The Welfare Cost of Inflation Risk*

Tomás E. Caravello   Marc de la Barrera
MIT   IESE Business School

July 2026

Preliminary and incomplete – comments welcome!

Abstract

When borrowing and lending occur in nominal terms, inflation shocks redistribute wealth among agents. The prospect of these shocks reduces welfare ex-ante by exposing both sides to risk and distorting financial trade. We show that heterogeneity and incomplete inflation indexation are necessary for inflation risk to matter for welfare through this channel, and that the cost is governed by two key moments: the size of nominal positions and the inflation risk premium. The same ingredients generate a positive inflation risk premium even when inflation and aggregate consumption are uncorrelated, consistent with recent U.S. data. We quantify this mechanism in a heterogeneous-agent economy calibrated to match key asset-pricing and macro moments. Eliminating inflation risk yields aggregate welfare gains of about 0.45% of perpetual aggregate consumption, comparable to estimates of the welfare costs of trend inflation due to transaction or price-setting frictions.

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 Dávila, Andrés Schneider, 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

A classic question in macroeconomics relates to the welfare costs of inflation. 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*. In particular, if borrowing and lending take place in nominal terms, inflation risk exposes both parties to risk, which reduces financial trade and welfare. We quantify the welfare cost of inflation risk through this channel, and find it to be on the same order of magnitude of the standard channels mentioned above.

A Back of the Envelope Estimate. In order to get a sense of the orders of magnitude involved, we start by asking: how costly would it be to hedge the existing nominal positions in the economy against inflation? If inflation risk is costly for welfare, then asset prices will reflect it (Alvarez and Jermann, 2004), and thus hedging existing nominal positions should be costly. Estimating this cost requires two basic inputs: the size of nominal positions, and the corresponding inflation risk premium (i.e the difference in expected returns between a nominal and a real bond). We approximate the first one by taking the stock of household and business debt, which is around 130% of GDP, and take the second one from the asset pricing literature (Haubrich et al., 2012; Fleckenstein et al., 2017) to be around 20 bps per year at the relevant maturity. Under these assumptions, the cost of hedging inflation risk would be around 0.26% of GDP per year, which is on the order of magnitude of recent estimates of trend inflation costs due to transaction or price-setting frictions (Cavallo et al., 2023; Benati and Nicolini, 2025). Of course, this calculation raises several conceptual issues: nominal positions in a closed economy are net out to zero; eliminating inflation risk changes equilibrium prices, making the baseline risk premium an invalid input; and it is unclear where the premium comes from in the first place. We turn to a simple general equilibrium model to address each of these.

Key Insights in a Simple Model. We consider a two-period Lucas (1978) economy populated by two agents, who may differ in their preferences or endowments. Agents have three assets at their disposal: trees, that pay an uncertain number of goods at time $t = 1$; real bonds, that pay one unit of good with certainty; and nominal bonds, which pay one dollar with certainty and thus are exposed to inflation risk. We assume prices are fully flexible and consider the cashless limit, so both standard sources of inflation costs are absent. There are two sources of uncertainty: both inflation and aggregate consumption at $t = 1$ are stochastic. We assume that aggregate consumption and inflation shocks are uncorrelated.

The key friction in the economy is that it is not possible to borrow in real bonds. In equilibrium, this implies that all borrowing and lending must take place in nominal terms, and thus is exposed to inflation risk. Our welfare experiment is to compare this economy to an otherwise identical economy in which the volatility of the inflation process is set to zero, holding fixed the aggregate consumption process and all preference/technology parameters. We derive three key results in this set-up.

First, we show that both heterogeneity and incomplete markets (in the form of imperfect indexation to inflation) are necessary for inflation risk to have welfare consequences. The intuition is straightforward: if agents are homogeneous then there is no reason to trade, whereas if they are heterogeneous but can trade in real bonds, agents will not use the nominal bond market, since it exposes both agents to inflation risk.

Our second result shows that our model has one key distinctive prediction: it can generate positive inflation risk premium even when inflation and aggregate consumption are uncorrelated at all leads and lags, which is hard to generate in representative agents models.[1] This is because the agent that is long in nominal bonds is always the one determining inflation risk premium. Therefore, inflation shocks represent bad news for nominal lenders and command a premium, even if they are uncorrelated with aggregate consumption. We further show that this lack of co-movement is exactly what we observe in the post 1985 US data: consumption growth and inflation are essentially uncorrelated at all leads and lags, but empirical estimates of inflation risk premium in this subsample, while smaller than in earlier samples, are still consistently positive.[2] We see this as additional evidence that our key assumptions are relevant at the macro level.

Our third set of results shows how to measure the welfare gains of marginally reducing inflation risk in the above model. In particular, we show that these gains can be decomposed into two sources: a direct channel, which comes from mechanically reducing the risk of existing positions, and an indirect channel, that comes from adjustments of portfolios and prices due to lower risk. While the strength of the latter channel depends on the specific assumptions about preferences and endowments, the strength of the direct channel only depends on the size of nominal positions and size of inflation shocks, regardless of why agents trade in the first place. We show that the direct gains are positive for both lenders and

[1] In those models, inflation risk premia arise from the covariance of inflation and consumption growth, either contemporaneously (standard time-separable CRRA preferences), with lagged consumption growth (habit models as Campbell and Cochrane (1999); Wachter (2006)) or future consumption growth (recursive utility, as Bansal and Yaron (2004); Piazzesi and Schneider (2006)).

