The Welfare Cost of Inflation Risk*

Tomás E. Caravello Marc de la Barrera

MIT IESE Business School

Preliminary and incomplete – comments welcome!

February 2026

Abstract

Inflation shocks create redistribution risk and distort financial markets, potentially reducing welfare. We develop a heterogeneous-agent model with incomplete markets in which agents cannot issue real bonds, so inflation risk cannot be fully insured and is priced in equilibrium. The model generates a positive inflation risk premium even when inflation and output shocks are uncorrelated. We show that the inflation risk premium is informative about the welfare cost of inflation risk. In a calibration that matches key macroeconomic and asset-pricing moments, eliminating inflation risk yields welfare gains for currently alive agents equivalent to a permanent 0.4% increase in aggregate income. These gains are mainly driven by general equilibrium effects: eliminating inflation risk allows agents to borrow and lend more efficiently, improving the allocation of aggregate risk. Gains are smaller for future cohorts, highlighting an important intergenerational dimension.

Keywords: unexpected inflation, inflation risk premium, incomplete markets, heterogeneous agents, welfare cost of inflation.

JEL codes: E31, G12, D52

*We thank Ricardo J. Caballero, Eduardo Dàvila, Stavros Panageas, Christian K. Wolf, and participants of the BSE Summer forum for useful comments and suggestions. All errors are the sole responsibility of the authors.

1 Introduction

Inflation is widely regarded as an important source of macroeconomic risk. A substantial literature studies the costs of *trend* inflation through transaction frictions or price dispersion, but much less is known about the welfare effects of inflation *risk*. This paper asks a simple question: How socially costly is inflation risk, and how does it compare quantitatively to traditional sources of welfare costs of inflation?

To answer this question we develop a model with heterogeneous agents facing incomplete markets, in the form of no inflation indexation. Our analysis builds on a simple observation: inflation shocks redistribute wealth across agents with nominal positions (Doepke and M. Schneider, 2006), and these redistributive shocks cannot be fully traded away when markets are incomplete. Heterogeneity creates motives for financial trade, while the absence of effective inflation hedges exposes agents to it in equilibrium. As a result, inflation risk adds volatility to consumption, distorts borrowing and lending decisions, and is priced in equilibrium, making it a potentially important determinant of welfare.

We calibrate the model to jointly match standard asset pricing moments, as well as macroeconomic moments describing the dynamics of aggregate consumption growth and inflation. A salient feature of recent data is these two variables are essentially uncorrelated at all leads and lags. With a representative agent, this fact would imply a zero inflation risk premium. Nevertheless, our model generates a positive, empirically relevant inflation risk premium as an untargeted moment: Market incompleteness exposes agents to inflation risk even in the absence of aggregate comovement.

Using the model, we show that eliminating inflation risk generates welfare gains equivalent to a permanent increase in aggregate income of around 0.4% for initially alive agents. Most of these gains come via general equilibrium effects: inflation risk reduces lending in equilibrium. When it is eliminated, agents can better share the risk of output shocks by trading the existing assets. This estimate is at least as large as canonical measures based on price dispersion and transaction frictions, suggesting that inflation risk is a significant and previously overlooked source of inflation-induced welfare losses.

The magnitude of the welfare losses is disciplined by observable moments: the equity risk premium, the zero correlation between inflation and aggregate consumption growth, the positive inflation risk premium, and the stock of nominal debt. The equity premium indicates that risk is quantitatively important for welfare, while the lack of inflation–output comovement implies that inflation risk premia stems from market incompleteness. The level of the inflation risk premium and the size of outstanding nominal debt then reveal the severity of this incompleteness and discipline the magnitude of welfare losses from inflation risk. By contrast, complete-markets models would imply zero welfare costs, yet such models cannot match the joint macroeconomic and asset-pricing evidence. Conceptually, our results echo Alvarez and Jermann (2004):

matching asset pricing moments is key for welfare assessments. In our case, the joint set of macro and asset pricing moments reveal, through the lens of our model, that inflation risk is socially costly.

Insights in a two period model. We start our theoretical analysis by presenting a simple two period model. There are two risk averse agents with potentially heterogeneous preferences and/or endowments. The agents are endowed with some amount of goods at time 0, and stocks, that pay risky dividends at time 1. The model has two sources of risk: a shock to aggregate output, and a shock to the price level. We assume a cashless economy with flexible prices, and that both shocks are uncorrelated. Agents can trade in stocks and bonds. There are two classes of bonds: nominal and indexed (real) bonds. Both classes of bonds are in zero net supply.

The key assumption is that agents cannot borrow in real terms, so in equilibrium all borrowing occurs in nominal bonds, which are exposed to inflation risk. We see this assumption as reasonable based on three observations: first, indexed instruments are rare in practice. Of course, this may just reflect that agents see little value in indexing contracts, but it could also reflect some underlying cost, for example complexity costs (Gabaix, 2025) or institutional costs.[1] The lack of indexation is often considered a puzzle, see for example (Fischer, 1975). For our analysis, we take the stance that trade in indexed bonds is costly and thus does not happen. Second, we see the no indexation assumption as putting the proposed mechanism in equal footing with typical analyses of other sources of welfare costs of inflation, where some sort of cost to holding money or changing prices is assumed, and higher trend inflation makes these costs worse. [2] Third, as we explain below, our assumptions regarding the degree of market incompleteness and processes for output and inflation generate an empirically reasonable inflation risk premium and outstanding nominal debt as untargeted moments. Thus, both the price and the quantity of inflation risk in the model are in line with empirical evidence.

In this setup, we derive three main results. First, if there is trade in nominal bonds in equilibrium, then the inflation risk premium is positive. Second, we show that the equilibrium is Pareto inefficient: if we eliminated inflation risk, all agents can be better off.[3] Third, the level of inflation risk premium and the size of nominal positions are directly informative about the welfare gains for lenders of reducing inflation risk.

Quantitative Model. We then introduce a quantitative model. The basic structure is identical to Gârleanu and Panageas (2015) and Andrés Schneider (2022): there are two types of agents with heterogeneous

[1] Examples would include the tax treatment of paper gains due to indexation, as well as outright prohibition of indexing.

[2] For transactional costs, it is usually assumed that money gives utility by itself (Sidrauski, 1967), or that money is required to make transactions, but acquiring and holding money is costly (Baumol, 1952; Tobin, 1956). For costs arising from price and wage stickiness, it is usually assumed that either prices change infrequently for exogenous reasons (Taylor, 1980; Calvo, 1983; Woodford, 2003), or that changing prices or wages requires paying some costs (Rotemberg, 1982; Golosov and Lucas, 2007; Nakamura et al., 2018). Therefore the relevant friction is typically either assumed outright, or microfounded by a direct cost.

[3] However, we show that the equilibrium is constrained efficient.

risk aversion and EIS. There is an underlying OLG structure that makes the stationary wealth distribution non-degenerate. The main differences are that we include inflation risk and relax complete markets. Agents can trade in stocks, a nominal bond with positive duration, and an instantaneous (real) bond. As before, the key constraint is that agents cannot borrow in real terms. In this economy, the price level is locally deterministic, but inflation is stochastic. [4] Since traded nominal bonds have positive duration, shocks to inflation generate capital gains or losses, and therefore expose its holders to inflation risk. The assumed duration directly indexes the degree of market incompleteness, ranging from zero (instantaneous bonds, effectively real) to severe (long-duration bonds, highly exposed to inflation shocks). We calibrate this duration externally to roughly match the duration of outstanding debt in the market.

We use the remaining preference and technology parameters to jointly match: i) the processes for aggregate consumption growth and inflation, with uncorrelated shocks, ii) standard asset pricing moments, such as short term rates and the equity premium. Our model also generates reasonable values for the inflation risk premium and outstanding stock of nominal debt as untargeted moments.

The Welfare Cost of Inflation Risk. With the calibrated model at hand, we can now turn to the main question of the paper. We face the usual challenge of measuring aggregate welfare with heterogeneous agents. In order to do so, we adopt a metric that follows in spirit that of (Baqaee and A. T. Burstein, 2025). In particular, we first focus on the initially alive generation and compute the uniform increase of output (in all times and states) such that, if a planner were able to make initial transfers, it could leave all agents exactly indifferent between the economy with inflation risk and higher output, and the economy without inflation risk. Note that this measure does not rely in any interpersonal comparisons of utility. [5]

We estimate that the gains from eliminating inflation risk for initially alive agents are equivalent to a 0.4% permanent increase in aggregate consumption. This number is significant, and at least as large as typical estimates of costs stemming from price dispersion or transaction frictions. Surprisingly, this large estimate is mainly driven by compensation required for less risk averse agents. In equilibrium, these agents have debt positions in nominal bonds. The presence of inflation risk makes these debt positions more expensive, and therefore they take smaller positions. Once inflation risk is eliminated, debt is cheaper, and that allows these agents to lever up more and have higher expected returns. On the flip side, the equilibrium effects of inflation risk in agents with higher risk aversion is mixed: They bear more risk, but they are getting a positive compensation for it. However, for initially alive agents, it is generally the case that eliminating inflation is beneficial for both agents. For example, evaluating the loss at the mode of the wealth distribution

[4] That is, the price level follows $\frac{dP}{P} = \pi_t dt$, and π is an Ito process.

[5] See the discussion in (Baqaee and A. T. Burstein, 2025), and also (Dávila and Schaab, 2025; Barcons, Dávila, and Schaab, 2026)

yields gains of 7.2% for less risk averse agents, and 0.5% for more risk averse agents.

We also show that these gains are partially driven by the initial transitional dynamics: eliminating inflation risk creates better investment opportunities, which the initially old generation capitalizes. We then compute the required compensation for unborn cohorts. Since cohorts born after the initial date do not benefit from this, and thus estimated welfare gains are lower for them. Overall, our exercise shows that this often overlooked source of welfare costs of inflation can be comparable to other traditional sources, though the gains are concentrated among currently alive agents and are partly offset by inter-generational redistribution.

Related Literature. This paper relates to several strands of literature in macroeconomics and finance. First, the question of the welfare costs associated with inflation is decades old and has given rise to a vast literature. See for example a detailed description of different sources in Fisher and Modigliani (1978). There are two main strands in this literature. The first branch has to do with the effects of inflation on the cost of holding money, which is necessary to transact. Earlier studies include Bailey (1956) and Friedman (1969). Although there is disagreement on its exact magnitude, the estimated costs through this channel are substantial. For example, Lucas (2000) finds that reducing inflation from 10% to 0% is equivalent to a gain of 1% in real income. More recent work finds smaller, although still substantial costs: (Benati and Nicolini, 2025) present a benchmark estimate of 0.35%.

An important related paper within the transaction frictions literature is (**allais2020welfare**): to the best of our knowledge, this is the only other paper that focuses on the effects of inflation *risk*. In contrast to us, they focus on an economy with idiosyncratic risk and money in the utility function. In that economy, the effects of inflation risk operate through different channels to ours: inflation risk generates volatility in real money holdings, which hurts welfare since utility is concave in money. In contrast, we focus on a cashless economy, and the welfare costs of inflation risk come through imperfect risk-sharing of aggregate shocks, as well as the stock market effects of inflation risk.

The second main strand of the literature on welfare costs of inflation has to do with New Keynesian costs of inflation, which are related to its effects on price dispersion, markups and costs of changing prices. The exact way in which pricing frictions are modeled is important for the welfare costs of inflation (A. Burstein and Wellwig, 2008; Nakamura et al., 2018): Calvo models tend to imply large costs, whereas menu cost models imply negligible costs. (Cavallo, Lippi, and Miyahara, 2023) calibrate a flexible pricing model that matches the distribution of price changes, and find that the total cost of increasing inflation from 0 to 10% is around 0.4 of GDP. Focusing on the cost of wage adjustment, (Guerreiro et al., 2025) find that workers are willing to give up around 1.75% of their wages to avoid renegotiation costs of readjusting their wage. Our

contribution is to focus on a different, so far essentially unexplored cost of inflation: that associated with its volatility, not its trend. We show that this can also lead to welfare losses, and our quantification shows that these can be as substantial as those from traditional sources.

While our paper tackles a classical macro question, the models and mechanisms we consider are deeply rooted in the macro-finance literature of asset pricing models with heterogeneous agents (He and Krishnamurthy, 2013; Brunnermeier and Sannikov, 2014; Gârleanu and Panageas, 2015). In particular, our model extends that of Gârleanu and Panageas (2015) to include inflation risk and incomplete markets. Our paper is most closely related to Andrés Schneider (2022) and Andres Schneider (2025). Relative to those papers, we feature incomplete markets for inflation risk, and we focus on the implied welfare cost. Both of those papers include an analysis of nominal yields, but inflation does not affect the equilibrium allocations. Instead, it affects the prices of nominal bonds, which are not traded by the agents in the model. In contrast, in our model inflation risk does affect real allocations and thus it can generate welfare losses.

Also in the intersection between macro and finance, we relate to Alvarez and Jermann (2004). That paper uses asset prices to directly measure welfare costs of business cycles with a sufficient statistics approach. Our paper focuses on the costs of inflation risk. It is also more conventional in that we uses asset pricing moments as targets to calibrate and validate a particular model, and then evaluate welfare under counterfactual scenarios.

While not our main focus, our paper also connects with the literature on sources of inflation risk premium. A first strand of the literature focuses on flexible price models, and calibrates the correlation between aggregate consumption and inflation. An important reference in this literature is Piazzesi and M. Schneider (2006). It considers a model with recursive preferences in which positive inflation shocks correlate with bad news about future consumption growth. Wachter (2006) consider an estimated habit model. Inflation risk-premium is positive because innovations to inflation are negatively correlated with innovations to aggregate consumption growth. Another branch of literature considers New Keynesian models with asset prices and inflation (Rudebusch and Swanson, 2012; John Y. Campbell, Pflueger, and Luis M. Viceira, 2020; Pflueger, 2025; Cieslak and Pflueger, 2023; Caballero and Simsek, 2023). These models endogenize the correlation between output and inflation as stemming from demand and cost-push shocks. Our main contribution to this literature is to focus in a mechanism that can generate positive inflation risk premium even if consumption and inflation are uncorrelated at all leads and lags, which we argue is the empirically relevant case for recent data. We also show that, given our mechanism, inflation risk premium is informative about the potential welfare losses stemming from inflation volatility. These losses would be absent in a representative agent framework.

Finally, we relate to a recent literature on how to perform welfare evaluations with heterogeneous agents,

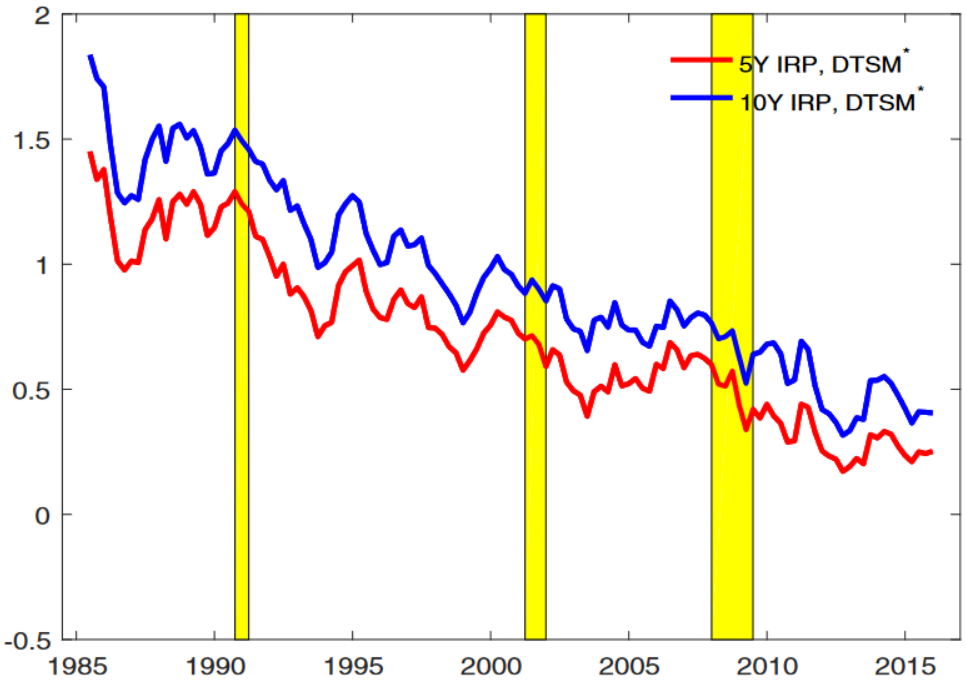


Figure 1: Inflation risk-premium for 5 and 10 year bonds derived from the model estimated by Ajello, Benzoni, and Chyruk (2019).

without having to take a stance on interpersonal utility comparisons (Baqaee and A. T. Burstein, 2025; Dávila and Schaab, 2025; Barcons, Dávila, and Schaab, 2026). Our proposed efficiency metric is quite similar to that of Baqaee and A. T. Burstein, 2025. Given that our question is about aggregate welfare and our mechanism requires heterogeneity, we see our paper as a natural application of the tools developed in this literature.

Roadmap. Section 2 briefly reviews some motivating empirical facts. Section 3 explains the main insights in a two period model. Section 4 introduces the quantitative model and its calibration. Section 5.1 describes the positive properties of the model, in terms of matching asset pricing moments. Section 5.2 presents the estimated welfare losses from inflation risk. Section 6 concludes.

2 Motivating Evidence

In this section, we present motivating empirical evidence on several key facts that our model matches, and that are important to generate welfare costs from inflation risk.

