any period $h \geq t$, this demand function corresponds to

$$y_{k,h|t} = \left(\frac{P_{k,t}}{P_{j,h}} \right)^{-\varepsilon_j} y_{j,h} \tag{78}$$

As a solution, the optimal reset price for a firm producing the intermediate good k is

$$P^*_{k,t} = \left(\frac{\varepsilon_j}{(1+\tau_k)(\varepsilon_j - 1)} \right) \sum_{h\geq t} \mathbb{E}_t \left[\frac{(\beta\theta_j)^{h-t} (c_{\mathrm{s},h})^{-\sigma} (x_{\mathrm{s},h})^{\omega} (P_{j,h})^{1+\varepsilon_j} (P_{c,h})^{-1} mc_{k,h} y_{j,h}}{\sum_{\tau\geq t} \mathbb{E}_t \left[(\beta\theta_j)^{\tau-t} (c_{\mathrm{s},\tau})^{-\sigma} (x_{\mathrm{s},\tau})^{\omega} (P_{j,\tau})^{1+\varepsilon_j} (P_{c,\tau})^{-1} y_{j,\tau} \right]} \right] \tag{79}$$

[32] As discussed earlier, since we assume that non-HtM households own firms' shares, for any period $h \geq t$, the appropriate stochastic discount factor corresponds to $\Lambda_{h|t} = \beta^{h-t} (\lambda_{\mathrm{s},h} / \lambda_{\mathrm{s},t})$, where $\lambda_{i,t}$ denotes the Lagrange multiplier associated with the budget restriction of agent i.

Since final good producers operate in a perfectly competitive environment, final good prices correspond to the Dixit-Stiglitz price aggregator of all the industry's varieties. Because of the symmetry of the optimization problem faced by intermediate goods producers within each production sector, all these resetting price firms charge the same optimal price $p^*_{j,t}$. As a result, the price level of each final good $j = \{d, c\}$ evolves according to the following law of motion

$$p_{j,t} = \left[(1-\theta_j)\left(P^*_{j,t}\right)^{1-\varepsilon_j} + \theta_j \left(P_{j,t-1}\right)^{1-\varepsilon_j}\right]^{\frac{1}{1-\varepsilon_j}} \tag{80}$$

A first-order log-linear approximation of the standard sectoral price equation around a zero-inflation steady state yields the following approximation to the sector-specific New Keynesian Phillips Curve

$$\hat{\pi}_{j,t} = \beta \mathbb{E}_t \left[\hat{\pi}_{j,t+1}\right] + \mu \widehat{mc}_{j,t} \tag{81}$$

where $j = \{d, c\}$ and $\mu = (1-\theta_j)(1-\beta\theta_j)/\theta_j$ captures the sensitivity of sectoral inflation to real marginal costs.