[2] If anything, the correlation is positive, which is consistent with a shift in the composition of shocks hitting the economy, with non-inflationary demand shocks being the main driver of output fluctuations Angeletos et al. (2020). Note that this correlation is negative in the 1960-1985 sample, consistent with Piazzesi and Schneider (2006).

borrowers: while inflation shocks produce winners and losers *ex-post* (Doepke and Schneider, 2006b), reducing inflation risk benefits both parties *ex-ante*. Moreover, the gains for lenders equal nominal positions times the inflation risk premium, providing a theoretical foundation for the back-of-envelope formula. Since this derivation is purely about direct effects in a stylized two-period setting we build a quantitative model that retains the main assumptions at its core, but is able to match several key moments of the data.

Quantitative Model. A key complication of quantitatively evaluating this question is that it requires heterogeneity, incomplete markets, and aggregate risk. Solving models with these ingredients is known to be computationally challenging. For that reason, we consider a model whose basic structure is identical to Gârleanu and Panageas (2015); Schneider (2022): there are two types of agents with heterogeneous risk aversion and Elasticity of Intertemporal Substitution (EIS). There is an underlying Overlapping Generations Structure (OLG) structure that makes the stationary wealth distribution non-degenerate. In this model, all heterogeneity is summarized in a single state variable and thus it can be solved globally in reasonable time.[3] The economy has three assets: trees, a nominal bond with positive duration, and an instantaneous (real) bond. As before, the key constraint is that agents cannot borrow in real terms. We assume 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 two relevant moments: inflation risk premium and size of nominal positions; even though we did not target them explicitly in the calibration.

Welfare Costs of Inflation Risk. Having calibrated a general equilibrium model that quantitatively matches the key asset pricing and macroeconomic moments, we turn to study the welfare consequences of reducing inflation risk. First, we follow Robert E. Lucas

[3] As we discuss in Section 6.3, the downside is that we abstract from several ingredients that potentially drive borrowing and lending in reality, such as hump-shaped life-cycle dynamics, self-insuring idiosyncratic income shocks, or buying large indivisible assets.

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

(1987) and ask each agent how much they would be willing to pay to eliminate inflation risk, in terms of a proportional consumption reduction. This gives us a histogram of welfare gains and losses for different agent types and cohorts. The gains from eliminating inflation risk are substantial and quite heterogeneous, with most gains accruing to initially-alive risk-tolerant agents, who hold large leveraged positions and gain the most from cheaper borrowing.

The downside of this exercise is that these gains cannot be aggregated to yield a single number, which we can compare to other costs of inflation. In order to obtain that, we follow Dávila and Schaab (2025); Barcons et al. (2026) and consider a marginal reduction in inflation risk. Once we measure all gains and losses in the same units, we can sum gains for all individuals and get a single efficiency number. Given that we have an infinity number of cohorts, summing compensations for all agents will generally diverge, so we focus in two numbers: the compensation required for cohorts born in the first 50 years, and then a long-run number, defined as the total compensation required for a cohort to be born very far out in the future. We find the gains for the first set of cohorts to be around 0.45% of perpetual aggregate consumption, while the compensation for far out cohorts is essentially zero. We show, however, that the compensation for these cohorts is mainly driven by the effects of inflation risk in very low probability events, where aggregate consumption is abnormally low. Thus, we focus on the effects in the initial set of cohorts as our headline number.

We compare the estimated magnitude of this cost with typical estimates of costs of *trend* inflation from two standard sources: transaction and price-setting frictions. These are the two main costs in the standard New Keynesian literature, with a vast literature that estimates their magnitudes and determinants, which we briefly review in Section 7. The main takeaway is the cost of inflation risk through our channel is similar to those of standard costs of *trend* inflation.

Related Literature. This paper relates to several strands of literature in macroeconomics and finance. First, a large and well established literature studies the welfare costs of inflation through different channels, see for example the exhaustive description of different costs in Fisher and Modigliani (1978). The idea that inflation exposes both borrowers and lenders to risk and that it inefficiently distorts financial markets is not new: it is mentioned in the above cited paper, as well as in Friedman (1977). Our contribution thus lies in describing what key moments of the data are important to identify this cost, and then provide a quantification of this cost in a general equilibrium model that matches those moments.

Our channel of interest relates the paper to a literature that measures the redistributional effects of unexpected inflation shocks, following the pioneering work of Doepke and

Schneider (2006b).[5] The key difference is that we focus on the *ex-ante* welfare losses if agents expect these shocks to occur repeatedly, while Doepke and Schneider (2006b) measures the effects after these shocks occur. This distinction is important because as we show below, the realized shock generates winners and losers, but the possibility of the shock *ex-ante* diminishes welfare for both net lenders and borrowers.

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 indexation. Our paper is most closely related to Schneider (2022, 2025). The present paper features incomplete inflation indexation, and we focus on the implied welfare cost of this friction. Schneider (2022, 2025) 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 argues that asset prices contain relevant information to measure welfare costs of business cycles. Our paper applies similar insights to measure the costs of inflation risk.

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 (Piazzesi and Schneider, 2006; Wachter, 2006). Another branch of literature considers New Keynesian models with asset prices and inflation (Rudebusch and Swanson, 2012; Campbell et al., 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 US 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, without having to take a stance on interpersonal utility comparisons

[5]Other related work includes Doepke and Schneider (2006a); Auclert (2019).

(Baqaee and Burstein, 2025; Dávila and Schaab, 2025; Barcons et al., 2026). 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 presents a simple calculation to get a sense of the potential order of magnitude of these gains. Section 3 presents the key insights of the paper in a two period model. Section 4 presents the quantitative model and its calibration. Section 5 shows that this model can match several key macro and asset pricing moments. Section 6 then uses the model to quantify the welfare gains of eliminating inflation risk. Section 7 compares these costs to estimates of inflation costs from transaction and price-setting frictions. Section 8 offers some final remarks.