2.1 Inflation risk-premium is positive, even in recent times

Estimating the inflation risk premium is hard. Earlier approaches relied on comparing TIPS and nominal yields. A pioneering contribution is John Y Campbell and Luis M Viceira (2001), how fit a two factor term structure model. They estimate inflation risk premia on the range of 35 to 100 basis points depending on the maturity. In most cases, an affine gaussian term structure model is proposed. Different data and modeling choices have been implemented. Haubrich, Pennacchi, and Ritchken (2012) fit a four factor model with stochastic volatility using data on nominal yields, surveys, realized inflation and swaps. They find a risk premium of around 40bps at a 30-year horizon. Ang, Bekaert, and Wei, 2008 fit a three factor model with regime switching. Abrahams et al. (2016) fit a six factor model, that explicitly accounts for the relative liquidity of TIPS. They find lower inflation risk premium, in the order of 10 basis points for a 10 years maturity. Among the papers that do not specify an affine model for the SDF, Fleckenstein, Longstaff, and Lustig (2017) uses data on swaps and options to fit a model on physical and risk-neutral inflation. They find the inflation risk premium for a 10 year maturity to be in the order of 20 basis points. Fang, Liu, and Roussanov (2025) uses the returns of several assets to compute the price of the inflation risk, as well as inflation betas, for both core and energy components. They find positive inflation risk-premium for core inflation even in the post-2000 subsample.[6]

Figure 1 is taken from Ajello, Benzoni, and Chyruk (2019) and depicts the inflation risk-premium at 5 and 10 year horizons in an estimated arbitrage-free affine model. Their model is estimated in the 1985Q1-2015Q4 sample and thus omits most of the large inflation episodes in the 70s and 80s. We can observe that, although declining throughout the sample, inflation risk-premium is consistently positive.

2.2 Inflation and aggregate consumption are uncorrelated

We compute the sample cross-correlogram between aggregate consumption growth and inflation, i.e $Corr(\Delta \log(c_t), \pi_{t-\ell})$ for $-20 \leq \ell \leq 20$. For the baseline results, we use real Personal Consumption Expenditures and CPI inflation. The time period used are quarters.

Figure 2 presents the results for two subsamples: 1960Q1-1984Q4 (left) and 1985Q1-2019Q4 (right). We observe that for the earlier subsample, there is a significantly negative correlation between consumption and inflation at short lags and leads. Based on this sample, one would conclude that inflation is indeed bad news in terms of aggregate consumption, and thus inflation risk premium may stem from covariance with aggregate consumption.

[6] In that paper, they report their finding as "only core inflation carries a negative risk-premium". This is purely linguistic: in their paper, this implies that an asset that pays more when inflation is high (i.e an inflation hedge) has lower expected return. Since our paper relates to the bond pricing literature, we refer to this as "positive inflation risk-premium", since this refer to the relation between expected returns of nominal and real bonds.

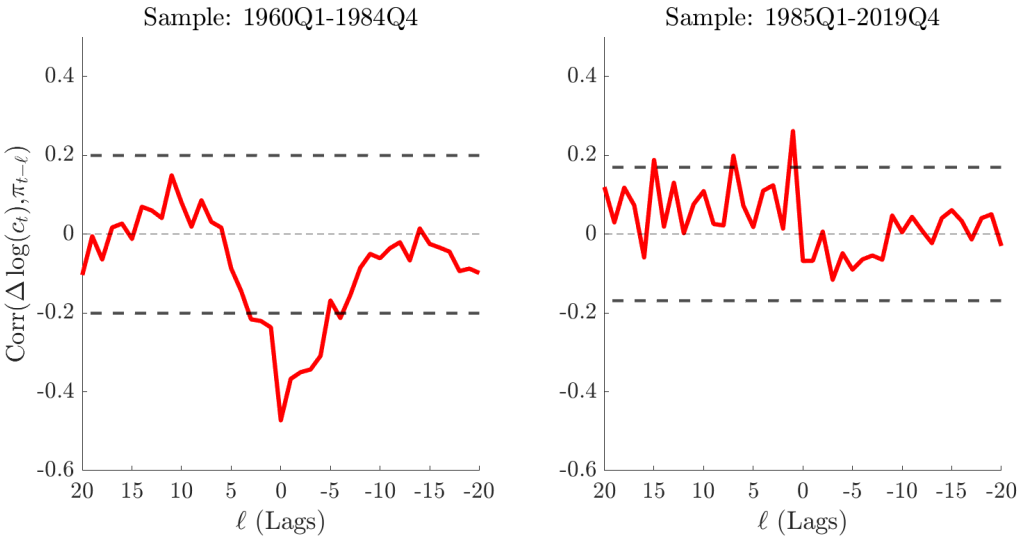


Figure 2: Cross-correlogram between PCE real consumption growth and CPI inflation for two different subsamples. Dashed lines are 95% confidence intervals.

This negative co-movement is no longer observed when we focus on the more recent sample, in the right panel of Figure 2. if anything, inflation is (mildly) positively correlated with consumption after 1985. This pattern is robust to using core inflation, and to including the Covid-19 period, see Appendix A.

In this world, a consumption-based representative agent would imply an essentially zero inflation risk-premium, since inflation is uncorrelated with aggregate consumption both contemporaneously (which would generate zero inflation risk-premium with time-separable preferences) and at lags and leads (which would generate zero risk premium in a model with habits or recursive preferences respectively). This is consistent with the widely held view that weakly inflationary demand shocks are the primary source of fluctuations in aggregate consumption in recent times.

2.3 Heterogeneous exposure to inflation risk

In order to get a sense on the degree of heterogeneity on net nominal positions, we use the Survey of Consumer Finances (SCF) which is a sample of 4,595 households representative of the US economy provided by the FED. Using this data we can observe nominal and real positions for different households. The exercise is similar to Doepke and M. Schneider (2006), which uses the same dataset but in 1989 and 2001. We show the results for the 2022 vintage.

In the SCF we can observe wealth and income components of households. We are interested in the distinction between nominal assets and real assets. We consider real assets are all the non-financial assets plus stocks, equity mutual funds. The remaining assets (mostly liquid assets and bonds and pension funds)

Income percentile Age group	[0,20)	[20,40)	[40,60)	[60,80)	[80,90)	[90,100]
<35	-0.30	-0.92	-0.58	-0.68	-0.89	0.05
35-44	-0.08	-0.79	-0.60	-0.54	-0.29	0.05
45-54	-0.12	-0.08	-0.10	-0.24	-0.12	0.11
55-64	0.09	0.02	-0.00	0.10	0.24	0.23
65-74	0.09	0.31	0.34	0.33	0.41	0.29
>=75	0.14	0.26	0.19	0.38	0.47	0.29

Panel A: Net nominal exposure, defined as nominal assets minus nominal debt over net worth.

Income percentile Age group	[0,20)	[20,40)	[40,60)	[60,80)	[80,90)	[90,100]
<35	1.30	1.92	1.58	1.68	1.89	0.95
35-44	1.08	1.79	1.60	1.54	1.29	0.95
45-54	1.12	1.08	1.10	1.24	1.12	0.89
55-64	0.91	0.98	1.00	0.90	0.76	0.77
65-74	0.91	0.69	0.66	0.67	0.59	0.71
>=75	0.86	0.74	0.81	0.62	0.53	0.71

Panel B: Real assets share, defined as total real assets over net worth.

Income percentile Age group	[0,20)	[20,40)	[40,60)	[60,80)	[80,90)	[90,100]
<35	0.00	0.00	0.01	0.01	0.00	0.01
35-44	0.00	0.00	0.00	0.01	0.01	0.05
45-54	0.00	0.00	0.01	0.01	0.02	0.09
55-64	0.01	0.01	0.01	0.03	0.02	0.17
65-74	0.01	0.02	0.02	0.03	0.04	0.16
>=75	0.01	0.02	0.02	0.02	0.02	0.11

Panel C: Net worth of each group as a fraction of the total net-worth of the economy.

Table 1: Source: Net nominal exposure, share of real assets and net worth for households sorted by age and income percentiles.

are considered nominal. Given the virtual lack of privately issued indexed debt, all debt is considered to be nominal.

In terms of aggregates, non-financial assets represent 65% of total assets, and 50% of those are primary housing and vehicles. Another 34% of the non-financial assets accrue to business, and the remaining are real state investments. For the 35% of financial assets, its three main categories of interest are nominal savings (29%) which includes cash, certificates of deposits and bonds, pension funds (31%) and equities (22.68%), held directly and through mutual funds or exchange-traded funds.

Focusing on heterogeneity, we split households by the age of the household head and the income percentile

they belong as in Doepke and M. Schneider (2006). We compute exposure to inflation as nominal assets minus nominal debt divided by total net worth, defined as the sum of total nominal assets and real assets minus nominal debt. We show the results on Panel A of Table 1. We observe substantial heterogeneity across age: young households tend to have negative nominal exposure, whereas older agents tend to have positive exposure. Young households have negative net nominal asset positions due to homeownership and vehicle ownership, which are bought using nominal debt. Those households take nominal debt to buy a real asset. On the other side, older, richer households have positive net exposure since they have fully paid their mortgages and have savings in terms of pension funds and bonds.

A second and related fact is that the share in real assets of those different household groups differ. The label of real assets includes both fiscal assets as houses and cars, as well as financial assets that are claims on those real assets, such as equity. Panel B shows how younger households have a higher real asset share, because they are levered on housing, while older and richer households have a smaller share of their net wealth in real assets. Overall, this shows that the mirror image of heterogeneity of nominal positions, in a world where debt is mainly nominal, is heterogeneity in real asset positions. Finally, Panel C shows the distribution of net worth across these two dimensions. As we can see, richer and older households hold most of the wealth.

2.4 Takeways

First, inflation risk-premium is positive even for recent sample, where cost-push shocks are thought to be relatively small. Consistent with this view, aggregate consumption growth and inflation are essentially uncorrelated at all leads and lags.

Second, there is large heterogeneity in nominal exposure across households. As documented by Doepke and M. Schneider (2006), age and income are two important factors that drive this heterogeneity, with younger, poorer households having negative nominal exposure and older, richer households having positive exposure. The flip side of that is heterogeneity in real asset positions: younger, poorer households are levered and thus have more than 100% of their net wealth in real assets, whereas older, richer households tend to have less nominal debt and nominal assets as a share of their net wealth. In what follows, we consider models that match these facts jointly, and used them to quantify the welfare cost of inflation risk.

3 Key insights in a simple model

3.1 Economic Environment

There are two time periods, $t = 0, 1$. Assume agents choose from three assets: a tree, a nominal bond, and a real bond. Their prices are denoted by $q_s, q_b^{\$}, q_b$ respectively, and are denominated in goods at time 0. The tree is in unit net supply while the other assets are in zero net supply. The tree has pay-off $y = \bar{y} + \sigma_y \varepsilon_y$ and the nominal bond has real pay-off $\Pi^{-1} - \sigma_\pi \varepsilon_\pi$, whereas the real bond pays 1 unit of final good with certainty.[7] Define $R^{\$} = \frac{\Pi^{-1} - \sigma_\pi \varepsilon_\pi}{q_b^{\$}}$ and $R = \frac{1}{q_b}$ being the realized real returns on the nominal and real bonds respectively. Assume ε_π and ε_y are mean zero and independent. This implies that aggregate consumption and inflation are uncorrelated. Agents face a borrowing constraint in the real bond.

Agents' Problem. Agents have time-separable, expected utility preferences over consumption at 0 and lotteries of consumption at $t = 1$, with instantaneous utility function u_i, which is increasing and concave for both agents. Agents are endowed with an initial amount of goods at time 0, y_{i0}, and a share of the tree, ω_i. Assume there are two agents $i = A, B$ with heterogeneous preferences and endowments. The optimization problem of agent i is:

$$\begin{aligned} \max_{c_{0i}, a_i, b_i, b_i^{\$}} \; & u_i(c_{0i}) + \mathbb{E}[u_i(c_{1i})] \\ \text{s.t. } & c_{0i} + q_s(a_i - \omega_i) + q_b b_i + q_b^{\$} b_i^{\$} = y_{0i} \\ & c_{1i} = a_i(\bar{y} + \sigma_y \varepsilon_y) + b_i^{\$}(\Pi^{-1} - \sigma_\pi \varepsilon_\pi) + b_i \\ & b_i \geq 0. \end{aligned}$$

Replacing the expression for c_{1i} and forming a Lagrangian we get the following optimality conditions:

$$\begin{aligned} \mathbb{E}[u_i'(c_{i1})(\Pi^{-1} - \sigma_\pi \varepsilon_\pi)] &= \lambda_i q_b^{\$} \\ \mathbb{E}[u_i'(c_{i1})(\bar{y} + \sigma_y \varepsilon_y)] &= \lambda_i q_s \\ u'(c_{0i}) &= \lambda_i \\ \mathbb{E}[u_i'(c_{i1})] &\geq \lambda_i q_b \quad (\text{with} = \text{ if } b_i > 0). \end{aligned}$$

The solution to this system implies consumption and asset demands as a function of prices $(q_s, q_b, q_b^{\$})$ and wealth $W_i = y_{0i} + q_s s_i$.

[7] We get this pay-off for the nominal bond if we assume that the (net) inflation is π, and choose the distribution of π such that $\frac{1}{1+\pi} = \bar{\Pi}^{-1} - \sigma_\pi \varepsilon_\pi$

Equilibrium Definition. A competitive equilibrium is an allocation $(c_{0i}, a_i, b_i, b_i^{\$})_{i=A,B}$ and a price vector $(q_s, q_b, q_b^{\$})$ such that:

1. Given the initial endowments and prices, the allocation solves the optimization problem for both agents.
2. Markets clear:

$$c_{0A} + c_{0B} = y_0, \quad c_{1A}(s_1) + c_{1B}(s_1) = y_1(s_1) \quad \forall s_1,$$
$$a_A + a_B = 1, \quad b_A + b_B = 0, \quad b_A^{\$} + b_B^{\$} = 0,$$

where $s_1 = (\varepsilon_\pi, \varepsilon_y)$ is the aggregate state at time 1.

Before proceeding with the characterization of the equilibrium, a few points are worth discussing.

No Inflation Indexation. The key assumption that delivers our results is the impossibility of borrowing in inflation indexed terms. As discussed in the introduction, we have three main justifications for this assumption: i) inflation indexed debt is very rare in reality, ii) other sources of inflation-related welfare losses also require assuming some cost that is made worse with inflation, iii) in Section 4, we externally calibrate the degree of this market incompleteness, and our model generates empirically reasonable inflation risk premium and outstanding private debt.

Cashless Limit. We stress that in this model money serves no transactional role: it is only a unit of account in which some contracts (nominal bonds) are denominated. Furthermore, the price level is exogenous and driven by shocks ε_π. Such an environment can be microfounded by adding money in the utility function, taking the cashless limit, and specifying monetary policy appropriately, see Appendix B.1. Thus, for interpretation purposes, we think of the ε_π shocks as coming from monetary shocks.

Equilibrium Selection. Given that there is a borrowing constraint in real terms, this can lead to equilibrium multiplicity: for any equilibrium price vector, if we increase the price of q_b (i.e lower the real rate), then both agents would like to borrow more in real terms. Given that this is not possible, the new vector could also be an equilibrium. In this case, following Werning (2015) we select the equilibrium with the highest real rate, which is the one that would survive if we allowed agents to borrow small amounts of indexed debt.

Note that in the selected equilibrium, the optimality condition of real bonds will hold with equality for one of the agents. That is, given the equilibrium prices, one of the agents is indifferent to consume a little less and save a little more in real bonds. From now on we refer to this agent as the *pricer*, given that it is a marginal investor in all three assets and thus "prices" them.

3.2 Positive Inflation Risk Premium with no Output-Inflation Correlation

We show below that, if heterogeneity is such that agents hold non-zero nominal positions in equilibrium, then the inflation risk premium is positive, i.e $\mathbb{E}[R^\$ - R] > 0$.

Proposition 1. *Assume that primitives are such that, in equilibrium,* $|b_A^\$| > 0$. *Then* $\mathbb{E}[R^\$ - R] > 0$.

Proof. In equilibrium, the optimality conditions of the "pricer" imply:

$$\mathbb{E}[u_i'(c_{i1})\left(R^\$ - R\right)] = 0$$
$$\mathbb{E}[R^\$ - R] = -Cov\left(u_i'(c_{i1}), R^\$ - R\right)/\mathbb{E}[u_i'(c_{i1})]$$

so the inflation risk premium crucially depends on $Cov\left(u_i'(c_{i1}), R^\$ - R\right)$. For inflation risk premium to be positive, we need this covariance to be negative, which obtains whenever c_{i1} has negative co-movement with inflation. In turn, this requires that the "pricer" is long in nominal bonds, $b^\$ > 0$.

We proceed by contradiction. Assume that the real bond is priced by the borrower (i.e the agent with $b_i^\$ < 0$). In that case, c_{i1} is increasing in inflation, so $Cov\left(u_i'(c_{i1}), R^\$ - R\right) > 0$ and the inflation risk premium would be negative. However, for this to be an equilibrium, it must be the case that the other agent is optimizing as well. If $R > \mathbb{E}[R^\$]$, the other agent (which has $b_{-i}^\$ > 0$) would find optimal to switch all his savings to real bonds: she gets a higher expected return and lower risk. In order to see that such portfolio would indeed be better, note that the original proposed allocation (where the agent invests in nominal bonds) implies consumption $c_{1i} = a_i y + b_i^\$(\Pi^{-1} - \sigma_\pi \varepsilon_\pi)$. The new allocation, where the agent switches to save in real, implies consumption $\tilde{c}_{1i} = a_i y + \frac{q_b^\$}{q^b} b_i^\$$. Since $R > \mathbb{E}[R^\$]$, this implies $\frac{q_b^\$}{q^b} > \mathbb{E}[\Pi^{-1} - \sigma_\pi \varepsilon_\pi] = \Pi^{-1}$. Thus:

$$\mathbb{E}[u_i(c_{1i})] < \mathbb{E}[u_i(a_i y + b_i^\$ \Pi^{-1})]$$
$$< \mathbb{E}[u_i(a_i y + \frac{q_b^\$}{q^b} b_i^\$)] = \mathbb{E}[u_i(\tilde{c}_{1i})],$$

where the first line follows from the fact that c_{1i} is a mean preserving spread of $a_i y + b_i^\$ \Pi^{-1}$ (since ε_π is independent of y), and the second line follows from the fact that $a_i y + \frac{q_b^\$}{q^b} b_i^\$$ first order stochastically dominates $a_i y + b_i^\$ \Pi^{-1}$. Thus, since the agent has increasing marginal utility and is risk averse, this agent prefers to lend in real, and this this cannot be an equilibrium. The above contradiction shows that in any equilibrium with trade in nominal bonds, we must have that the agent who has $b^\$ > 0$ prices the real bond and thus $\mathbb{E}[R^\$ - R] > 0$. □

In summary, once we consider an equilibrium where there is non-zero trade in nominal bonds, the inflation risk premium is determined by the agent whose consumption declines with inflation, and thus inflation risk

premium positive. However, the argument in the proof is more general that the current set-up. Appendix B.2 shows that a version of Proposition 1 holds in several extensions of the model, such as adding idiosyncratic income shocks, recursive preferences, allowing inflation to be correlated with output, and allowing non-traded income to be correlated with inflation. In all cases, the main result remains: the inflation risk premium is determined by an agent whose consumption falls with inflation, and thus the risk premium is positive. In section 4, we generate non-zero nominal positions by considering agents with heterogeneous risk aversion and elasticity of intertemporal substitution, along with a constraint on indexed debt issuance.