2 Motivation: The Cost of Hedging Nominal Positions

In order to get a broad sense at the order of magnitude of this cost, we ask: how costly would it be to hedge the all nominal positions in the economy from inflation risk? If inflation risk is detrimental for welfare, the agents who bear that risk will be willing to pay a significant amount of money to eliminate it, and market prices will reflect that. In order to hedge a portfolio of nominal bonds, we would need to swap it for a portfolio of real bonds with the same duration. This means that, for each nominal bond we hedge, we lose an expected real return $i - E[\pi]$ (the nominal rate minus expected inflation), but we obtain for a guaranteed real return r. The expected cost for each bond is thus $i - E[\pi] - r =$ Inflation Risk Premium. Multiplying this by the total number of nominal assets to be hedged, the total cost of inflation risk can be estimated as:

$$\text{Total Cost of Inflation Risk} = \frac{\text{Nominal Positions}}{\text{GDP}} \times \text{Inflation Risk Premium.} \tag{2.1}$$

Having introduced this simple formula, we turn to discuss how to measure the two key terms in the data.

2.1 Approximating the Key Terms

Nominal Positions to GDP and Maturity. Given that the purpose of the exercise is purely illustrative, we treat all positions as an aggregate with a fixed maturity. We measure the size of these positions in the US using the Flow of Funds, using the 2026 Q1 release. Focusing on debt outstanding at the end of 2025, the data shows this is 21.1 trillion dollars for households, 22.6 trillion dollars for businesses, and 38.2 for federal and state governments

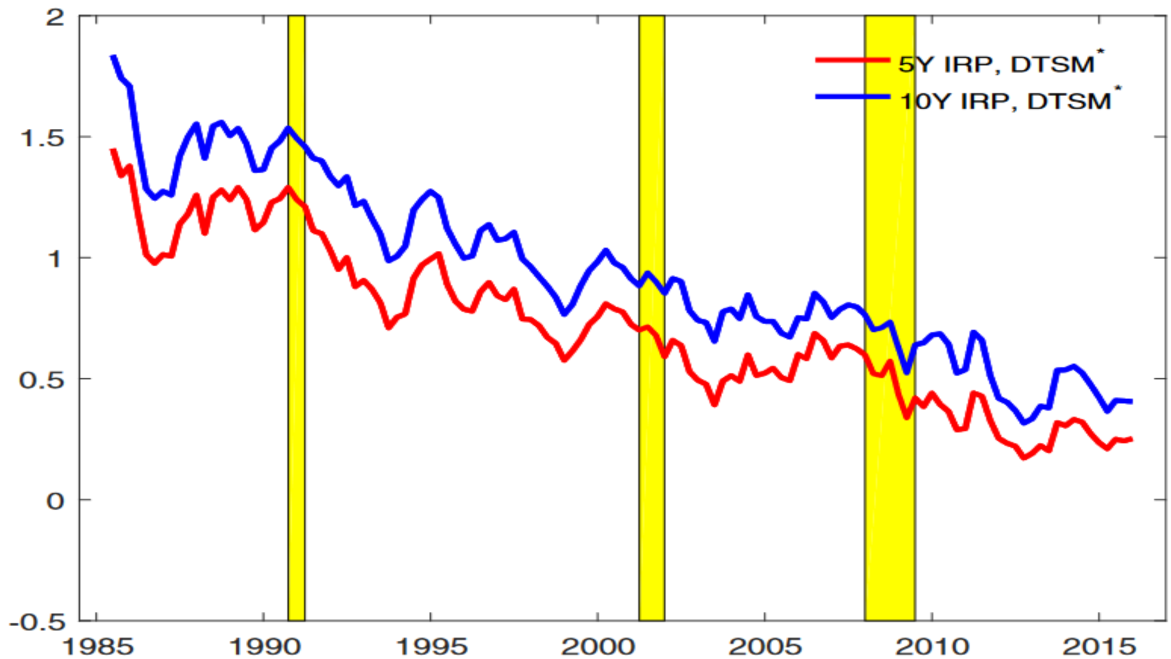

Figure 2.1: Inflation risk-premium for 5 and 10 year bonds derived from the model estimated by Ajello et al. (2019).

combined. Since nominal GDP stood at 31.4 trillion dollars at the end of 2025, private sector debt (combined households and businesses) sits around 139% of GDP, whereas government debt is around 121% of GDP, and thus total debt is around 260% of GDP.

Having these magnitudes in mind, we must take a stance on two key things: i) what is a reasonable average maturity of this debt, ii) how much of it is already indexed, and thus not exposed to this risk. Regarding the first issue, government debt is usually calibrated to have a duration of around 5 years (Auclert et al., 2020; Bianchi et al., 2023). Measuring the duration of private sector debt is much harder. For example, Doepke and Schneider (2006b) estimate the duration on net nominal positions between 4 and 5 years for 2004. Given that recent data features a long period of low nominal interest rates (which mechanically increase maturity) and inflation risk premium is usually reported at 5 years, we keep the maturity to 5 years. Measuring the degree of indexation is also quite hard, however, all private debt is issued in nominal terms, and just a tiny share of government debt is indexed. In order to be conservative, we will use a nominal positions to GDP number of 130%, which is half of the total debt measured. This accounts for potential indexation, as well as the uncertain path of future taxes (Doepke and Schneider, 2006b). In fact, the models we write below only have private debt.

Inflation Risk Premium. Estimating the inflation risk premium is quite hard. Earlier approaches relied on comparing TIPS and nominal yields. However, TIPS are known to be less liquid than nominal bonds, which biases the calculation. Haubrich et al. (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 17bps at a 5-year horizon, 45bps at 10 years, and 100 bps at 30 years. Fleckenstein et al. (2017) uses data on swaps and options to fit a model on physical and risk-neutral inflation. Their sample is from 2008 to 2015. They find the inflation risk premium for a 5 year maturity to be in the order of 20 basis points. Figure 2.1 is taken from Ajello et al. (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, with values of around 20bps for a 5 year maturity at the end of the sample. Given this evidence, we set 20bps as the relevant measure of inflation risk premium.

2.2 Results and Discussion

We can use the above numbers in equation (2.1) to estimate a total inflation risk of

$$0.26\% \text{ GDP per year} = 130\% \text{ of GDP} \times 0.2\% \text{ per year} \tag{2.2}$$

Even if we only used only household debt, which sits at around 70% of GDP, we would obtain a cost of 0.14% per year. These are sizable numbers compared to other typical costs of inflation. To cite a few recent estimates, Benati and Nicolini (2025) find the cost of a 5% inflation rate to be 0.35% of GDP due to transaction frictions. Cavallo et al. (2023) find that the increasing trend inflation from 0 to 10% generates a welfare loss of around 0.4% of GDP due to price-setting frictions.[6] Our simple calculation indicates that, in principle, the cost due to inflation risk can be in the same order of magnitude However, reaching that conclusion at this point is clearly precipitated, since it is unclear why (2.1) is a valid measure of the *welfare* cost of inflation risk in general equilibrium.

Besides refining our admittedly quite raw measurement exercise, there are several conceptual issues with taking the above computation at face value. First, for each lender there is a borrower. We could have equally taken the stock to be negative (since there are also agents that are short in these positions), and find that inflation risk *increases* welfare. Second, it is unclear what experiment the above exercise responds. Ideally, one would like to compare

[6] We review the literature on the quantification of typical welfare costs of inflation in Section 7.

an economy with inflation risk with one without it. But in general equilibrium, eliminating inflation risk will alter asset prices, so using the inflation risk premia of a baseline equilibrium is not valid. Third, it is unclear where this risk premium is coming from, and whether it matters for the exercise. We turn below to a simple model where we clarify these issues.

3 Key Insights in a Simple Model

This section considers this question in a two-period, cashless economy with flexible prices. We make three main points. First, inflation risk matters for welfare only if agents are heterogeneous and nominal contracts are imperfectly indexed. Second, under these assumptions, positive inflation risk premia can arise even when inflation is uncorrelated with aggregate consumption. Third, the direct welfare gain from reducing inflation risk is disciplined by nominal positions and the inflation risk premium.

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. 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. Agents face a borrowing constraint in the real bond. Assume that ε_y and ε_π are mean zero and independent, this is done by simplicity, but we show empirical evidence in support of this assumption below. In this model, we are interested in the welfare effects of lowering σ_π: that is, we lower inflation risk but leave the process for aggregate output untouched.

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 strictly concave for both agents. Agents are endowed with an initial amount of goods at time 0, y_{i0}, and a number of trees, ω_i. Assume there are two agents $i = A, B$ with heterogeneous preferences and endowments. The optimization problem of agent i is:

$$\max_{c_{0i},a_i,b_i,b_i^{\$},c_{1i}} u_i(c_{0i}) + \mathbb{E}[u_i(c_{1i})] \tag{3.1}$$
$$\text{s.t. } c_{0i} + q_s(a_i - \omega_i) + q_b b_i + q_b^{\$} b_i^{\$} = y_{0i}$$
$$c_{1i} = a_i(\bar{y} + \sigma_y \varepsilon_y) + b_i^{\$}(\Pi^{-1} - \sigma_\pi \varepsilon_\pi) + b_i$$
$$b_i \geq 0.$$

The solution to this system implies consumption and asset demands as a function of prices $(q_s, q_b, q_b^{\$})$ and initial endowments. Replacing the expression for c_{1i} and forming a Lagrangian we get the following optimality conditions:

$$\mathbb{E}[u_i'(c_{i1})(\Pi^{-1} - \sigma_\pi \varepsilon_\pi)] = \lambda_i q_b^{\$} \tag{3.2}$$
$$\mathbb{E}[u_i'(c_{i1})(\bar{y} + \sigma_y \varepsilon_y)] = \lambda_i q_s \tag{3.3}$$
$$u'(c_{0i}) = \lambda_i \tag{3.4}$$
$$\mathbb{E}[u_i'(c_{i1})] \geq \lambda_i q_b \quad (\text{with } = \text{ if } b_i > 0). \tag{3.5}$$

In conjunction with the constraints, these characterize the solution to the problem.

Definition 1. A competitive equilibrium is an allocation $(c_{0i}, c_{1i}(s_1), a_i, b_i, b_i^{\$})_{i=A,B}$ and a price vector $(q_s, q_b, q_b^{\$})$ such that i) Given the initial endowments and prices, the allocation solves the optimization problem for both agents; ii) 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 characterizing the equilibrium, a few points are worth discussing.

No Inflation Indexation. As we show below, a key assumption is the impossibility of borrowing in inflation indexed terms. 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.[7] The

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

lack of indexation is often considered a puzzle, (Fischer, 1975).[8] For our analysis, we take the stance that trade in indexed bonds is prohibitively costly and thus does not happen.