3.3 Pareto Gains and Constrained Efficiency

Having derived an important positive property of the equilibrium, we now turn to welfare properties. Our first lemma shows that if there is trade in nominal bonds, the equilibrium is trivially Pareto inefficient.

Lemma 1. *If $|b_i^{\$}| > 0$, then the equilibrium is Pareto inefficient.*

Proof. Construct a Pareto improvement as $c_{i,0}^{PI} = c_{i,0}$ and $c_{i,1}^{PI} = \mathbb{E}[c_{i,1}|y]$.□

Intuitively, an almighty planner can take the same allocations but we eliminated the inflation risk (i.e the planner gives both agents the equilibrium expected return of nominal bonds but without the risk). Both agents improve, so inflation risk trivially generates Pareto inefficiencies.

On the other hand, we also may ask if the equilibrium is constrained efficient, that is if a planner could improve upon the market allocation by choosing a different set of portfolios subject to the constraints faced by agents. In particular, assume the planner can choose among any consumption plans in $t = 1$ that can be spanned as $c_{1i} = ay + b^{\$}(\Pi^{-1} - \sigma_\pi \varepsilon_\pi)$, for any a, b satisfying asset market clearing. The proposition below shows that the equilibrium is constrained efficient.

Proposition 2. *The equilibrium is constrained efficient.*

Proof. The proof follows from deriving the first order conditions of the planner's problem, and showing that they are satisfied by the equilibrium allocation. See Appendix B.3 □.

Note however that this is not a general property of these models: as shown by (Geanakoplos and Polemarchakis, 1986), if the economy had more than one good in period 2, the equilibrium with incomplete asset markets is generally constrained inefficient. For this reason, in the present paper we abstain from optimal policy (since by assumption the model has no role for it) and focus only on measuring the welfare costs of inflation risk.

3.4 Inflation Risk Premium and Welfare

Finally, we show how inflation risk premium is related to the welfare effects of changing inflation risk. We consider a marginal reduction in σ_π maintaining the equilibrium portfolios. We do so in order to separate the direct effects of changes in inflation risk from further changes due to agents reoptimizing and general equilibrium forces. More concretely, recall that in this economy the consumption allocation at time $t = 1$ is:

$$c_{i1} = a_i(\bar{y} + \sigma_y \varepsilon_y) + b_i^{\$}(\Pi^{-1} - \sigma_\pi \varepsilon_\pi) + b_i.$$

Thus, taking portfolio allocations as fixed, we can get the marginal welfare gain (in terms of $t = 0$ goods) of a change in inflation volatility as:

$$\frac{\partial \mathbb{E}[u_i(c_{i1})]}{\partial \sigma_\pi} / (u_i'(c_0)) = \frac{\mathbb{E}[u_i'(c_{i1})(-\varepsilon_\pi)]}{u'(c_0)} b_i^{\$}.$$

An interpretation of this magnitude is as follows: if at $t = 0$ we let the agent to pay a cost $\bar{c}(\sigma_\pi^* - \sigma_\pi)$ to reduce σ_π to σ_π^*, $\frac{\partial \mathbb{E}[u_i(c_{i1})]}{\partial \sigma_\pi} / (u_i'(c_0))$ is the value of $\bar{c}$ that makes the agent choose $\sigma_\pi^* = \sigma_\pi$. In order words, this is the willingness to pay to marginally reduce inflation risk at the equilibrium allocation, in terms of time 0 goods.

Since $\mathbb{E}[u_i'(c_{i1})(\varepsilon_\pi)] > (<) 0$ for the lender (borrower), the above partial derivative is negative for both agents: a reduction in σ_π increases welfare for both agents, which means that both agents are willing to pay positive amounts to reduce inflation risk. This shows that, if we marginally reduced inflation risk, there is potential for Pareto gains.

In order to quantify these potential gains, from the FOCs of the "pricer" imply:

$$\begin{aligned} \mathbb{E}[u_i'(c_{i1})(\Pi^{-1} - \sigma_\pi \varepsilon_\pi)] &= \lambda_i q_b^{\$} \\ \frac{\mathbb{E}[u_i'(c_{i1})(-\varepsilon_\pi)]}{u'(c_0)} b_i^{\$} &= -q_b \frac{E[R^{\$}] - R}{\sigma_\pi / q_b^{\$}} b_i^{\$}. \end{aligned}$$

This formula shows that, fixing the asset allocation and focusing on the unconstrained agent, we can evaluate the effects in welfare of a marginal reduction in inflation risk by multiplying the price of the inflation risk, $\frac{E[R^{\$}] - R}{\sigma_\pi / q_b^{\$}}$, times the exposure to inflation risk, $b_i^{\$}$ and appropriately discounted over time by q_b. This is intuitive: under the proposed experiment, marginally lowering σ_π reduces the inflation risk faced by the agent by $b_i^{\$}$, and thus multiplying this number by the price of the inflation risk we obtain the willingness to pay to reduce this risk. This shows that, in this model, inflation risk premium is informative about the size of potential welfare gains from reducing inflation risk.

We emphasize that this exercise does not directly show what would happen in equilibrium if we reduced σ_π since it ignores agents re-optimizing and general equilibrium forces. Furthermore, the model considered so far is highly stylized. Thus, in order to show that the preceding mechanism can quantitatively match the data and to derive the implied welfare costs of inflation risk, the next section presents a quantitative model.

4 Quantitative Model

In this section, we present a quantitative version of the model. The model is similar to Gârleanu and Panageas (2015), with a crucial difference: as in the previous section, markets are incomplete and agents must trade using assets that are exposed to nominal risk.

4.1 Economic Environment

Exogenous Processes. Time is continuous and indexed by $t \in [0, +\infty)$. There is a single consumption good. Aggregate output is denoted by y_t. A fraction $(1-\omega)$ of aggregate output is labor income. We assume this goes directly to hand-to-mouth agents, that immediately consume it. [8] The rest is dividend income, produced by trees which are owned by investors. Following Panageas (2020), there is a continuum of trees that pay dividends in units of the consumption good. Letting $s \leq t$ denote the time of creation of a tree, we denote its time-t dividends by $y_{t,s} = \omega\delta_q e^{-\delta_q(t-s)} y_t$, where y_t is aggregate output, which follows a geometric Brownian motion

$$\frac{dy_t}{y_t} = \mu_y dt + \sigma_y dW_{y,t}, \tag{1}$$

where $W_y = \{W_{y,t} \in \mathbb{R}; \mathcal{F}_t, t \geq 0\}$ is an aggregate Brownian motion in the probability space $(\Omega, \mathbb{P}, \mathcal{F})$ representing permanent shocks to aggregate output. The price level p_t satisfies $\frac{dp_t}{p_t} = \pi_t dt$ where π_t is inflation, which follows a Cox, Ingersoll, and Ross (1985) process (also used in Andrés Schneider (2022)):

$$d\pi_t = \lambda_\pi(\bar{\pi} - \pi_t) + \tilde{\sigma}_\pi \sqrt{\pi_t - \pi_{\min}} dW_{\pi,t} \tag{2}$$

where $W_\pi = \{W_{\pi,t} \in \mathbb{R}; \mathcal{F}_t, t \geq 0\}$ is an aggregate Brownian motion in the probability space $(\Omega, \mathbb{P}, \mathcal{F})$ representing shocks to the inflation rate. The parameters $(\lambda_\pi, \bar{\pi}, \pi_{\min})$ are real numbers and are associated with the persistence, mean, and lower bound on inflation. Importantly, the exogenous process π_t is stationary. The volatility term implies that inflation risk increases with the level of inflation.

It will be useful to write the model in matrix notation, so we let $dW_t = [dW_{y,t}, dW_{\pi,t}]^T$ the vector of

[8] These agents play no role in the equilibrium. We introduce this in order to and additional parameter (ω) that allows us to match traded wealth to GDP.

Brownian shocks. Importantly, $\text{corr}(dW_{y,t}, dW_{\pi,t}) = 0$. This assumption differentiates our model from the rest of the literature, (e.g Piazzesi and M. Schneider (2006) or Andrés Schneider (2022)) which obtain a positive inflation risk premium from the assumed negative correlation between inflation and output.

Assets and returns. There are effectively three assets in the market: short term bonds, long-term bonds, and trees. Since the price level is locally deterministic, short term bonds (even if denominated in nominal terms) are risk-free. However, just like in the model in Section 3, agents cannot borrow using this bond. Given that this bonds are in zero net supply, this implies that in equilibrium neither agent can use this bond. As we explain below, the main purpose of including this asset is to have a well defined real risk-free rate.

Agents can also invest in long-term, nominal bonds. A bond issued at time t pays $P_t \delta_b e^{-\delta_b(s-t)}$ at time s, where P_t is the price level and δ_b is a measure of the duration of the bond. Taking $\delta_b \to \infty$ yields an instantaneous bond, and $\delta_b \to 0$ corresponds to a consol. The advantage of modeling long term bonds this way is that one bond issued at t is equivalent to $\frac{P_t}{P_{t+s}} e^{-\delta_b(t-s)}$ bonds issued at time $s > t$, so we can summarize everything in terms of the equivalent units of time t bonds, and refer to this as "bonds" moving forward. The instantaneous return of holding a bond is

$$\frac{\delta_b}{q_t^b} dt + \frac{dq_t^b}{q_t^b} - (\delta_b + \pi_t) dt = \mu_{t,b} dt + \sigma_{b,t}^T dW_t. \tag{3}$$

In particular, investors receive the interest payment δ_b, bond prices change by dq_t^b/q_t^b, they mature at rate δ_b, and since they are nominal its value gets eroded at rate π_t.

Since the bonds are long term and nominal, they are subject to inflation risk: a shock that increases inflation reduces the real value of all remaining coupons. The longer the duration (lower δ_b), the larger the inflation risk bonds are exposed to. The constraint in short-term debt issuance forces investors to save and borrow assets subject to inflation risk. It also shows that, once we model price dynamics as we are, inflation risk premium is intrinsically linked to the duration of the assets that are actually traded. This also implies that the degree of market incompleteness in this economy is directly tied to assumed duration on the traded nominal bond, a point to which we come back below when calibrating the model.

Finally, investors can also trade trees. Note that holding one unit of a tree created at time s has exactly the same pay-offs as holding $e^{\delta_q(s'-s)}$ units of a tree created at time s'. Given this equivalence, we can summarize the full portfolio of trees investors have by converting everything to the equivalent number of time t trees an investor owns at time t. From now on refer to this as trees, and treat it as a single asset.

Letting q_t be the price at time t of a tree, the instantaneous return of holding a tree is

$$\omega\delta_q \frac{y_t}{q_t}dt + \frac{dq_t}{q_t} - \delta_q dt = \mu_{q,t}dt + \sigma_{q,t}^T dW_t, \tag{4}$$

for $\mu_{q,t}$, $\sigma_{q,t}$ and $H_{q,t}$ to be determined in equilibrium. Owning a tree yields $\omega\delta_q y_t$ in dividends, its price changes by dq_t/q_t, and it depreciates at rate δ_q. The expected return of a tree is $\mu_{q,t}$, and $\sigma_{q,t}^T = [\sigma_{q,t}^y, \sigma_{q,t}^\pi]$ is how much returns are affected by shocks to output and inflation. For notational simplicity, we write returns in matrix form

$$\begin{bmatrix} \frac{\omega\delta_q y_t}{q_t}dt + \frac{dq_t}{q_t} - \delta_q dt \\ \frac{\delta_b}{q_t^b}dt + \frac{dq_t^b}{q_t^b} - \delta_b dt - \pi_t dt \end{bmatrix} = \mu_t dt + \Sigma_t dW. \tag{5}$$

Agents and Demographics. There are two types of investors with recursive preferences as in Duffie and Epstein (1992). Agents' utility is

$$\mathcal{U}_t^i = \mathbb{E}_t\left[\int_t^\infty f^i(c_s^i, \mathcal{U}_s^i)ds\right], \tag{6}$$

where

$$f^i(c,\mathcal{U}) = \frac{1-\gamma}{1-1/\psi^i}\mathcal{U}\left[\left(\frac{c}{((1-\gamma^i)\mathcal{U})^{\frac{1}{1-\gamma^i}}}\right)^{1-\frac{1}{\psi^i}} - (\rho+\varphi)\right] \quad .$$

Risk aversion is denoted by γ^i and ψ^i is the elasticity of intertemporal substitution. The time discount rate is ρ and agents die at rate φ. Throughout the paper we maintain the convention that $0 \leq \gamma_A \leq \gamma_B$. We refer of type-$A$ agents as risk tolerant and type-B agents as risk averse.

We introduce deaths and births as a way to generate a stationary wealth share across agents (Gârleanu and Panageas, 2015). At rate φ agents are born, and with probability $\bar{\alpha}$ they are of type A and with probability $1-\bar{\alpha}$ they are of type B. Each newborn is endowed with $\frac{\delta_q}{\varphi}$ new trees, so in aggregate, new trees are created at rate δ_q. Since existing trees depreciate at rate δ_q, the mass of trees is constant and equal to 1. Agents can insure against the risk of dying in actuarily fair annuity markets. In particular, an agent with wealth n_t^i gets paid $\varphi n_t dt$ if they stay alive for an extra instant, but if they die they give up all their wealth. Given this, the wealth evolution of agent i conditional on not dying is:

$$\frac{dn_t^i}{n_t^i} = [r_t dt + (\theta_t^i)^T((\mu_t - r_t)dt + \Sigma_t dW)] - \frac{c_t^i}{n_t^i}dt + \varphi dt, \tag{7}$$

where $(\theta_t^i)^T = [\theta_{q,t}^i, \theta_{b,t}^i]$ is the portfolio allocation of the agent in risky assets. The first term captures the (stochastic) return on the agent's portfolio, the second term captures consumption, and the last term is the return on annuities. Investors choose consumption c_t^i and portfolios θ^i to maximize (6) subject to (7) and the no short-term borrowing constraint, $\theta_{q,t}^i + \theta_{b,t}^i \leq 1$. In Appendix C.1 we show the HJB of the agent and

its optimality conditions as a function of their wealth $n^i_{t,s}$ and the aggregate state of the economy x_t, which we describe below.

4.2 Recursive Formulation

The model admits a recursive representation. There are two aggregate states in this economy, one which is exogenous, the inflation rate π_t, and one endogenous, the wealth share of type-A investors, α_t, defined as:

$$\alpha_t = \frac{n_{A,t}}{n_{A,t} + n_{B,t}}.$$

Let $x = (\alpha, \pi) \in [0,1] \times [\pi_{min}, +\infty)$ be the state vector. Write its law of motion in matrix form

$$dx = \begin{bmatrix} d\alpha \\ d\pi \end{bmatrix} = \mu_x(x)dt + \Sigma_x(x)dW.$$

In Appendix C.2, we derive the law of motion of $d\alpha_t$ as a function of market returns and portfolio choices of the agents.

Homotheticity and Scaling. An important feature of our formulation is that homotheticity allows to write the value function as

$$\mathcal{U}^i(n,x) = \frac{n^{1-\gamma^i}}{1-\gamma^i} V^i(x)^{\frac{1-\gamma^i}{1-\psi^i}}. \tag{8}$$

Furthermore, all agents of the same type take the same portfolio decisions and consumption to wealth ratio:

$$c^i_t(n,x) = h^i(x)n$$
$$\theta^i(n,x) = \theta^i(x).$$

This allows us to aggregate consumption and assets demands within each group. For example, consumption of agents type A satisfies $\bar{\alpha}\int_{-\infty}^{t} \varphi e^{-\varphi(t-s)} c^A_{t,s} ds = n^A_t h^A(x)$, where n^A_t is total wealth of type A.

Furthermore, given that output shocks dW_y are permanent, we can rescale the prices of trees by y_t as: $q_t = y_t q(x)$, where $q(x)$ is the aggregate price-dividend ratio. Having defined the aggregate states, value, and policy functions, we are ready to define the recursive equilibrium.

Equilibrium definition. A recursive competitive equilibrium is a set of value functions $\{V^i(x)\}_{i=A,B}$, policy functions $\{h^i(x), \theta^i(x)\}_{i=A,B}$ and prices $\{q(x), q_b(x), r(x)\}$ such that:

1. Taking prices as given, V^i, h^i, and θ^i satisfy the optimality conditions and HJB of agent $i = A, B$.

2. For all x, goods and asset markets clear:

$$\alpha\theta_q^A(x) + (1-\alpha)\theta_q^B(x) = 1 \qquad \text{(Trees)}$$

$$\alpha\theta_b^A(x) + (1-\alpha)\theta_b^B(x) = 0 \qquad \text{(Bonds)}$$

$$(\alpha h^A(x) + (1-\alpha)h^B(x))q(x) = 1. \qquad \text{(Goods)}$$

Just like in the model of Section 3, due to the borrowing constraint there is a continuum of real rates that are possible in equilibrium. Again, we select the equilibrium with the highest one, which corresponds to the one in which the optimality condition of real bonds for one of the agents holds with equality.