Second, we show in Section 3.3 that, when there is trade in nominal bonds, this assumption generates a distinct testable prediction: even if inflation and output are uncorrelated, inflation risk premium is positive. We show below that this prediction is supported in recent US data. Furthermore, in Section 4 we externally calibrate the relevance of this friction, yet the model is able to produce empirically reasonable inflation risk premium and nominal positions as untargeted moments.

Third, 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. 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), or that changing prices or wages requires paying some costs (Rotemberg, 1982; Golosov and Lucas, 2007). Therefore the relevant friction is typically either assumed outright, or microfounded by a direct cost.

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, inflation 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. This also stresses that, in our model, eliminating inflation risk is a policy counterfactual. Thus, in principle, the gains could be materialized by different conduct of policy, as opposed to other sources of uncertainty, which are fundamental and usually seen as unrelated to policy.[9]

Equilibrium Selection and Pricing Real Bonds. 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 price vector could also

[8]See also Doepke and Schneider (2017) for a theory of the emergence of different units of account.

[9]For example, Alvarez and Jermann (2004) make this argument to focus on the gains from eliminating fluctuations at business cycle frequencies, as opposed at gains from eliminating all uncertainty in consumption.

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 On the Necessity of Heterogeneity and Incomplete Indexation

Our first result is that both heterogeneity and incomplete indexation are essential for there to be a cost of inflation risk: if we remove one of those, then the equilibrium allocations do not depend on σ_π, and hence the welfare cost of inflation risk is zero.

Proposition 1. ***Necessity of Heterogeneity and Incomplete Indexation.*** *Assume that either i) agents are homogeneous or ii) both agents can freely borrow in real terms. Then equilibrium allocations do not depend on ε_π, and thus the welfare cost of inflation risk is zero.*

Proof. Note that the only what in which ε_π can affect allocations is if $|b_i^\$| > 0$ in equilibrium, so we need to show that, absent any of the assumptions, there is no trade in nominal bonds. Consider first the case with homogeneity. In that case, both agents have the same preferences and endowments. We can construct a symmetric equilibrium by setting consumption equal to endowments for both agents, which implies $b_i = b_i^\$ = 0$, and finding the supporting prices from (3.2)-(3.5). To see this equilibrium is unique, note that in the individual optimization problem (3.1) given that constraints are a convex set and the objective function is strictly concave, the solution is unique. Given that both agents face the same prices, they must choose the same quantities, so an asymmetric equilibrium cannot exist.

If both agents can freely borrow in real terms, then it is optimal to have all trade in that market, since it is not exposed to inflation risk. To see this more clearly, using (3.2), (3.4) and(3.5) with equality, we get:

$$\mathbb{E}\left[\frac{u'(c_{1i})}{u'(c_{0i})}\left(R^\$ - R\right)\right] = 0$$

which we can rewrite as usual:

$$\mathbb{E}\left[R^\$ - R\right] = -R^{-1}\text{Cov}\left(\frac{u'(c_{1i})}{u'(c_{0i})}, R^\$ - R\right) \tag{3.6}$$

and this equation has to hold for both agents even though the left hand side does not depend on the agent. If $|b_i^{\$}| > 0$, that means that there one agent whose consumption comoves positively with inflation, and the other agent has negative comovement. Due to the montonocity in u', that implies that for $j \neq i$ we must have:

$$\text{Cov}\left(\frac{u'(c_{1i})}{u'(c_{0i})}, R^{\$} - R\right) < 0 < \text{Cov}\left(\frac{u'(c_{1j})}{u'(c_{0j})}, R^{\$} - R\right) \tag{3.7}$$

which contradicts (3.6). Thus, an equilibrium with $|b_i^{\$}| > 0$ cannot exist.

Note that both proofs generalize to the case where $\varepsilon_\pi = d_y \varepsilon_y + \varepsilon_z$ where ε_z is mean zero and independent of ε_y and, i.e there is some correlation between inflation and output, but inflation has some additional risk. Thus, the assumption of independent output and inflation shocks is just for expositional convenience. □

This result shows that those two assumptions are necessary to get the equilibrium to depend on pure inflation shocks, and thus to obtain welfare gains from eliminating them. There is plenty of micro evidence for both assumptions. It is uncontroversial that there is large heterogeneity across households. In terms of nominal exposure, Doepke and Schneider (2006b) show that there is large heterogeneity in nominal positions, and thus nominal shocks have large redistributional effects. While it is very clear that both features are present in the data, some readers may be skeptical that this meaningfully affects the equilibrium. For that reason, we derive below a key distinctive implication of our set-up: we can obtain positive inflation risk premium even when aggregate output and inflation are uncorrelated.

3.3 Positive Inflation Risk Premium with no Output-Inflation

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 $\mathbb{E}[R^{\$} - R] > 0$ even if inflation and output shocks are independent.

Proposition 2. *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:

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

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. This shows that in any equilibrium with trade in nominal bonds, the agent who has $b^{\$} > 0$ is also the marginal investor on 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 is positive. We note that the argument in the proof is more general that the current set-up. The intuition of the result is simple: although inflation and aggregate output are uncorrelated, inflationary shocks are bad news for bondholders, which are the agents who price both nominal and real bonds. Due to this reason, they will demand extra compensation to hold nominal bonds, and thus a positive inflation risk premium emerges. Appendix B.2 shows that versions of Proposition 2 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.