Numerical Solution We solve the model globally, using finite differences to approximate the necessary derivatives. We define a grid for α and π, and guess the following equilibrium objects $V^i(x), \Sigma(x), q_b(x)$. Given these, we can use the rest of the equilibrium equations to get other equilibrium objects, as well as 7 equations that must be satisfied at the equilibrium. Thus for each point in the state space we have 7 unknowns ($V^i(x)$ for both agents, all 4 elements of $\Sigma(x)$ and $q_b(x)$). We are left with a system of $n_\alpha \times n_\pi \times 7$ unknowns to be solved globally using Newton method. Appendix C.6 provides more details of the implementation of the numerical solution.

4.3 Pricing Zero Coupon Bonds and the Inflation Risk Premium

In order to compare our model with the yield curves and inflation risk premium observed in the data, we need to price zero coupon real and nominal bonds. In order to do so, note that once we found the equilibrium value and policy functions, we can compute the stochastic discount factor of each agent. Since markets are incomplete, in general these will differ across agents. Following Schröder and Skiadas (1999), the evolution of the SDF of agent i is given by:

$$\frac{dm_t^i}{m_t^i} = f_{\mathcal{U}}^i + \frac{df_c^i}{f_c^i} = -r^i(x)dt - (\kappa^i(x))^T dW_t, \tag{9}$$

where $f_c^i(c,\mathcal{U}) = (\rho+\varphi)\left((1-\gamma^i)\mathcal{U}\right)^{\frac{1/\psi^i-\gamma^i}{1-\gamma^i}} c^{-\frac{1}{\psi^i}}$ is the marginal utility of consumption of agent i. In the Appendix we provide the formulas to compute $r^i(x), \kappa^i(x)$ directly from the solution to the equilibrium.

Both SDFs are valid to price trees and long term nominal bonds because these are freely traded in equilibrium. For pricing all other assets, we use the SDF of agent who has $\theta_b > 0$. In our calibration this is agent B, which has higher risk aversion. Note that in the selected equilibrium, this is a valid SDF, in the sense that it can price all available assets, including the short term bond.

Equipped with the stochastic discount factor and the law of motion of the state we can obtain $Q^{\$}(x,T)$ by solving the following PDE:

$$\begin{aligned} -Q^{\$}_T(x,T) + \mathcal{L}[Q(x,T)] - (r(x)+\pi)Q^{\$}(x,T) - (\nabla Q^{\$})^T \Sigma_x(x)\kappa(x) = 0 \\ Q^{\$}(x;0) = 1. \end{aligned} \tag{10}$$

Here $Q^{\$}_T$ represents the derivative of the bond with respect to the maturity, $\mathcal{L}[f] = (\nabla f)\mu_x + \frac{1}{2}Tr\left(\Sigma_x \Sigma_x^T (\nabla^2 f)\right)$ is the infinitesimal generator where $\nabla Q^{\$T} \equiv [Q^{\$}_\alpha, Q^{\$}_\pi]$ is the gradient of the bond prices with respect to the state and similarly $\nabla^2 Q^{\$}$ is the Hessian. As usual, yields and prices are related as $y^{\$} = -\frac{1}{T}\log(Q^{\$}(x,T))$. Similarly, we can find the price of a *real* zero coupon bond by solving:

$$\begin{aligned} -Q_T(x,T) + \mathcal{L}[Q(x,T)] - r(x)Q(x,T) - (\nabla Q)^T \Sigma_x(x)\kappa(x) = 0 \\ Q(x;0) = 1. \end{aligned} \tag{11}$$

In order to compute term and inflation risk premia, we follow Haubrich, Pennacchi, and Ritchken (2012). In particular, we create counterfactual zero risk premia nominal and real bond prices by solving equations (10)- (11) with $\kappa(x) = 0$ for all x. Denoting these prices by $Q^{\$,rf}(x,T)$ and $Q^{rf}(x,T)$ and their corresponding yields by $y^{\$,rf}, y^{rf}$. Then we compute the real and nominal term premia, and inflation risk-premia as:

$$RTP(x,T) = y(x,T) - y^{rf}(x,T \tag{12}$$

$$NTP(x,T) = y^{\$}(x,T) - y^{\$,rf}(x,T) \tag{13}$$

$$IRP(x,T) = NTP(x,T) - RTP(x,T). \tag{14}$$

4.4 Calibration

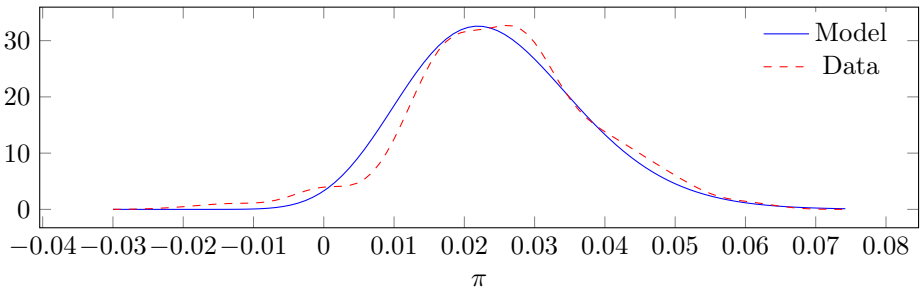


Figure 3: Stationary density of inflation, model vs. Data (sample: 1985 to 2019).

We calibrate the model at the annual frequency. The parameters are shown in Table 2. Given that our results depend crucially on the sources inflation risk premium, we do not target it in the calibration. Instead,

we ensure our model matches other asset pricing moments like the real rate and the equity risk premium. We then add an externally estimated process for inflation, and evaluate the resulting inflation risk premium as an untargeted moment.

We use Gârleanu and Panageas (2015) specification of the output process, with an annual growth rate of 2%, and $\sigma = 0.041$ so that time-integrated data from our model can roughly reproduce the first two moments of annual consumption growth. An instantaneous volatility of $\sigma = 0.041$ corresponds to a volatility of 0.033 for model-implied, time-integrated, yearly consumption data (John Y. Campbell and Cochrane, 1999). We set ω is set to 0.15. This generates a (traded) wealth-to-GDP ratio of around 4 on average, which is roughly consistent with what is observed in the data.[9]

For inflation, we estimate a Cox, Ingersoll, and Ross (1985) (CIR) model by maximum likelihood, with data from 1985 to 2019. We set $\pi_{\min}$ to -2.5% and estimate the rest of the parameters. Figure 3 shows the stationary distribution of inflation and compare it with the empirical one. The overall fit is good. The empirical distribution is somewhat more right-skewed than the empirical one, and it also has fatter tails. We note that the persistence of inflation is low, $\lambda_\pi = 0.546$, which means that the half life of the shock is around five quarters. Thus, our empirical process for inflation captures fluctuations in inflation at business cycle frequencies, and not longer term regime changes, such as changes in the inflation target. The value of σ_π implies that the volatility of inflation at $\bar{\pi}$ is 1.4%.

Preferences are calibrated to match the risk free rate and the equity risk premium as in Gârleanu and Panageas (2015). As in that paper, having recursive preferences with heterogeneity in both risk aversion and the EIS is crucial to match the asset pricing moments. In particular, heterogeneous risk aversion is required to have heterogeneous portfolios, and agent B has a very low EIS. The capital depreciation δ_q is chosen to match the real rate. [10]

As explained above, the duration of nominal bonds is the key parameter that regulates the severity of market incompleteness, and therefore the potential magnitude of inflation risk premium due to the proposed channel. With this in mind, long-term bonds are calibrated to have an average maturity of 5 years. This is in line with durations used for government debt in the macro literature (see Auclert (2020) and Bianchi, Faccini, and Melosi (2023)). Less evidence is available for the average duration of private sector debt. Since in the US most of outstanding household debt is mortgages, the actual duration of private sector debt is likely higher, so we see our number as conservative.

[9] For example, Auclert (2020) calibrate a wealth to GDP ratio of 3.82. For reference, the ratio between household net worth (FRED: TNWBSHNO) and GDP fluctuates around 3.8 and 4.8 between 1985 to date.

[10] In Gârleanu and Panageas (2015), agents face a decreasing path of income that incentivizes saving for retirement and thus puts downward pressure on the real rate. Solving this model in an incomplete markets setting like ours would be unfeasible and this is why we resort to capital depreciation to lower the real rate. See Panageas (2020) for more details.

Parameter	Description	Value
	Inflation Process	
λ_π	Persistence of inflation	0.546
$\bar{\pi}$	Inflation target	0.027
π_{min}	Minimum inflation	-0.03
σ_π	Volatility	0.062
	Output Process	
μ	Output growth rate	0.02
σ	Volatility of output	0.041
ω	Dividend share	0.150
	Agents	
ρ	Discount rate	0.001
φ	Death rate	0.02
$\bar{\alpha}$	Share of type-A investors	0.01
γ^A	Risk aversion of type-A agent	0.9
ψ^A	EIS of type-A agent	1.1
γ^B	Risk aversion of type-B agent	10
ψ^B	EIS of type-B agent	0.030
	Assets	
δ_q	Capital depreciation	0.03
δ_b	Bond duration	0.2

Table 2: Calibrated Parameters. See Section 4.4 for details.

5 Quantiative Results

5.1 Asset Pricing implications

Moment	Baseline	No Inflation Risk	Data
Real Rate	1.13%	1.24%	0.94%
Real Rate Volatility	1.09%	0.62%	1.22%
Equity Risk Premium	5.08%	4.98%	4.64%
Equity Volatility	20.09%	20.67%	12.21%
IRP 10-yr	0.17%	0.00%	0.60%

Table 3: Model moments vs. Data. “Baseline" indicates the model solution under the calibration in Table **??**. “No Inflation Risk" uses the same calibration except for σ_π, which is set to zero. Regarding data sources, the real rate is taken from FRED as the difference between the 1 year treasury yield (DTB1YR) and expected inflation (EXPINF1YR). Equity risk premium is from Martin (2017), the return volatility is taken from Schiller dataset and IRP is from Ajello, Benzoni, and Chyruk (2019). Except for the equity risk premium and IRP which are taken directly from other papers, the sample is 1985-2019.

Having described the model and its calibration, we turn to analyze its implications. In this section we focus on implied asset pricing moments, whereas the next section explores the welfare cost of inflation risk.

Table 3 shows several implied asset pricing moments on average and compares them to the data. Figure

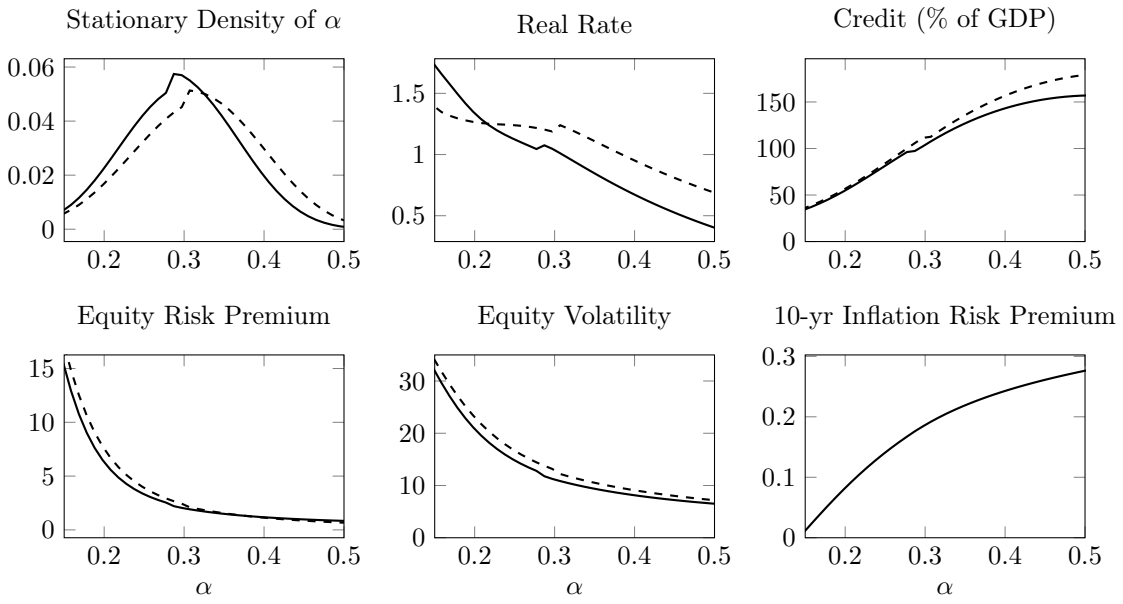


Figure 4: Selected equilibrium variables as a function of α, averaging across values of π using the stationary distribution. Panels "Real Rate", "Equity Risk Premium" and "10-yr inflation risk premium" are reported in percentage points. "Credit" reports the stock of bonds traded as a share of GDP. "Equity volatility" reports σ_q. Solid Line: baseline calibration. Dashed Line: no inflation risk ($\sigma_\pi = 0$).

4 reports those moments across the distribution of α, averaging across all inflation states. We comment on the main points below.

Standard Asset Pricing Moments. As we can see in Table 3, the model roughly matches the average level of the real rate. It also produces a sizable equity risk premium, which is broadly in line with the data. The model is also able to produce quite high equity volatility, which in fact is too high relative to the data.[11]

The mechanism is the same as Gârleanu and Panageas (2015): heterogeneous agents generate a time-varying distribution of wealth. Risk-tolerant agents are levered and when a bad output shock hits, it redistributes towards risk averse agents (α goes down). The particular combination of risk aversion and EIS chosen makes it so that changes in real rates and equity premium tend to reinforce each other, at least for a significant range of α. This generates significant volatility in price-dividend ratios, which in turn increases the equity premium.

Figure 4 shows this is in more detail. The equilibrium value of risk premium distribution is highly non-linear in α. In normal times (when α is close to its mode), the equity risk premium is low and around 3% a year. However, when bad shocks hit and the wealth share of risk-tolerant investors decreases, the equity risk premium increases sharply to values around 10% or higher. Intuitively, leverage of the risk-tolerant

[11] Note however that we focus on a recent subsample. Gârleanu and Panageas, 2015 use the full sample and obtain a volatility of returns of 18.2%, much closer to ours.

agents increases as its wealth share decreases, and thus the volatility of the stock market also increases as α decreases. This behavior is consistent with Martin (2017), who documents large spikes in equity risk premium using options data.

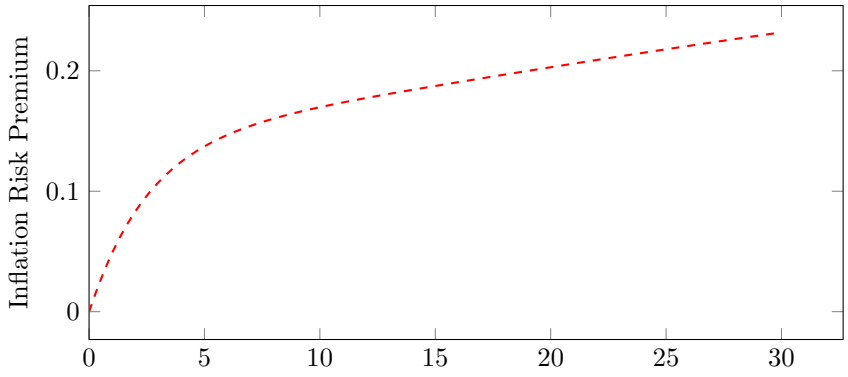
Nominal Debt Outstanding. We refer to the nominal debt outstanding as a fraction of GDP as Credit. This is computed as Credit $= (1-\alpha)\theta_B(\alpha,\pi)\tilde{q}(\alpha,\pi)$. The top right panel in Figure 4 depicts it as a function of α. As we can see, credit varies significantly with α. Intuitively, the closer wealth is an even split, there is more room for both agents to trade, and thus the total volume of credit increases. On the flip side, if wealth is very concentrated in one type of agents, those agents have effectively no counterpart to trade with, and thus equilibrium credit decreases.

In quantitative terms, the average Credit to GDP ratio under our calibration is around 100% of GDP. Defining the empirical counterpart of this is tricky, since our model is deliberately stylized. To get a sense of the magnitudes, we use debt statistics from the Flow of Funds for 2025Q3. In particular, we take private sector debt to be the sum of household and non-financial business debt. The final figure is around 137% of GDP. [12] Thus, our model generates a stock of debt that is roughly consistent with the data.

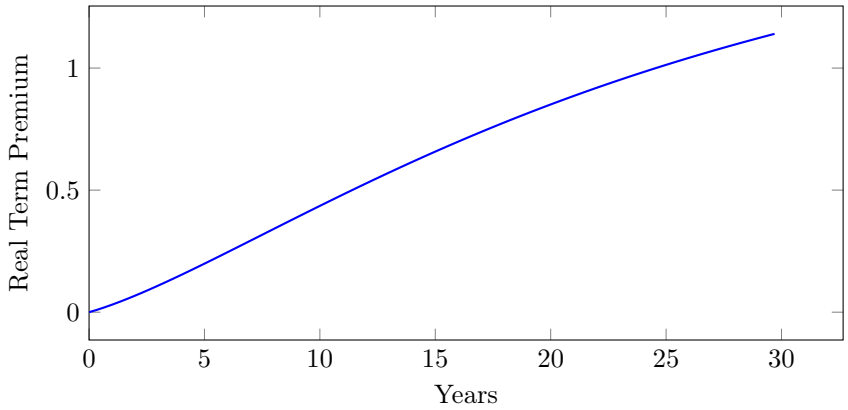
Inflation Risk Premium. Relative to Gârleanu and Panageas (2015), the novelty of our model is possibility of generating positive inflation risk premium out of uncorrelated inflation and output shocks. The basic mechanism is simple: since agents have different risk aversions, they want to have different exposure to stocks. In equilibrium, this requires that they trade in the nominal bonds market. Then the intuition of Section 3 applies directly: since there is non-zero exposure to this risk in equilibrium, the inflation risk premium is positive. We can also see that inflation risk-premium increases with α. The intuition is that, as α grows, there is more trade in the bond market (see panel "Credit" depicted above). This makes agents more exposed to inflation risk, and thus inflation risk-premium increases.