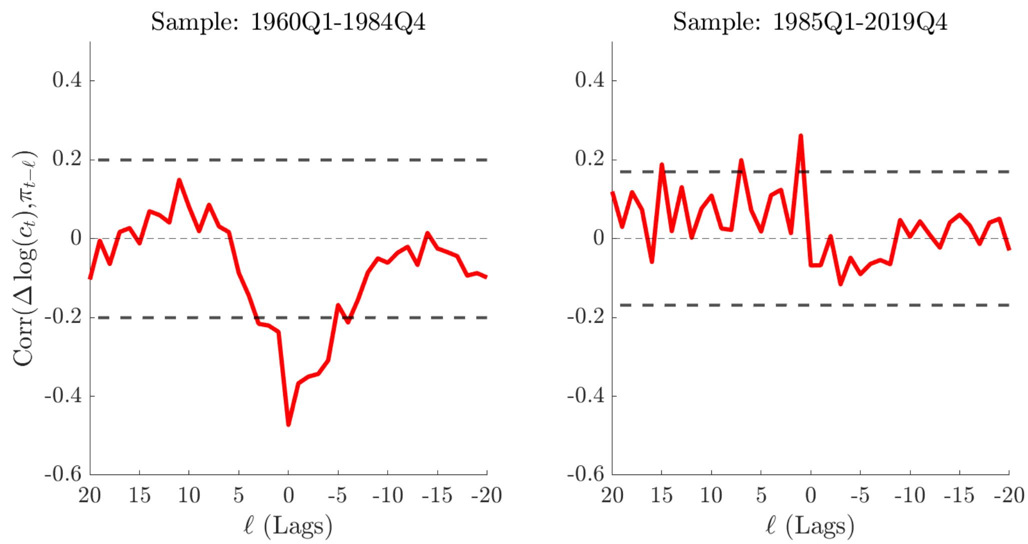

Figure 3.1: Cross-correlogram between PCE real consumption growth and CPI inflation for two different subsamples. Dashed lines are 95% confidence intervals.

Empirical Evidence. While the assumption of zero correlation may sound very knife edge, we argue that this is a reasonable representation of recent data. Figure 3.1 presents the sample cross-correlogram between quarterly aggregate consumption growth and inflation 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. This negative co-movement is no longer observed when we focus on the more recent sample, in the right panel of Figure 3.1. If anything, inflation is (mildly) positively correlated with consumption growth after 1985.[10] This is consistent with the view that, in recent times demand shocks are the primary source of aggregate consumption fluctuations, but their impact in inflation is quite weak (Angeletos et al., 2020).

Under the recent correlation structure, 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). In contrast, once we allow for heterogeneity and incomplete indexation, we can reconcile this correlation structure with a positive inflation risk premium even in recent times, which is a finding in several studies with data post 1980s (Fleckenstein et al., 2017; Ajello et al., 2019). While not definitive, we see this as additional evidence in favor of the two key assumptions in which our analysis hinges upon.

[10]This pattern is robust to using core inflation, and to including the Covid-19 period, see Appendix A.

3.4 Inflation Risk Premium, Nominal Positions, and Welfare

Having explained what are the key ingredients, we now turn to study the welfare costs of inflation risk in this economy. In particular, we follow Dávila and Schaab (2025) and consider a marginal reduction in σ_π. Using that at time $t = 1$ is $c_{i1} = a_i(\bar{y} + \sigma_y \varepsilon_y) + b_i^{\$}(\Pi^{-1} - \sigma_\pi \varepsilon_\pi)$, simple differentiation gives:

$$\frac{dc_{i1}}{d\sigma_\pi} = \underbrace{-\varepsilon_\pi b_i^{\$}}_{\text{Direct Effect}} + \underbrace{\frac{da_i}{d\sigma_\pi}(\bar{y} + \sigma_y \varepsilon_y) + \frac{db_i^{\$}}{d\sigma_\pi}(\Pi^{-1} - \sigma_\pi \varepsilon_\pi)}_{\text{Indirect effect}}. \tag{3.8}$$

This perturbation has two effects in time 1 consumption: first, given portfolios, agents are mechanically exposed to less inflation risk. Second, in general equilibrium, asset prices will change, so agents will optimally re-optimize their savings and portfolio policies in response.

In what follows, we show that a computation similar in spirit to Section 2 will essentially measure the welfare effects only through the direct channel. While incomplete, this is important because the strength of the direct channel only depends on the size of nominal positions, *regardless of why agents trade.*[11] Once $|b_i^{\$}| > 0$, reducing inflation risk directly reduces consumption risk, regardless on whether agents have nominal positions due to life-cycle motives, preference heterogeneity, etc. On the flip-side, predicting the size and direction of the indirect effect will depend on the underlying general equilibrium foundations.

In particular, consider a planner that marginally reduces σ_π, but maintains the equilibrium portfolios and prices as fixed. Thus, c_{i0} is unaffected, and $\frac{dc_{i1}}{d\sigma_\pi}$ only changes through the direct effect described above. In order to get the welfare consequences of this in terms of $t = 0$ goods, we use the basic formula in Dávila and Schaab (2025) to obtain:

$$\frac{dU_i/d\sigma_\pi}{\lambda_i} = \frac{\mathbb{E}[u_i'(c_{1i}) \times \frac{dc_{i1}}{d\sigma_\pi}]}{u'(c_{i0})} = \frac{\mathbb{E}[u_i'(c_{i1})(-\varepsilon_\pi)]}{u'(c_{i0})} b_i^{\$}.$$

This measures the willingness to pay to marginally reduce inflation risk in terms of time 0 goods. For example, if this measure equal -2, agent i values a Δ reduction in σ_π as much as 2Δ units of time-0 good.

Assume that this economy features non-zero nominal trade. This implies that there is a lender, with $b_i^{\$} > 0$, and a borrower. From the budget constraint, we see that positive inflation shocks reduce the consumption of the lender and increase the one of the borrower. Given that marginal utility is decreasing, we have that $\mathbb{E}[u_i'(c_{i1})(\varepsilon_\pi)] > (<)0$ for the lender (borrower). This implies that the above partial derivative is negative for both agents: a

[11] In a sense, this is ex-ante version of the exercises done in Doepke and Schneider (2006b).