The average values are shown in the bottom part of Table 3 and in the bottom right panel in Figure 4. As we can see, our model is able to produce a sizable inflation risk-premium, about half of what is found in empirical studies, and close to the values found in the pre-Covid years. For reference, the estimated inflation risk premium is about 40bps in Haubrich, Pennacchi, and Ritchken (2012), and it gets to a similar value at the end of the sample in Ajello, Benzoni, and Chyruk (2019). This is remarkable given that our calibration left inflation risk premium as an untargeted moment. Therefore, our analysis shows that a non-trivial part of the observed inflation risk premium may come from our proposed mechanism.

[12] According to the above source, total household debt in the US was 20.7 trillion dollars and non-financial business debt was 22.1 trillion. Given that US nominal GDP stood around 31 trillion, household debt stood at around 66% of GDP, whereas non-financial business debt was 71%, and thus total debt is around 137% of GDP.



(a) Inflation Risk Premium



(b) Real Term Premium

Figure 5: Term Structure of Risk Premia.

Term Structure Implications. Figure 5 depicts the term structure for inflation and real term premium, averaging across all states with the stationary distribution. As we can see, inflation risk premium is zero at low maturities, and it's term structure is upward sloping. It reaches 18 bp at 10 years and around 24 bps at 30 years. Overall, the upward-sloping shape of the term structure of inflation risk premium is consistent with empirical studies (Haubrich, Pennacchi, and Ritchken, 2012; Ajello, Benzoni, and Chyruk, 2019), although our model generates an inflation risk-premium that is somewhat lower than the generated by those models.

Regarding the real term premium, which is positive, upward sloping, and larger than the inflation risk premium. The mechanism here is similar to the one stressed in Andrés Schneider (2022): in these models, negative output shocks reduce α and increase real rates. This makes real bonds fall when output falls, so they command a positive risk premium. However, the size of the real term premium is about half of what is typically found in empirical studies, such as Figure 7 in Haubrich, Pennacchi, and Ritchken (2012) or Figure 3 in Andrés Schneider (2022).

Overall, the model predicts upward sloping curves for inflation and term premium, which are consistent

with the data. However, the prices of these risks in our model are somewhat smaller (about half) of what is found in the literature, yet not irrelevant by any means. Given that higher risk premia would likely increase our estimated welfare costs (Alvarez and Jermann, 2004), we proceed with this calibration and interpret our results as a lower bound.

Effects of inflation risk. The dashed line in Figure 4 shows what happens if we shut down inflation risk in the model, i.e if we set $\sigma_\pi = 0$. In that case, the stationary distribution of α shifts to the right, and the level of credit increases, specially for high values of α. The idea is that, as inflation risk falls, borrow becomes more attractive for type-A agents. They use it to take larger positions in trees, which increases the drift in their wealth and thus alters the stationary distribution.

Regarding asset prices as a function of α, the real rate increases somewhat, indicating higher desire to borrow from type-A agents. The equity risk premium does not change much. However, given that the stationary distribution shifted to the right, the economy with σ_π has a lower average equity premium.

Overall, eliminating inflation risk makes borrowing more appealing. This increases credit and makes type-A agents lever up more. In equilibrium, this increases the wealth share of type-A agents, and since more trees are held by these type of agents, the risk premium decreases.

5.2 The Welfare Costs of Inflation Risk

We next turn our attention to the welfare costs of inflation risk. Right off the bat, the presence of relevant heterogeneity makes welfare evaluations tricky. This problem is made even worse by the presence of overlapping generations of agents (Barcons, Dávila, and Schaab, 2026). Below we present efficiency measures that do not require to take a stance on interpersonal welfare comparisons, at least for some subsets of agents.

5.2.1 Initially Alive Agents

First, we focus on an efficiency measure that takes into account the welfare effects of eliminating inflation risk in the initially alive agents. Consider an economy where the volatility of inflation is zero, while all other parameters are the same as in Table 2. We want to compare this economy with an economy with inflation risk, where σ_π is set at the value calibrate above. Our first measure of the welfare cost of inflation risk is by how much aggregate output needs to increase so that, we appropriate transfers, a planner can to make both agents indifferent between an economy with inflation risk and higher output, or an economy with no inflation risk.

In particular, assume there is a fictitious planner in the economy with inflation, that can create trees, and also redistribute trees at will among initially alive agents *before* any trade happens. In other words,

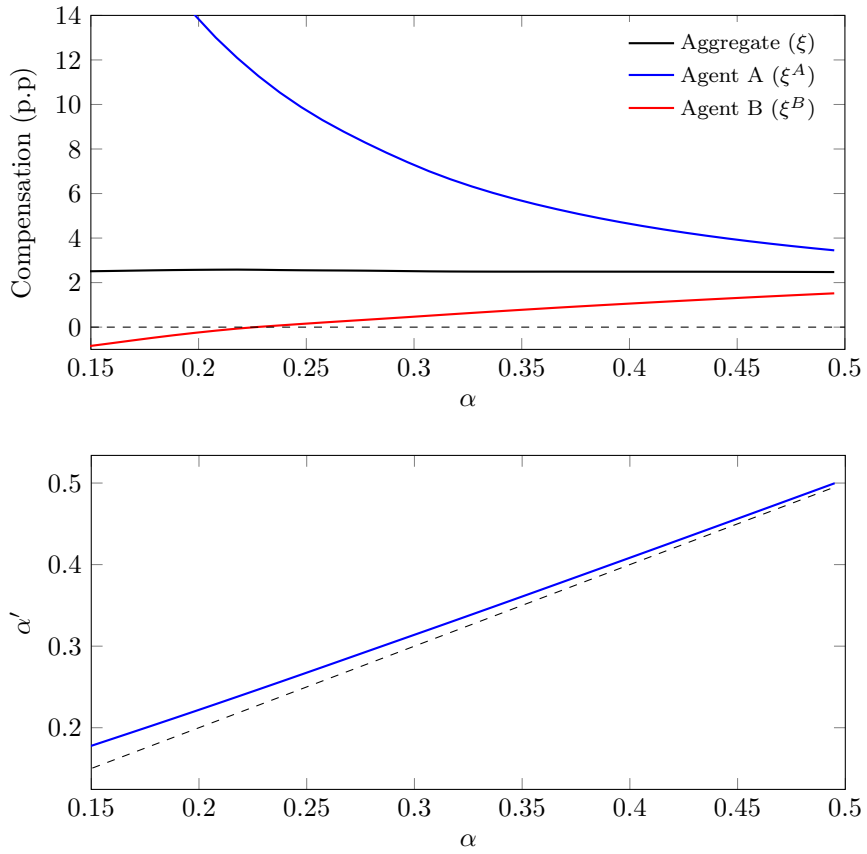


Figure 6: Top: Compensation required to leave agents indifferent, in percentage points of *traded* income. The aggregate number is computed as $\xi = \alpha\xi^A + (1-\alpha)\xi^B$ following the discussion in the text. Bottom: α', i.e implied value of α after initial redistribution takes place.

at $t = 0$ given an initial state $x_0 = (\alpha_0, \pi_0)$, and before any trade happens, the planner can i) create $\xi(x_0)$ trees, ii) allocate $\eta^A(x_0)$ trees to type-A investors and $\eta^B(x_0)$ trees to type-B investors. Note that there are also hand-to-mouth agents, but these are not affected by our exercise: they consumption is not tied in any way to inflation. We thus focus on investors to compute the gains in terms of traded income, but then scale by ω, the fraction of income that is traded, to get everything in terms of aggregate output. The resource constraint of the planner is

$$\xi(x_0) = \bar{\alpha}\eta^A(x_0) + (1-\bar{\alpha})\eta^B(x_0). \tag{15}$$

Let $\mathcal{U}_{0_}^{i,\sigma_\pi=0}(n,x)$ be the lifetime utility of an initially alive agent (encoded in the subscript $0_$) of type i in an economy without inflation risk (σ_π), with wealth n when the aggregate state is x. Given that there are $\bar{\alpha}$ agents of type A and $1-\bar{\alpha}$ agents of type B, if the initial wealth share is α when each individual type i agent has wealth $n^i = \chi_i(\alpha_0) y_0 \tilde{q}^{\sigma_\pi=0}(x_0)$ where $\chi_A(\alpha) = \frac{\alpha_0}{\bar{\alpha}}$ and $\chi_B(\alpha) = \frac{1-\alpha_0}{1-\bar{\alpha}}$.[13]

[13]Note that this follows from the fact that at time 0 all agents are endowed with trees so the aggregate wealth in the economy is $y_0\tilde{q}^{\sigma_\pi=0}(x_0)$. Furthermore, if the total wealth share of agents type A is α_0 but there are only $\bar{\alpha}$ such agents, then each agent has a $\chi_A(\alpha_0)$ slice of the total wealth.

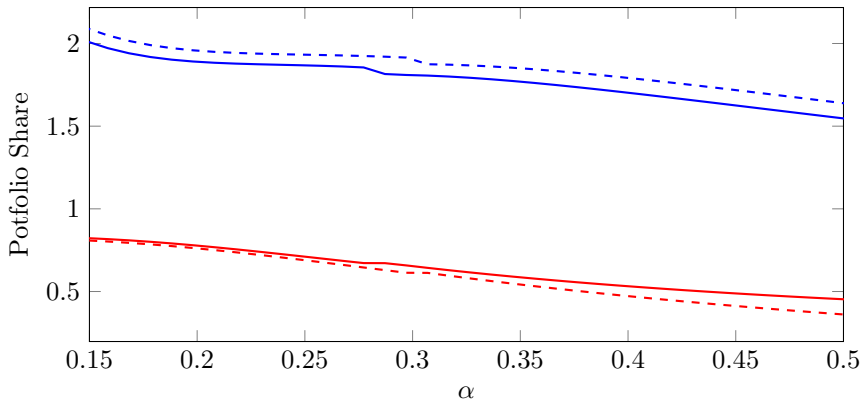


Figure 7: Portfolio shares for trees, θ_q^i. Blue: Agent A. Red: agent B. Solid: baseline economy with inflation risk. Dashed: no inflation risk.

The planner wants to make both agents *exactly* indifferent between the economy with inflation risk and the economy without inflation risk. However, any initial redistribution will change the initial state α_0, which we take into account. After redistribution takes place, the economy starts at a state $x'(x_0)$ where $\alpha'(x_0) = \frac{\bar{\alpha}\eta^A(x_0)}{\xi(x_0)}$ (and π is unchanged). Thus the condition that both agents are indifferent between the economy with no inflation risk at state x_0 and the economy with inflation risk is

$$\mathcal{U}^i_{0_}(\eta^i(x_0)y_0\tilde{q}(x'_0(x_0)), x'_0(x_0)) = \mathcal{U}^{i,\sigma_\pi=0}_{0_}(\chi_i(x_0)y_0\tilde{q}^{\sigma_\pi=0}(x_0), x_0) \qquad i = A, B \tag{16}$$

where $\mathcal{U}^i_{0_}(n, x)$ denote lifetime utility for an initially alive agent of type i with initial wealth n an initial state x. The solution to this fictitious planning problem gives us a well defined measure of efficiency, at least when comparing welfare across initially alive agents. In particular, ξ represents the aggregate efficiency loss of inflation risk: if $\xi(x_0) > 1$, this means we would need to have more trees in the economy to make all agents indifferent between an economy with inflation risk and $\xi(x_0)$ trees, or an economy with $\sigma_\pi = 0$ and 1 tree. Having found $\xi(x_0)$ Then $\xi^i(x_0) \equiv \frac{\eta^i(x_0)}{\chi^i(x_0)}$ represents the compensation each agent requires to be indifferent between the two economies.

Results. Figure 6 shows the efficiency effects for the relevant range of α, with a constant $\pi = \bar{\pi}$. The top panel shows ξ, η^A and η^B, while the bottom panel shows the implied α' after redistribution. First, the efficiency costs of inflation risk are very stable across initial α, and quite high: the economy would need around 2.5 permanently higher traded income to leave everyone indifferent between an economy with inflation risk, and one without it and lower income. Since traded income is only 15% of aggregate income in this economy, this yields a cost of 0.38% of aggregate GDP, still a considerable amount.

Second, we see that there is vast heterogeneity in the compensation required for each type of agent.

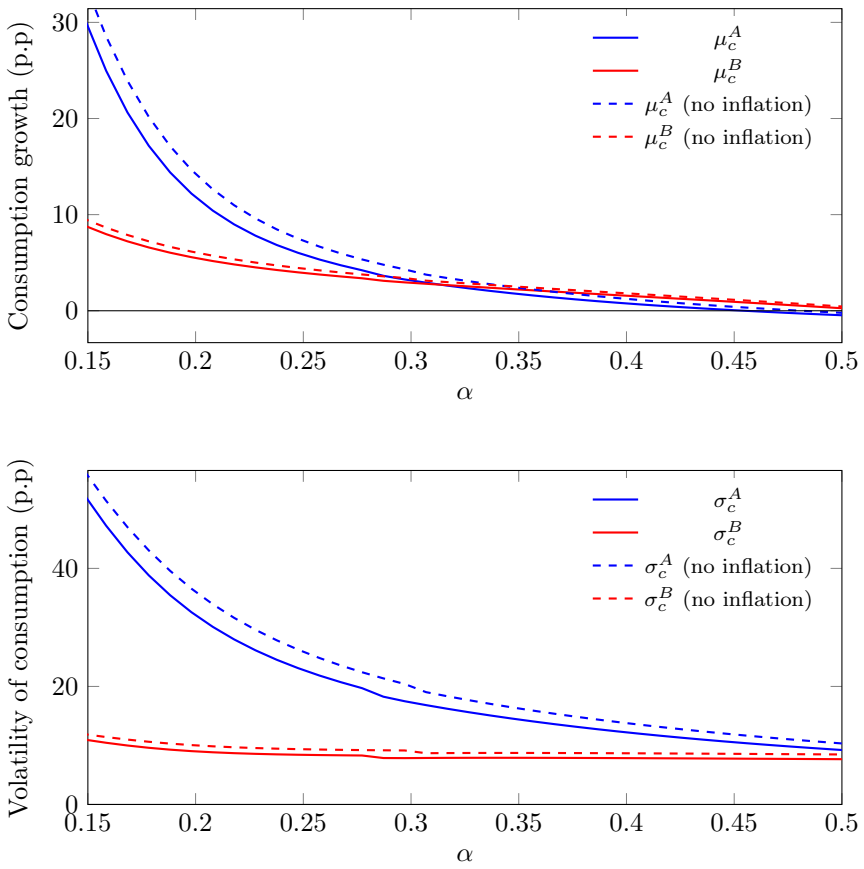


Figure 8: Top: Drift of consumption the process. Bottom: Instantaneous volatility of the consumption process. In both cases, we fix $\pi = 2\%$ and vary α. Blue: Agent A. Red: agent B. Solid: baseline economy with inflation risk. Dashed: no inflation risk.

Risk-tolerant agents (type A) require very high compensation to be indifferent, whereas more risk averse agents require a smaller compensation. For reference, at the mode of the distribution, type A agents require a 7.8% compensation to be indifferent, whereas type B agents require around 0.4%. Furthermore, the degree of required compensation for each type varies greatly with the initial value of α we are comparing to. On first glance, this looks contradictory, since one would have expected more risk averse agents to care about this risk more. However, this reasoning ignores two forces: i) investment opportunities and optimal portfolios change when $\sigma_\pi = 0$, ii) we are accounting for the redistribution that asymmetric compensations generate.

Regarding the first force, Figure 7 plots the optimal portfolio shares of each agent. Lower inflation risk makes borrowing cheaper for type A agents, who now can lever up more aggressively. In other words, inflation risk impairs the ability of agents to share aggregate output risk through borrowing and lending, and this impaired risk-sharing is the primary source of welfare losses. This is especially important when α is low: in that case, expected returns are exceptionally high, so each individual type A agent levers up even more aggressively. Conversely, type B agents lever up less, for market clearing to hold.

In order to understand the implications for consumption, Figure 10 plots the equilibrium drift and volatility of the consumption process for living agents of each type.[14] For given α, the economy with no inflation risk features higher consumption growth for both agents, specially for type A agents. Notice that, due to the OLG structure of this economy, the aggregate consumption growth rate of living agents can differ from the aggregate growth of the endowment. Indeed, the first one is highly variable with α, whereas the second one is constant and equal to μ. As explained in Panageas (2020), the difference is driven by the entry of new agents. We come back to this point below, when we analyze welfare for unborn cohorts.

The economy without inflation risk also features *higher* consumption volatility as a function of α, which may seem counterintuitive at first. This is all due to the increase in borrowing explained above: given that it is cheaper to borrow, type A agents lever up more, which generates higher expected returns for them, which increases their welfare. [15] In turn, this makes α more exposed to output shocks, which makes equity (and thus wealth) more volatile. This force is stronger than the reduction in consumption volatility due to the absence of inflation shocks, and thus consumption volatility goes up. In terms of welfare, the increase in expected consumption growth dominates and thus we get welfare gains. This is particularly important for agent A: the additional consumption volatility does not affect her welfare much due to low risk aversion, and higher consumption growth meaningfully affect her welfare due to high EIS. This explains the large required compensation for type A agents. The above discussion also shows that the bulk of the welfare consequences of inflation risk comes through general equilibrium forces, in particular due to the effect of inflation risk in equilibrium leverage and risk premia.

Regarding the redistribution force, since $\xi^A > \xi^B$, $\alpha' > \alpha$, i.e α would be higher in the compensated scenario. By itself, this makes the situation more attractive to individual B types: for a given wealth, they now have effectively more A types to trade with, so prices are more favorable to them. Thus, all else equal, the redistribution force makes types A less happy with the compensated scenario, and thus they require a larger compensation. On the flip side, the initial redistribution makes types B happier with the compensated scenario, since they due to it they face more favorable prices. When α is low, the jump in α' is so large that $\xi^B(x_0)$ may fall below 1, i.e type B agents would be willing to start with a little less trees if the redistribution were to happen. [16]

[14] These are obtained by applying Ito's Lemma to the consumption policy functions, i.e we compute $\frac{dc_t^i}{c_t^i} = \mu_c(x)dt + \Sigma_c(x)dW$.

[15] Quantitatively speaking, while the change in portfolio weights observed in Figure 7 does not seem very large, the resulting effects on consumption growth are important because the equity premium is large.

[16] Note that this does not imply that they will have less *wealth*, since the price of the trees are endogenous to σ_π and α. Indeed, they could be in a situation where they have less trees, but they are worth so much more that they end up with higher wealth.

5.2.2 Unborn Cohorts

One issue with the above measure is that it completely ignores unborn cohorts. In this section, we compute the required compensation for this agents. In order to do so, consider the utility at time 0 (with aggregate state x_0 and initial output y_0) of an agent that is going to be born at time t. We follow Barcons, Dávila, and Schaab (2026) and model this by saying that the preferences of this agent have the form:

$$\mathcal{U}^i_{t,0} = \mathbb{E}_0\left[\int_t^\infty f^i(c^i_s, \mathcal{U}^i_s)ds\right],$$

i.e the flow of utility is 0 until the time the agent is born, and from then on it has the same functional form as above.[17] Given that the utility at the time of birth can be computed using the obtained value function and that the wealth at birth is $\frac{\delta_q}{\varphi} y_t \tilde{q}(x_t)$, the utility at time 0 of this agent conditional on the initial states can be written as:

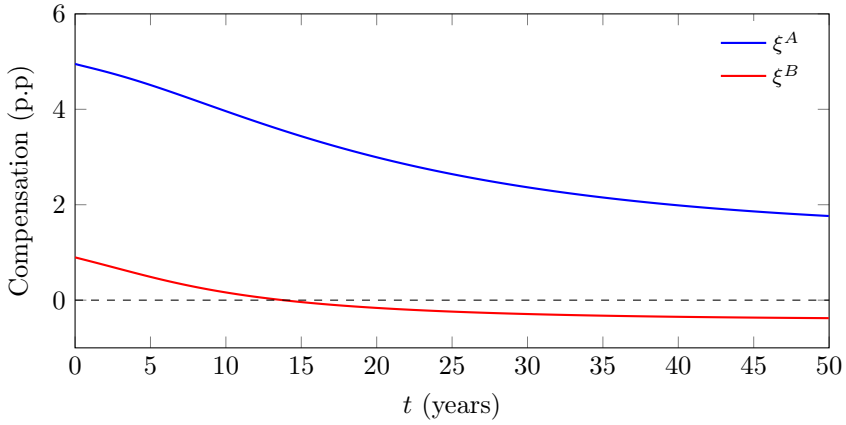


Figure 9: Compensation for unborn cohorts. The horizontal axis indicates the year of birth of the corresponding cohort. Initial state x_0 is fixed at $\pi_0 = 2\%$ and α_0 equal to its meadian.

$$U^{i,t}(x_0, y_0) = \mathbb{E}\left[\frac{\left(\frac{\delta_q}{\varphi} y_t \tilde{q}(x_t)\right)^{1-\gamma_i}}{1-\gamma_i}\left(V^i(x_t)\right)^{\frac{1-\gamma_i}{1-\psi_i}} \middle| x_0, y_0\right] \tag{17}$$

Similar to before, we can compute the number of trees that an unborn agent must be promised at time

[17] In Appendix C.8 we show that is correct under out specific functional form for f.

t=0 (before any uncertainty realizes) to be indifferent between being born in either economy as:[18]

$$\xi^{i,t}(x_0) = \left(\frac{\mathbb{E}\left[(y_t \tilde{q}^{\sigma_\pi=0}(x_t))^{1-\gamma_i} (V^{i,\sigma_\pi=0}(x_t))^{\frac{1-\gamma_i}{1-\psi_i}} | x_0, y_0 \right]}{\mathbb{E}\left[(y_t \tilde{q}(x_t))^{1-\gamma_i} (V^{i}(x_t))^{\frac{1-\gamma_i}{1-\psi_i}} | x_0, y_0 \right]} \right)^{\frac{1}{1-\gamma_i}}.$$

Results. Figure 9 shows the results. As we can see, the compensation is decreasing in the time to birth, even turning negative for type B for cohorts that are far out. The reason is again due to general equilibrium effects. As discussed above, shutting down inflation risk increases the equilibrium consumption growth rates for *living* agents of both types. Since aggregate consumption growth is constant, the only way this can happen is by reducing the consumption growth of future cohorts.

In terms of the mechanism through which this happens, recall that eliminating inflation risk increases risk-taking for type-A agents. For a given initial wealth share, this increases expected consumption growth for living agents, as shown in Figure 10. Higher risk-taking also means that the wealth share of type-A agents has a higher upward drift. Indeed, as shown in Figure 4, the stationary distribution of α shifts to the right when $\sigma_\pi = 0$. Thus, for a given initial α, eliminating inflation risk implies that the wealth share of risk-tolerant agents is expected to increase faster as time goes by. Since valuations are increasing in the wealth share of risk-tolerant agents, growth in α implies an increase in asset valuations, which generates high realized returns for the initially old agents and thus delivers high consumption growth. On the flip side, once valuations are high, this implies lower expected returns moving forward. This implies that cohorts that are born many years after the regime change face worse investment opportunities than the initially old generation, and consequently their welfare gains from eliminating inflation risk are lower. In summary, the initial old generation benefits from a transition phase with good investment opportunities, which they capitalize.

The above discussion underscores an important inter-generational redistribution dimension in the welfare effects of inflation risk: eliminating inflation risk is particularly beneficial to the initially old generation, because it benefits from better investment opportunities. However, the benefits are smaller (and may even turn negative for some agents) for unborn cohorts, which do not benefit from the initial transition.

6 Conclusion

Our central finding is that inflation risk distorts financial markets and risk-sharing in a way that is detrimental to welfare. Quantitatively, the implied welfare loss is sizable—comparable to well-studied costs associated

[18] In order to compute this, the key object is something like $\mathbb{E}\left[(y_t \tilde{q}(x_t))^{1-\gamma_i} (V^{i}(x_t))^{\frac{1-\gamma_i}{1-\psi_i}} | x_0, y_0 \right]$. Appendix C.8 shows how to do so using Feyman-Kac.

with *trend* inflation. In addition, we show that a positive inflation risk premium can arise even when inflation and output shocks are uncorrelated. In our framework, heterogeneity and incomplete markets are sufficient to generate empirically relevant inflation risk premia, in a model consistent with observed macroeconomic and asset-pricing moments.

To maintain tractability and the ability to solve the model globally, we abstract from several potentially important features of the data. A large fraction of nominal debt consists of mortgages, yet we do not model housing explicitly. Similarly, much nominal borrowing is intermediated by banks, which are absent from our framework. Incorporating housing and financial intermediation may alter the magnitude (and possibly the distribution) of the welfare costs we quantify. More broadly, our results hinge on the response of leverage and consumption volatility to changes in inflation uncertainty. Additional empirical evidence on these margins would help assess the realism of the model's implied adjustments and guide future refinements. Finally, we abstract from nominal rigidities and price stickiness, the core ingredients of New Keynesian models. Extending the framework to incorporate such frictions would allow researchers to study the interaction between inflation volatility and price dispersion costs, as well as the implications for optimal monetary policy, which is trivial in our environment.

References

Abrahams, Michael et al. (2016). "Decomposing real and nominal yield curves". In: *Journal of Monetary Economics* 84, pp. 182–200. ISSN: 03043932.

Ajello, Andrea, Luca Benzoni, and Olena Chyruk (Sept. 2019). "Core and 'Crust': Consumer Prices and the Term Structure of Interest Rates". In: *The Review of Financial Studies* 33.8, pp. 3719–3765. ISSN: 0893-9454. eprint: `https://academic.oup.com/rfs/article-pdf/33/8/3719/33509189/hhz094.pdf`.

Alvarez, Fernando and Urban J. Jermann (2004). "Using Asset Prices to Measure the Welfare Cost of Business Cycles". In: *American Economic Review* 94.5, pp. 1223–1256.

Ang, Andrew, Geert Bekaert, and Min Wei (2008). "The Term Structure of Real Rates and Expected Inflation". In: *The Journal of Finance* 63.2, pp. 797–849. ISSN: 0022-1082.

Auclert, Adrien (2020). "Monetary Policy and the Redistribution Channel". In: *American Economic Review* 110.9, pp. 2800–2851.

Bailey, Martin J. (1956). "The Welfare Cost of Inflationary Finance". In: *Journal of Political Economy* 64.2, pp. 93–110.

Baqaee, David Rezza and Ariel T. Burstein (2025). "Welfare and Output with Income Effects and Taste Shocks". In: *American Economic Review* 115.1, pp. 1–46.

Barcons, Sergi, Eduardo Dávila, and Andreas Schaab (Jan. 2026). *Intergenerational Welfare Assessments*. NBER Working Paper.

Baumol, William J. (1952). "The Transactions Demand for Cash: An Inventory Theoretic Approach". In: *Quarterly Journal of Economics* 66.4, pp. 545–556.

Benati, Luca and Juan Pablo Nicolini (Oct. 2025). *The Welfare Costs of Inflation Reconsidered*. Staff Report 675. Federal Reserve Bank of Minneapolis.

Bianchi, Francesco, Renato Faccini, and Leonardo Melosi (Nov. 2023). "A Fiscal Theory of Persistent Inflation". In: *The Quarterly Journal of Economics* 138.4, pp. 2127–2179.

Brunnermeier, Markus K. and Yuliy Sannikov (Feb. 2014). "A Macroeconomic Model with a Financial Sector". In: *American Economic Review* 104.2, pp. 379–421.

Burstein, Ariel and Christian Wellwig (2008). "Welfare Costs of Inflation in a Menu Cost Model". In: *The American Economic Review* 98, pp. 438–443.

Caballero, Ricardo and Alp Simsek (2023). "A Monetary Asset Pricing Model". In: *NBER Working Paper*.

Calvo, Guillermo A. (1983). "Staggered Prices in a Utility-Maximizing Framework". In: *Journal of Monetary Economics* 12.3, pp. 383–398.

Campbell, John Y and Luis M Viceira (2001). "Who Should Buy Long-Term Bonds?" In: *THE AMERICAN ECONOMIC REVIEW* 91.1, p. 47.

Campbell, John Y. and John H. Cochrane (1999). "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior". In: *Journal of Political Economy* 107.2, pp. 205–251.

Campbell, John Y., Carolin Pflueger, and Luis M. Viceira (Aug. 2020). "Macroeconomic Drivers of Bond and Equity Risks". In: *Journal of Political Economy* 128.8, pp. 3148–3185.

Cavallo, Alberto, Francesco Lippi, and Ken Miyahara (June 2023). "Inflation and Misallocation in New Keynesian Models". In: *Proceedings of the ECB Forum on Central Banking 2023*. First draft: May 2023. European Central Bank. Sintra, Portugal.

Cieslak, Anna and Carolin Pflueger (2023). "Inflation and Asset Returns". In: *Annual Review of Financial Economics* 15, pp. 433–448.

Cox, John C., Jr. Ingersoll Jonathan E., and Stephen A. Ross (1985). "A Theory of the Term Structure of Interest Rates". In: *Econometrica* 53.2, pp. 385–407.

Dávila, Eduardo and Andreas Schaab (2025). "Welfare Assessments with Heterogeneous Individuals". In: *Journal of Political Economy* 133.9, pp. 2918–2961.

Doepke, Matthias and Martin Schneider (2006). "Inflation and the Redistribution of Nominal Wealth". In: *Journal of Political Economy* 114 (6).

Duffie, Darrell and Larry G Epstein (1992). "Stochastic Differential Utility". In: *Econometrica* 60 (2), pp. 353–394.

Fang, Xiang, Yang Liu, and Nikolai Roussanov (July 2025). "Getting to the Core: Inflation Risks Within and Across Asset Classes". In: *The Review of Financial Studies*. Advance article, published 16 July 2025, hhaf050.

Fischer, Stanley (1975). "The Demand for Index Bonds". In: *Journal of Political Economy* 83.3, pp. 509–534. eprint: `https://doi.org/10.1086/260339`.

Fisher, Stanley and Franco Modigliani (1978). "Towards an Understanding of the Real Effects and Costs of Inflation". In: *Weltwirtschaftliches Archiv* 114.4, pp. 810–833.

Fleckenstein, Matthias, Francis A. Longstaff, and Hanno Lustig (2017). "Deflation Risk". In: *The Review of Financial Studies* 30.8, pp. 2719–2760. ISSN: 0893-9454, 1465-7368.

Friedman, Milton (1969). "The Optimum Quantity of Money". In: *The Optimum Quantity of Money and Other Essays*.

Gabaix, Xavier (2025). "A Theory of Complexity Aversion". Working paper.

Gârleanu, Nicolae and Stavros Panageas (2015). "Young, Old, Conservative, and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing". In: *Journal of Political Economy* 123.3, pp. 670–685.

Geanakoplos, John and Herakles M. Polemarchakis (1986). "Existence, Regularity, and Constrained Suboptimality of Competitive Allocations When the Asset Market Is Incomplete". In: *Essays in Honor of Kenneth J. Arrow* 3, pp. 65–95.

Golosov, Mikhail and Robert E. Lucas (2007). "Menu Costs and Phillips Curves". In: *Journal of Political Economy* 115.2, pp. 171–199.

Guerreiro, João et al. (2025). "Why Do Workers Dislike Inflation? Wage Erosion and Conflict Costs". Working paper.

Haubrich, Joseph, George Pennacchi, and Peter Ritchken (Mar. 2012). "Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps". In: *The Review of Financial Studies* 25.5, pp. 1588–1629. ISSN: 0893-9454. eprint: `https://academic.oup.com/rfs/article-pdf/25/5/1588/24431472/hhs003.pdf`.

He, Zhiguo and Arvind Krishnamurthy (Apr. 2013). "Intermediary asset pricing". In: *American Economic Review* 103 (2), pp. 732–770. ISSN: 00028282.

Lucas, Robert E. (2000). "Inflation and Welfare". In: *Econometrica* 68.2, pp. 247–274.

Martin, Ian (2017). "What Is the Expected Return on the Market?" In: *The Quarterly Journal of Economics* 132.1, pp. 367–433.

Nakamura, Emi et al. (2018). "The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation". In: *Quarterly Journal of Economics* 133.4, pp. 1933–1980.

Panageas, Stavros (2020). "The implications of heterogeneity and inequality for asset pricing". In: *Foundations and Trends in Finance* 12 (3), pp. 199–275. ISSN: 15672409.

Pflueger, Carolin (May 2025). "Back to the 1980s or Not? The Drivers of Inflation and Real Risks in Treasury Bonds". In: *Journal of Financial Economics* 167, p. 104027.

Piazzesi, Monika and Martin Schneider (2006). "Equilibrium Yield Curves". In: *NBER Macroeconomics Annual* 21, pp. 389–472.

Rotemberg, Julio J. (1982). "Sticky Prices in the United States". In: *Journal of Political Economy* 90.6, pp. 1187–1211.

Rudebusch, Glenn D. and Eric T. Swanson (2012). "The Bond Premium in a DSGE Model with Inflation Targeting". In: *American Economic Review* 102.7, pp. 3111–3144.

Schneider, Andrés (Aug. 2022). "Risk-Sharing and the Term Structure of Interest Rates". In: *Journal of Finance* 77 (4), pp. 2331–2374. ISSN: 15406261.

Schneider, Andres (2025). "Financial Intermediaries and the Yield Curve". Working paper, SSRN (posted 8 September 2025; last revised 8 January 2026).

Schröder, Mark and Costis Skiadas (1999). "Optimal Consumption and Portfolio Selection with Stochastic Differential Utility". In: *Journal of Economic Theory* 89.1, pp. 68–126.

Sidrauski, Miguel (1967). "Inflation and Economic Growth". In: *Journal of Political Economy* 75, pp. 796–810.

Taylor, John B. (1980). "Aggregate Dynamics and Staggered Contracts". In: *Journal of Political Economy* 88.1, pp. 1–23.

Tobin, James (1956). "The Interest-Elasticity of Transactions Demand for Cash". In: *Review of Economics and Statistics* 38.3, pp. 241–247.

Wachter, Jessica A. (2006). "A consumption-based model of the term structure of interest rates". In: *Journal of Financial Economics* 79.2, pp. 365–399. ISSN: 0304-405X.

Werning, Iván (2015). "Incomplete Markets and Aggregate Demand".

Woodford, Michael (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

A Empirical Evidence

Figure 11 shows robustness to the results to alternative choices. The left panel shows what happens if we used core inflation instead of headline CPI since in principle the distinction can be important (Ajello, Benzoni, and Chyruk, 2019; Fang, Liu, and Roussanov, 2025). The right panel shows what happens if we extend the sample to include the Covid period. The finding of a zero or weakly positive correlation is robust to both choices.