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 may seem at odds with the standard intuition whereby inflation shocks produce winners (borrowers) and losers (lenders), e.g Doepke and Schneider (2006b). While this is also true in this model, the above result shows that reducing the size of these shocks increases *ex-ante* welfare of *both* lenders and borrowers, since they are both exposed to less risk.

In order to quantify these potential gains, manipulating the FOCs of the "pricer" yield:

$$\begin{aligned}\frac{\mathbb{E}[u_i'(c_{i1})(-\varepsilon_\pi)]}{u'(c_{i0})}b_i^\$ &= -q_b\frac{E[R^\$]-R}{\sigma_\pi/q_b^\$}b_i^\$ \\ &= -q_b\kappa_\pi b_i^\$.\end{aligned}$$

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, $\kappa_\pi = \frac{E[R^\$]-R}{\sigma_\pi/q_b^\$}$, 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. Assuming $q_b \approx 1$, we can estimate the welfare gains for the pricer of completely eliminating inflation risk as:

$$\frac{dU_i/d\sigma_\pi}{\lambda_i}|_{i:b_i^\$>0} \approx \underbrace{\text{Nominal Positions}}_{=|b_i^\$|}\times\underbrace{\text{Inflation Risk Premium}}_{=\sigma_\pi\kappa_\pi} \tag{3.9}$$

which exactly the formula we used above. Note, however, that this only captures the gains for one of the agents. We can account for the other agent by noting that, given the marginal nature of the perturbation, using the arguments in Dávila and Schaab (2025) we can compute the aggregate efficiency gain of marginally reducing inflation risk, Ξ_E, by summing the $\frac{dU_i/d\sigma_\pi}{\lambda_i}$ terms for both agents (since the perturbation is marginal and both are measured in the same units) as:

$$\Xi_E = \sum_{i=A,B} b_i^\$\frac{\mathbb{E}[u_i'(c_{i1})(-\varepsilon_\pi)]}{u'(c_{i0})} = |b_i^\$|\times\left(\sum_{i=A,B}|\frac{\mathbb{E}[u_i'(c_{i1})(-\varepsilon_\pi)]}{u'(c_{i0})}|\right) \geq |b_i^\$|\sigma_\pi\kappa_\pi$$

where we used the above observation about the signs of $b_i^\$$ and $\mathbb{E}[u_i'(c_{i1})(-\varepsilon_\pi)]/u'(c_{i0})$ to write everything in absolute values. Thus, our earlier "back of the envelope" formula was actually underestimating the welfare gains from eliminating inflation risk since we were only capturing the gains for one side of the market.

Takeaways. Overall, this section justifies the back of the envelope calculation in Section 2 using a simple two period model. We showed that, if there is heterogeneity and incomplete indexation, inflation risk will have welfare costs. While there is ample macro evidence of the prevalence of these two features, we also show that our mechanism generates a distinct implication: inflation risk premium can be positive even if aggregate inflation and output shocks are orthogonal. Finally, we showed that we can approximate the direct welfare gains of reducing inflation risk by multiplying the inflation risk premium times the absolute size of the nominal positions. However, a key shortcoming is that this exercise ignores what would happen in equilibrium if we reduced σ_π, since it agents will re-optimize and prices will adjust in general equilibrium. Furthermore, the model considered so far is highly stylized. Thus, in order to properly quantify the welfare costs of inflation risk, the next section presents a quantitative model.

4 Quantitative Model

In this section we present a quantitative general equilibrium model. The model is similar to Gârleanu and Panageas (2015), which features heterogeneity in preferences. Our key addition is that, as Section 3, markets are incomplete and agents must trade using assets that are exposed to nominal risk. Once we have calibrated the model, we can then evaluate the welfare consequences of inflation risk by resolving the model under the assumption that inflation is deterministic.

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.[12] 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. Thus, the number of goods produced by a tree decreases with the age of the tree. This depreciation rate allows us to

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

capture a decreasing pattern of lifetime earnings, which pushes down the real rate. This is conceptually similar to Gârleanu and Panageas (2015), where agents face a decreasing path of income that incentivizes saving and lowers the real rate. Embedding this in an incomplete markets setting like ours would be unfeasible, so we use tree depreciation as a tractable way of capturing the same force. See Panageas (2020) for more details.

Aggregate output y_t follows a geometric Brownian motion

$$\frac{dy_t}{y_t} = \mu dt + \sigma dW_{y,t}, \tag{4.1}$$

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 an OU process:

$$d\pi_t = \lambda_\pi(\bar{\pi} - \pi_t) + \sigma_\pi dW_{\pi,t} \tag{4.2}$$

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. We assume that $\lambda_\pi > 0$ so the exogenous process π_t is stationary. We assume that $\text{corr}(dW_{y,t}, dW_{\pi,t}) = 0$. As discussed above, this assumption differentiates our model from the rest of the literature, (e.g Piazzesi and Schneider (2006) or Schneider (2022)) which obtain a positive inflation risk premium from the assumed negative correlation between inflation and output. It will be useful to write the model in matrix notation, so we let $dW_t = [dW_{y,t}, dW_{\pi,t}]^T$.