B Simple model

B.1 Cashless Limit

In order to explicitly add money, consider a government that issues nominal money and taxes the agents in lump sum. Agents derive utility from holding money, with utility function $v(m) = a\log(m)$ where $m = \frac{M}{P}$ is the real quantity of money. The problem of agent i now reads:

$$\begin{aligned}
\max_{c_{0i},a_i,b_i,b_i^{\$}} \; & u_i(c_{0i}) + a\log(m_0) + \mathbb{E}[u_i(c_{1i}) + a\log(m_{1i})] \\
\text{s.t. } & c_{0i} + m_{0i} + q_s(a_i - \omega_i) + q_b b_i + q_b^{\$} b_i^{\$} = y_{0i} \\
& c_{1i} + m_{1i} = a_i(\bar{y} + \sigma_y \varepsilon_y) + (b_i^{\$} + m_{0i})(\Pi^{-1} - \sigma_\pi \varepsilon_\pi) + b_i \\
& b_i \geq 0.
\end{aligned}$$

The FOCs now read:

$$\begin{aligned}
\mathbb{E}[u_i'(c_{i1})(\Pi^{-1} - \sigma_\pi \varepsilon_\pi)] &= \lambda_i q_b^{\$} \\
\mathbb{E}[u_i'(c_{i1})(\bar{y} + \sigma_y \varepsilon_y)] &= \lambda_i q_s \\
\mathbb{E}[u_i'(c_{i1})] &\geq \lambda_i q_b \quad (\text{with} = \text{ if } b_i > 0). \\
u'(c_{0i}) &= \lambda_i \\
\frac{a}{m_{i0}} + \mathbb{E}[u_i'(c_{i1})(\Pi^{-1} - \sigma_\pi \varepsilon_\pi)] &= \lambda_i \\
\frac{a}{m_{i1}} - u_i'(c_{i1}) &= 0
\end{aligned}$$

From second to last equation, using the FOC for $b^\$$ we get:

$$\frac{a}{m_{i0}} = \lambda_i(1 - q_b^\$)$$

assuming that Π^{-1} and σ_π are such $q_b^\$ < 1$ (i.e inflation is high enough such that having money has a positive opportunity cost), we can rearrange and aggregate to get:

$$\frac{M_0/a}{P_0} = \frac{\lambda_A + \lambda_B}{1 - q_b^\$} \tag{18}$$

thus, for any a we can choose any P_0 by appropriately specifying a monetary policy rule of the form 18. Similarly, for m_1 we have:

$$\frac{M_1/a}{P_1} = \frac{1}{u'_A(c_{A,1})} + \frac{1}{u'_A(c_{B,1})} \tag{19}$$

so the same argument applies. Finally, note that as $a \to 0$, the equilibrium allocation satisfies $m_i \to 0$ for both i, and thus the equilibrium characterization of the main text holds.

B.2 Extensions

B.2.1 Idiosyncratic Income and Recursive Preferences

Assume that agents have preferences $U = W_i(c_{0i}, \mathbb{E}[v_i(c_{1i})])$ (i.e non-time separable), with $v_i(c_{1i})$ increasing and concave, and W strictly increasing and weakly conacave. Assume also that in addition to the tree, they have some non-traded idiosyncratic income at time 1, $w_{i1} = \bar{w}_i + \sigma_{w,i}\varepsilon_w$ with $\sum \sigma_{w,i} = 0$ and ε_w has mean zero and is independent of the other two shocks. This captures both life cycle motives (in $\bar{w}_i$) as well as idiosyncratic income risk in ε_w. Assume that the asset structure remains the same. Now the FOCs of each agent are:

$$\begin{aligned}
\frac{\partial W_i}{\partial c_1}\mathbb{E}[v'_i(c_{i1})(\Pi^{-1} - \sigma_\pi\varepsilon_\pi)] &= \lambda_i q_b^\$ \\
\frac{\partial W_i}{\partial c_1}\mathbb{E}[v'_i(c_{i1})(\bar{y} + \sigma_y\varepsilon_y)] &= \lambda_i q_s \\
\frac{\partial W_i}{\partial c_0} &= \lambda_i \\
\frac{\partial W_i}{\partial c_1}\mathbb{E}[v'_i(c_{i1})] &\geq \lambda_i q_b (\text{with} = \text{ if } b_i > 0).
\end{aligned}$$

and the goods market clearing conditions now read:

$$\begin{aligned} c_{0A} + c_{0B} &= y_0 \\ c_{1A}(s_1) + c_{1B}(s_1) &= \bar{w}_A + \bar{w}_B + y_1(s_1) \quad \forall s_1. \end{aligned}$$

Assume that we consider an equilibrium in which there is trade, ie $|b_i^{\$}| \neq 0$. Again, we want to show that the agent that lends (or more generally, that has positive net exposure to inflation and thus her consumption falls when inflation rises) is the one pricing the nominal bond. Let's go again by contradiction. Assume that the borrower prices the risk free bond. In that case proceeding similar to before:

$$\begin{aligned} \mathbb{E}[v_i'(c_{i1})\left(R^{\$} - R\right)] &= 0 \\ \mathbb{E}[R^{\$} - R] &= -Cov\left(v_i'(c_{i1}), R^{\$} - R\right)/\mathbb{E}[v_i'(c_{i1})], \end{aligned}$$

which assuming v is concave again implies that $\dfrac{\partial v_i}{\partial c_1}(c_{0,1}c_{i1})$ is monotonically decreasing in c_{1i} and thus inflation risk premium is negative, $R > \mathbb{E}[R^{\$}]$. Then, the same argument as before applies: the other agent can improve welfare by substiting all her nominal savings for real savings. Since $v_i(c_{1i})$ is increasing and concave, we have that $\mathbb{E}[v_i(c_{1i})]$ rises under this alternative portfolio, and since W is strictly increasing in c_1, this implies W would increase in this alternative scenario. Thus, an allocation where the nominal borrower is also in her Euler Equation for real bonds can't be an equilibrium. Therefore, the other agent must price the real bond. This implies $\mathbb{E}[R^{\$} - R] > 0$ as long as there is some unhedged exposure to inflation.

B.2.2 Inflation correlated with output

Go back to the baseline model, and assume now that inflation is correlated to ouput. In particular, assume both shocks are normal and $corr(\varepsilon_y, \varepsilon_\pi) = \rho$. Given that, we can use a Cholesky decomposition and write $\varepsilon_\pi = \rho\varepsilon_y + \sqrt{1-\rho^2}\varepsilon_z$, where ε_z is uncorrelated with ε_y. Due to the normality assumption, this implies it is independent.[19] Given this, assume the equilibrium is such that both agents have non-zero net exposure to ε_z (this is generically the case when agents cannot trade in the real bond). In that case due to market clearing, we will have an agent with positive exposure to inflation, and another with negative exposure. We want to show that the agent with positive exposure is the one that prices the real bond.

[19] This normality assumption is important. Generally speaking, we would need to purge out the contribution of ε_y as $\varepsilon_\pi = \mathbb{E}[\varepsilon_\pi|\varepsilon_y] + u$ where $\mathbb{E}[u|\varepsilon_y] = 0$. As it will be clear befloe, the problem with this is that it would be impossible to construct a portfolio that fully hedges the portion of risk that comes from $\mathbb{E}[\varepsilon_\pi|\varepsilon_y]$, since this is a non-linear function of ε_y, whereas we can only do "linear" trades in ε_y. However, in a continous time set-up, everything will be locally linear and thus the intuition derived under normality actually applies.

As before, let's proceed by contradiction: assume the agent with negative exposure is the one that prices the real bond. Define a portfolio that is created by having one unit of the nominal bond and $\frac{\sigma_\pi}{\sigma_y}\rho$ units of the tree. Notice that the pay-off of that portfolio is:

$$\Pi^{-1} - \sigma_\pi(\rho\varepsilon_y + \sqrt{1-\rho^2}\varepsilon_z) + \frac{\sigma_\pi}{\sigma_y}\rho(\bar{y} + \sigma_y\varepsilon_y) \quad = \quad \Pi^{-1} + \frac{\sigma_\pi}{\sigma_y}\rho\bar{y} - \sigma_\pi\sqrt{1-\rho^2}\varepsilon_z$$

so this portfolio has only exposure to ε_z. Let R_z be the real return of this portfolio. As before, if the agent who prices the real bond has negative exposure to this risk ($Cov(c_{1i}, \varepsilon_z) > 0$), we have $R > E[R_z]$. Consider now the other agent, who is positively exposed to this risk. Again, that agent will want to decrease her exposure to this risk to zero (by appropriately selling the portfolio we just created and buying real bonds). By the same argument as before, this is welfare improving, since the agent is increasing her expected pay-off and eliminating a mean preserving spread of her consumption (since ε_z has zero conditional mean).

B.2.3 Idiosyncratic income correlated with inflation

Assume a slighlty extended version of the baseline model where, in addition to the tree, agents have an endowment $w_{i1}(\varepsilon_\pi)$, that may correlate with inflation. Assume also that this is uncorrelated to aggregate output, so $w_{A1}(\varepsilon_\pi) + w_{B1}(\varepsilon_\pi) = \bar{w}$, constant. Assume that both agents have ex-ante the same endowments, but differ in risk aversion. In this case, it is still true that the agent with positive exposure to inflation risk (i.e the agent whose consumption falls if inflation goes up) is the one determining inflation risk premium (and thus, it is weakly positive). In this case, the issue is whether agents can fully hedge their nominal exposure by trading the nominal bond. In particular, their consumption at time $t = 0$ is

$$c_{1i} = a_i(\bar{y} + \sigma_y\varepsilon_y) + b_i^{\$}(\Pi^{-1} - \sigma_\pi\varepsilon_\pi) + w_{i1}(\varepsilon_\pi)$$

so if $w_{A1}(\varepsilon_\pi)$ is non-linear, we cannot kill the exposure. Assume $w_{i1}(\varepsilon_\pi) = \bar{w} + \sigma_{iw}\varepsilon_\pi$ with $\sigma_{Aw} + \sigma_{Bw} = 0$. In that case, we can fully hedge both agents from inflation by setting $b_i^{\$} = \frac{\sigma_i}{\sigma_\pi}$. However, there is no guarantee that this is optimal, and in general it will not be.

B.3 Proof of Proposition 2

We check for constrained Pareto efficiency by studying a planner that takes the equilibrium allocation, and we ask if we could increase the welfare of agent A without lowering the equilibrium expected utility of agent B, U_B^{eq}. The planner's problem is thus:

$$\begin{aligned}
\max_{\{c_{0i},a_i,b_i^{\$}\}_{i=A,B}} \quad & u_A(c_{0A}) + \mathbb{E}[u_A(c_{1A})] \\
& c_{1i} = a_i(\bar{y} + \sigma_y\varepsilon_y) + b_i^{\$}(\Pi^{-1} - \sigma_\pi\varepsilon_\pi) \\
& u_B(c_{0B}) + \mathbb{E}[u_B(c_{1B})] \geq U_B^{eq} \\
& c_{0A} + c_{0B} = y_0 \\
& a_A + a_B = 1 \\
& b_A^{\$} + b_B^{\$} = 0
\end{aligned}$$

Forming the Lagrangian:

$$\begin{aligned}
\mathcal{L} \quad = \quad & u_A(c_{0A}) + \mathbb{E}[u_A(a_A(\bar{y} + \sigma_y\varepsilon_y) + b_A^{\$}(\Pi^{-1} - \sigma_\pi\varepsilon_\pi))] + \gamma[u_B(c_{0B}) + \mathbb{E}[u_B(a_B(\bar{y} + \sigma_y\varepsilon_y) + b_B^{\$}(\Pi^{-1} - \sigma_\pi\varepsilon_\pi))]] + \\
& -\gamma_y[c_{0A} + c_{0B} - y_0] - \gamma_a[a_A + a_B - 1] - \gamma_{b^{\$}}[b_A^{\$} + b_B^{\$}]
\end{aligned}$$

taking FOCs:

$$\begin{aligned}
u'_A(c_{0A}) &= \gamma_y \\
\gamma u'_B(c_{0A}) &= \gamma_y \\
\mathbb{E}[u'_A(c_{1A})(\bar{y} + \sigma_y\varepsilon_y)] &= \gamma_a \\
\gamma\mathbb{E}[u'_B(c_{1B})(\bar{y} + \sigma_y\varepsilon_y)] &= \gamma_a \\
\mathbb{E}[u'_A(c_{1A})(\Pi^{-1} - \sigma_\pi\varepsilon_\pi)] &= \gamma_{b^{\$}} \\
\gamma\mathbb{E}[u'_B(c_{1B})(\Pi^{-1} - \sigma_\pi\varepsilon_\pi)] &= \gamma_{b^{\$}}
\end{aligned}$$

and the constraints. Taking the ratio for the FOCs of different agents for the same asset:

$$\frac{u'_A(c_{0A})}{u'_B(c_{0A})} = \frac{\mathbb{E}[u'_A(c_{1A})(\bar{y} + \sigma_y\varepsilon_y)]}{\mathbb{E}[u'_B(c_{1B})(\bar{y} + \sigma_y\varepsilon_y)]} = \frac{\mathbb{E}[u'_A(c_{1A})(\Pi^{-1} - \sigma_\pi\varepsilon_\pi)]}{\mathbb{E}[u'_B(c_{1B})(\Pi^{-1} - \sigma_\pi\varepsilon_\pi)]} \quad = \quad \gamma$$

and notice that from the FOCs of the agents, taking the same ratio yields:

$$\gamma = \frac{\lambda_A}{\lambda_B}$$

which is indeed constant. The rest of the equations hold as well due to market clearing, and since we choose U_B^{eq} to be the expected utility in equilibrium of agent B, this shows that the equilibrium is constrained efficient: if we had a planner that were to optimize the allocation of assets subject to the same constraints of the agents, the equilibrium allocation would exactly satisfy the FOCs of the planner and thus the planner will not deviate.

C Solving the model

C.1 The HJB of the investors

Given the wealth of an agent and the state $x = (\alpha, \pi)$, the value function can be written in recursive form $\mathcal{U}^i(n; x)$ and satisfies the following HJB equation

$$\begin{aligned} 0 = \max_{c^i, \theta^i} & \frac{1-\gamma^i}{1-1/\psi^i}\mathcal{U}^i \left[\left(\frac{c^i}{((1-\gamma^i)\mathcal{U}^i)^{\frac{1}{1-\gamma^i}}} \right)^{1-\frac{1}{\psi^i}} - (\rho+\varphi) \right] + \\ & + \mathcal{U}_n^i n \left(r + \left(\theta^i\right)^T (\mu - r) - \frac{c^i}{n} + \varphi \right) + \frac{1}{2}\mathcal{U}_{nn}^i n^2 \left(\theta^i\right)^T \Sigma\Sigma^T \theta^i \\ & + \left(\nabla_x \mathcal{U}^i\right)^T \mu_x + \frac{1}{2} Tr(\Sigma_x^T \mathcal{H}_x \mathcal{U}^i \Sigma_x) \\ & + n \left(\nabla_x \mathcal{U}_n^i\right)^T \Sigma_x \Sigma^T \theta^i \\ \text{s.t.} \quad & \left(\theta^i\right)^T \begin{bmatrix} 1 \\ 1 \end{bmatrix} \leq 1 \\ & n' = n(1 + (\theta^i)^T H_t) \end{aligned} \tag{20}$$

The first constraint implicitly assumes the no short-selling of real, short-term bonds that we impose. A levered portfolio in stocks $\theta_q^i > 1$, where none of the borrowing is done in nominal terms $\theta_b^i = 0$ and instead is done in real terms, is not allowed. The second condition represents the capital gain or loss in the event of a rare disaster.

Next we guess and verify that we can write $\mathcal{U}(n; x) \equiv \frac{n^{1-\gamma}}{1-\gamma} V^{\frac{1-\gamma}{1-\varphi}}$. From now on, we omit the i subscript.

$$
\begin{aligned}
0 = \max_{c/n,\theta} & \frac{1}{1-1/\psi}\left[\left(\frac{c/n}{V^{\frac{1}{1-\psi}}}\right)^{1-\frac{1}{\psi}} - (\rho+\varphi)\right] + \\
& + \left(r + (\theta)^T(\mu - r) - c/n + \varphi\right) - \frac{1}{2}\gamma\theta^T\Sigma\Sigma^T\theta \\
& + \frac{1}{1-\psi}\frac{(\nabla_x V)^T}{V}\mu_x + \frac{1}{2}\frac{1}{1-\psi}Tr\left(\Sigma_x^T\left(\left(\frac{1-\gamma}{1-\psi}-1\right)\frac{\nabla_x V(\nabla_x V)^T}{V^2} + \frac{H_x V}{V}\right)\Sigma_x\right) \\
& + \frac{1-\gamma}{1-\psi}\frac{(\nabla_x V)^T}{V}\Sigma_x\Sigma^T\theta^i \\
\text{s.t.} \quad & \theta^T\begin{bmatrix}1\\1\end{bmatrix} \leq 1
\end{aligned}
\tag{21}
$$

The first order condition with respect to the consumption to wealth ratio gives

$$
c/n = V \tag{22}
$$

this implies that the term $\left(\frac{c/n}{V^{\frac{1}{1-\psi}}}\right)^{1-\frac{1}{\psi}}$ simplifies to V in the HJB.

$$
\sum \alpha^i (c/n)^i = \frac{1}{n}\sum c^i = \frac{\omega y}{q} = \frac{1}{\tilde{q}} \tag{23}
$$

The first order conditions with respect to porfolio shares read

$$
\mu_t - r - \gamma\Sigma\Sigma^T\theta + \frac{1-\gamma}{1-\psi}\left(\frac{(\nabla_x V)^T}{V}\Sigma_x\Sigma^T\right)^T = \gamma\begin{bmatrix}1\\1\end{bmatrix} \tag{24}
$$

where γ_t is the Lagrange multiplier on the condition that $\theta^T\begin{bmatrix}1\\1\end{bmatrix} \leq 1$. Now, we multiply the FOC by the vector $[1,-1]$

$$
\mu_{q,t} - \mu_{b,t} = [1\ 1]\left(\gamma\Sigma\Sigma^T\theta - \frac{1-\gamma}{1-\psi}\left(\frac{\nabla_x V}{V}\Sigma_x\Sigma^T\right)^T\right)
$$

The left-hand side of the equation is the differential return of equity relative to nominal, long-term bonds. Since both agents can trade these assets without restrictions, the right-hand side of both agents has to coincide. Moreover, since issuing real risk free debt is not allowed, by market clearing we have that $\theta_q^i + \theta_b^i = 1$ for $i = A, B$, and market clearing imposes $\alpha\theta_q^A + (1-\alpha)\theta_q^B = 1$.

C.2 The law of motion of α_t

We start by deriving the law of motion of the wealth share of type-A investors. Abusing notation, we let $\alpha_t = \alpha_t^A$ but when α_t is treated as a vector, it means $\alpha_t = [\alpha_t^A,\ \alpha_t^B]^T$. In C.1 we show how the value function of investors is homothetic in their wealth, and thus all type-A investors make the same porfolio share and consumption share decisions. We look for an expression of the form

$$d\alpha_t = \mu_{\alpha,t}dt + \sigma_{\alpha,t}^T dW_t \tag{25}$$

and determine $\mu_{\alpha,t}, \sigma_{\alpha,t}$ and $H_{\alpha,t}$. We start by the law of motion of the wealth of type-A agents

$$\frac{dn_t^A}{n_t^A} = (\theta_t^A)^T (\mu_t dt + \Sigma_t dW) - \frac{c_t^A}{n_t^A}dt + \delta_q \frac{\bar{\alpha}}{\alpha_t^A}dt \tag{26}$$

Here, $\theta_t^A = [\theta_{q,t}^A,\ \theta_{b,t}^A]^T$ is the porfolio choice of the investor. The first term of the equation is the financial returns of type-A investors, the second term is the amount they consume, and finally the last term is the wealth of newborn agents. To see where it comes from, type-A agents are born at rate $\varphi\bar{\alpha}$, and each agent receives $\frac{\delta_q}{\varphi}$ trees, valuead at q_t. Thus, $q_t\frac{\delta_q}{\varphi}\varphi\bar{\alpha}$. Now, the expression divide by n_t^A and use that $n_t^A = \alpha_t^A n_t$, and total wealth is the price of capital q_t.

Next, the evolution of total wealth satisfies

$$\frac{d(n_t^A + n_t^B)}{n_t^A + n_t^B} = \begin{bmatrix}1\\0\end{bmatrix}^T (\mu_t dt + \Sigma_t dW) - \alpha_t^T (c/n)_t dt + \delta dt. \tag{27}$$

Here only the return to capital matters because bonds are in zero net supply. Wealth decreases as agents consume and increases due to the arrival of new trees. The depreciation of existin trees is already taken into account inside the term μ_t.

Applying the Itô rule for divisions and with some rearangement, we get that

$$\mu_{\alpha,t} = \alpha_t^A(1-\alpha_t^A)\left(\theta_q^T \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}\left(\mu - \Sigma_t(\Sigma_t)^T \begin{bmatrix}1\\0\end{bmatrix}\right) - [1-1]\,(c/n)_t\right) + \delta_q\left(\bar{\alpha} - \alpha_t^A\right) \tag{28}$$

$$\sigma_{\alpha,t} = \alpha_t^A(1-\alpha_t^A)\theta_q^T \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}\Sigma_t \tag{29}$$

C.3 The law of motion of capital

Equiped with the law of motion of the state $x = (\alpha, \pi)$, we can obtain the law of motion of any other equilibirium object applying Itô's lemma, conditional on $dM = 0$.

$$\frac{d\tilde{q}}{\tilde{q}} = \left(\frac{\nabla \tilde{q}^T}{\tilde{q}} \mu_x + \frac{1}{2} \frac{Tr(\Sigma_x^T \mathcal{H} \tilde{q} \Sigma_x)}{\tilde{q}} \right) dt + \frac{\nabla \tilde{q}^T}{\tilde{q}} \Sigma_x dW \tag{30}$$

$$\frac{dq}{q} = \frac{d(\tilde{q}y)}{\tilde{q}y} = \left(\mu_{\tilde{q}} + \mu_y + \frac{\nabla \tilde{q}^T}{\tilde{q}} \Sigma_x \Sigma_z \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) dt + \frac{\nabla \tilde{q}^T}{\tilde{q}} \Sigma_x dW + \sigma_y dW_y. \tag{31}$$

Similarly for bond prices

$$\frac{dq^b}{q^b} = \left(\frac{\nabla q^{bT}}{q^b} \mu_x + \frac{1}{2} \frac{Tr(\Sigma_x^T \mathcal{H} q^b \Sigma_x)}{q^b} \right) dt + \frac{\nabla q^{bT}}{q^b} \Sigma_x dW. \tag{32}$$

From the definition of returns, we have that

$$\begin{aligned}
\mu_q &= \frac{\omega}{\tilde{q}} - \delta_q + \frac{\nabla \tilde{q}^T}{\tilde{q}} \mu_x + \frac{1}{2} \frac{Tr(\Sigma_x^T \mathcal{H} \tilde{q} \Sigma_x)}{\tilde{q}} + \mu_y + \frac{\nabla \tilde{q}^T}{\tilde{q}} \Sigma_x \Sigma_z \begin{bmatrix} 1 \\ 0 \end{bmatrix} \\
\sigma_q^T &= \frac{\nabla \tilde{q}^T}{\tilde{q}} \Sigma_x + \begin{bmatrix} 1 & 0 \end{bmatrix} \Sigma_z \\
\mu_b &= \frac{\delta_b}{q^b} + \left(\frac{\nabla q^{bT}}{q^b} \mu_x + \frac{1}{2} \frac{Tr(\Sigma_x^T \mathcal{H} q^b \Sigma_x)}{q^b} \right) - \delta_b - \pi \\
\sigma_b^T &= \frac{\nabla q^{bT}}{q^b} \Sigma_x
\end{aligned} \tag{33}$$

We can stack the vector of asset prices $p = [\tilde{q},\ q^b]^T$

C.4 The Stochastic Discount Factor

From (9), after some rearrangement we have that Schröder and Skiadas (1999)

$$f_c(c, \mathcal{U}) = (\rho + \varphi)((1-\gamma)\mathcal{U})^{\frac{1/\psi - \gamma}{1-\gamma}} c^{\frac{1}{\psi}} \tag{34}$$

$$f_{\mathcal{U}}(c, \mathcal{U}) = \left(\frac{-\gamma + \frac{1}{\psi}}{1-\gamma} \right) \frac{f(c, \mathcal{U})}{\mathcal{U}} - \left(\frac{\rho + \varphi}{1-\gamma} \right) \tag{35}$$

We can rewrite it in terms of V and n using the guess and the policy function to get that

$$f_c = \frac{1}{1-\gamma} n^{-\gamma} V^{\frac{1-\gamma}{1-\psi}} \tag{36}$$

$$f_{\mathcal{U}} = \frac{\rho + \varphi}{1-\gamma} \left(\frac{1/\psi - \gamma}{1 - 1/\psi} \left[(\rho + \varphi)^{\psi - 1} V - 1 \right] - 1 \right) \tag{37}$$

Now, we have that

$$m = \exp\left(\int_0^t f_{\mathcal{U}}(c,\mathcal{U})ds\right) f_c \tag{38}$$

which means

$$\frac{dm}{m} = f_{\mathcal{U}}dt + \frac{df_c}{f_c} \tag{39}$$

C.5 Pricing Bonds and Equities

C.6 Numerical Solution

C.7 Equity Price

In the model, agents trade trees that depreciate at rate δ_q, but in order to compare our measures of equity risk premium to other papers and the data, we consider the price of an asset whose dividends are equal to the output. That is, a tree that does not depreciate.

C.8 Ex-Ante Welfare of Unborn Cohorts

Consider the EZ preferences as specified in the Appendix of Garleanu and Panageas. Once we transform them as they do, we get their equation D.8:

$$\frac{1-\gamma}{\alpha}\left(\rho^{\frac{1-\gamma}{\alpha}}(1-\gamma)V_t\right)^{\frac{\alpha}{1-\gamma}} = (\frac{1-\gamma}{\alpha})(1-e^{-\rho\Delta})c_t^{\alpha} + e^{-(\rho+\varphi)\Delta}\frac{1-\gamma}{\alpha}\left(\rho^{\frac{1-\gamma}{\alpha}}(1-\gamma)\mathbb{E}_t[V_{t+\Delta}]\right)^{\frac{\alpha}{1-\gamma}}, \tag{40}$$

which holds whenever the agent is alive. Assume that before the agent is born, the c_t^{α} term is zero, and the probability of dying and discounting is also zero. This is exactly the normalization used in (Barcons, Dávila, and Schaab, 2026). In that case, the recursion simplifies to $V_t = \mathbb{E}_t[V_{t+\Delta}]$, so if we iterate backwards from the time of birth z to zero, we obtain that $V_0 = \mathbb{E}_0[V_z]$ so taking the expectation is justified even when preferences are recursive given our specification of the aggregator and the certainty equivalence operator.

C.8.1 Value Function

In order to compute the value of unborn agents, first define the function $U(y_t, x_t, \tau) = \mathbb{E}\left[\frac{(y_{t+\tau}\tilde{q}(x_{t+\tau}))^{1-\gamma^i}}{1-\gamma_i}(V^i(x_{t+\tau}))^{\frac{1-\gamma^i}{1-\psi^i}}|x_t,\right.$
Note that this function can be written as:

$$U(y_t, x_t, \tau) \quad = \quad \frac{(y_t)^{1-\gamma_i}}{1-\gamma_i}u(x_t, \tau)$$

where $u(x_t,\tau)=\mathbb{E}\left[(\frac{y_{t+\tau}}{y_t}\tilde{q}(x_{t+\tau}))^{1-\gamma^i}(V^i(x_{t+\tau}))^{\frac{1-\gamma^i}{1-\psi^i}}|x_t\right]$, which does not depend on y_t since $y_{t+\tau}/y_t$ only depends on brownian increments after t, which are independent of y_t, x_t. Then, fixing a cohort being born at time T, note that:

$$\begin{aligned} U(y_t,x_t,T-t) &= \mathbb{E}_t\left[0\times dt+\mathbb{E}_{t+dt}[\frac{(y_T\tilde{q}(x_T))^{1-\gamma^i}}{1-\gamma_i}(V^i(x_T))^{\frac{1-\gamma^i}{1-\psi^i}}]\right] \\ &= U(y_t,x_t,T-t)+\mathbb{E}_t\left[dU\right] \end{aligned}$$

which implies $\mathbb{E}_t\left[dU\right]=0$. But using Ito's Lemma to compute this, and using homogeneity in U:

$$dU = d(\frac{(y_t)^{1-\gamma_i}}{1-\gamma_i})u+du\frac{(y_t)^{1-\gamma_i}}{1-\gamma_i}+d(\frac{(y_t)^{1-\gamma_i}}{1-\gamma_i})du$$

but

$$\begin{aligned} d(\frac{(y_t)^{1-\gamma_i}}{1-\gamma_i}) &= y_t^{-\gamma_i}dy-\frac{1}{2}\gamma_i y_t^{-\gamma_i-1}(dy)^2 \\ &= y_t^{1-\gamma_i}(\frac{dy}{y}-\frac{\gamma_i}{2}(\frac{dy}{y})^2) \\ &= y_t^{1-\gamma_i}((\mu-\gamma_i\frac{\sigma^2}{2})dt+\Sigma_y^T dW) \end{aligned}$$

where $\Sigma_y=[\sigma,0]'$ is a 2×1 vector. From Ito:

$$du = (-u_\tau+(\nabla_x u)^T\mu_x+\frac{1}{2}Tr(\Sigma_x\Sigma_x^T\nabla_x^2 u))dt+(\nabla_x u)^T\Sigma_x dW$$

Thus we have:

$$\begin{aligned} dU &= y_t^{1-\gamma_i}((\mu-\gamma_i\frac{\sigma^2}{2})dt+\Sigma_y^T dW)u+\frac{(y_t)^{1-\gamma_i}}{1-\gamma_i}[(-u_t+(\nabla_x u)^T\mu_x+\frac{1}{2}Tr(\Sigma_x\Sigma_x^T\nabla_x^2 u))dt+(\nabla_x u)^T\Sigma_x dW)] \\ &\quad +y_t^{1-\gamma_i}(\nabla_x u)^T\Sigma_x\Sigma_y dt \end{aligned}$$

and therefore $\mathbb{E}_t\left[dU\right]=0$ implies:

$$u_\tau=(1-\gamma_i)\left(\mu-\gamma_i\frac{\sigma^2}{2}\right)u+(\nabla_x u)^T\mu_x+\frac{1}{2}Tr(\Sigma_x\Sigma_x^T\nabla_x^2 u)+(\nabla_x u)^T\Sigma_x\Sigma_y$$

which we can solve forwards in τ (which is for older and older cohorts) from the intial conditions $u(x,0) = (\tilde{q}(x))^{1-\gamma^i}(V^i(x))^{\frac{1-\gamma^i}{1-\psi^i}}$. The equation is defined analogously for the alternative economy with no inflation, we just need to replace the appropriate equilibrium law of motions for the states.

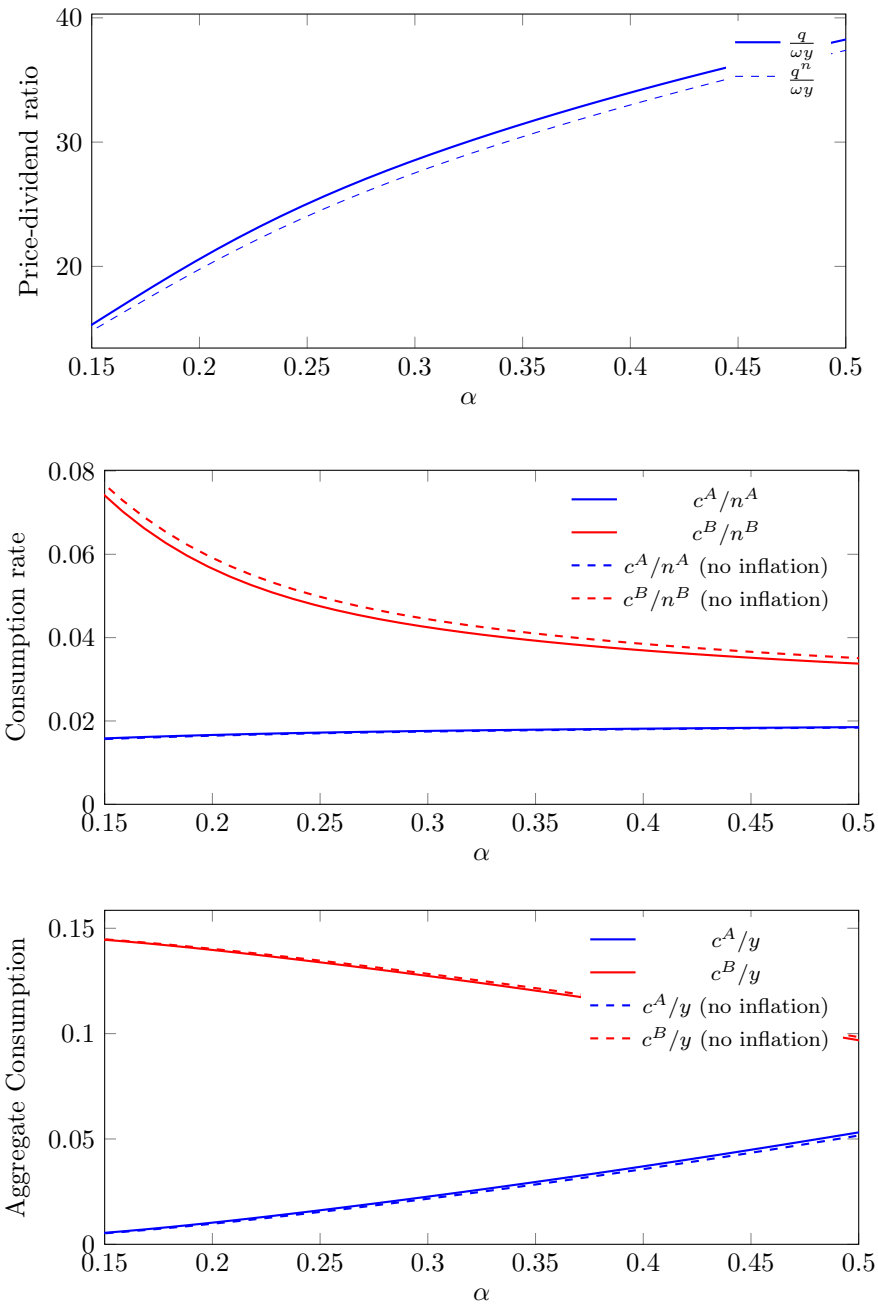


Figure 10: Top: price of capital. Middle: consumption ratio. Bottom: Consumption. In both cases, we fix $\pi = \bar{\pi}$ and vary α. Blue: Agent A. Red: agent B. Solid: baseline economy with inflation risk. Dashed: no inflation risk.

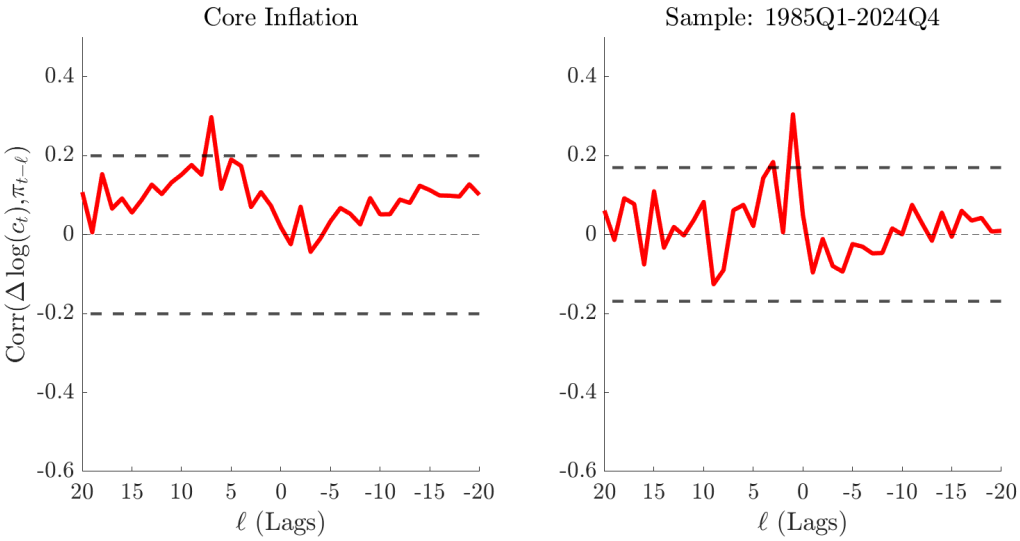


Figure 11: Cross-correlogram between PCE real consumption growth and inflation. Left: use core CPI instead of headline CPI. Right: extend the sample to 2024Q4.Dashed lines are 95% confidence intervals.