Assets and returns. There are 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, which in equilibrium implies that neither agent holds short term bond. 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 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 nominal perpetuity or 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

$$\frac{\delta_b}{q_t^b}dt + \frac{dq_t^b}{q_t^b} - (\delta_b + \pi_t)dt = \mu_{t,b}dt + \sigma_{b,t}^T dW_t. \tag{4.3}$$

In particular, investors receive the interest payment δ_b/q_t^b per dollar invested, bond prices change by dq_t^b/q_t^b, the coupons decay 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. This implies that the degree of market incompleteness in this economy is directly tied to assumed duration on the traded nominal bond. Since the constraint in short-term debt issuance forces investors to save and borrow assets subject to inflation risk, the equilibrium inflation risk premium is intrinsically linked to the duration of the assets that are actually traded, since this governs to how much inflation risk agents are exposed to.

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 an aggregate portfolio of trees of all maturities. From now on refer to this as trees, and treat it as a single asset. Thus, owning a tree yields $\omega = \int_{s=-\infty}^{t} y_{t,s} = \omega y_t$ in dividends. Letting q_t be the price at time t of a tree, the instantaneous return of holding a tree is

$$\omega\frac{y_t}{q_t}dt + \frac{dq_t}{q_t} - \delta_q dt = \mu_{q,t}dt + \sigma_{q,t}^T dW_t, \tag{4.4}$$

since the tree pays ωy_t in dividends, its price changes by dq_t/q_t, and depreciate 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

$$\begin{bmatrix} (\frac{\omega y_t}{q_t} - \delta_q)dt + \frac{dq_t}{q_t} \\ (\frac{\delta_b}{q_t^b} - \delta_b - \pi_t)dt + \frac{dq_t^b}{q_t^b} \end{bmatrix} = \mu_t dt + \Sigma_t dW. \tag{4.5}$$

where $\mu_t = [\mu_{q,t}, \mu_{b,t}]^T$ is the vector of instantaneous expected returns, and $\Sigma_t = [\sigma_{q,t}, \sigma_{b,t}]^T$ is a 2×2 matrix that captures the exposure of both assets to each shock.

Agents and Demographics. There are two types of investors with recursive preferences as in Duffie and Epstein (1992). Agents' utility at time t is

$$\mathcal{U}_t^i = \mathbb{E}_t \left[\int_t^\infty f^i(c_s^i, \mathcal{U}_s^i) ds \right], \tag{4.6}$$

where

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

is the corresponding aggregator. 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. Regardless of the type, each newborn is endowed with $\frac{\delta_q}{\varphi}$ new (aggregate) trees, so in aggregate, new trees are created at rate δ_q.[13] 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:

$$\frac{dn_t^i}{n_t^i} = [r_t dt + (\theta_t^i)^T ((\mu_t - r_t)dt + \Sigma_t dW)] - \frac{c_t^i}{n_t^i} dt + \varphi dt, \tag{4.7}$$

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 (4.6) subject to (4.7) 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 $n_{t,s}^i$ and the aggregate state of the economy x_t, which we describe below.

[13] That is, agents born at t get $1/\varphi$ date-t trees, which have are equivalent to δ_q/φ units of an aggregate portfolio of trees of all ages.

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 $x = (\alpha, \pi) \in [0, 1] \times \mathbb{R}$ 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. We show in Appendix C.1 that homotheticity allows to write the value function as

$$\mathcal{U}^i(n, x) = \frac{n^{1-\gamma^i}}{1-\gamma^i} V^i(x)^{\frac{1-\gamma^i}{1-\psi^i}}. \tag{4.8}$$

Furthermore, we show this implies that all agents of the same type take the same portfolio decisions and consumption-to-wealth ratio:

$$\begin{aligned} c_t^i(n, x) &= V^i(x)n \\ \theta^i(n, x) &= \theta^i(x). \end{aligned}$$

where in particular $V^i(x)$ is the same function that enters in 4.8. This linearity in wealth allows us to aggregate consumption and assets demands within each group. For example, consumption of agents type A satisfies $\bar{\alpha} \int_{-\infty}^{t} \varphi e^{-\varphi(t-s)} c_{t,s}^A ds = n_t^A V^A(x)$, where n_t^A is total wealth of type A.

Given that output shocks dW_y are permanent, we can rescale the prices of trees by y_t as: $q_t = y_t \tilde{q}(x)$, where $\tilde{q}(x)$ is the aggregate price-dividend ratio, which also equals to the aggregate net-worth. 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 policy functions $\{V^i(x), \theta^i(x)\}$ and prices $\{\tilde{q}(x), q_b(x), r(x)\}$ such that:

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

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

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

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

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

Note that we drop the value functions from the definition of the equilibrium because they are directly linked to the consumption-wealth ratios V^i. 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 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 marginal utility of agent i satisfies:

$$\frac{dM_t^i}{M_t^i} = f_{\mathcal{U}}^i + \frac{df_c^i}{f_c^i} = -(r^i(x) + \varphi)dt - (\kappa^i(x))^T dW_t. \qquad (4.9)$$

Following Gârleanu and Panageas (2015), the corresponding SDF has drift $r^i(x)$ and price of the risk $\kappa^i(x)$.[14] In Appendix C.4 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

[14] That is, the extra φ introduced by the risk of dying is ommited. This is because we assumed perfect annuity markets.

nominal bonds because these are freely traded in equilibrium. For pricing all other assets, we use the SDF of agent who has $\theta_b^i > 0$. In our calibration this is agent B, which has higher risk aversion. Note that in the selected equilibrium, this is a valid aggregate SDF, in the sense that it can price all available assets including the short term bond, since this investor is marginal for all assets and thus has access to dynamically complete markets.

Denote the the price of a zero coupon nominal bond with maturity τ by $Q^{\$}(x,\tau)$. Equipped with the stochastic discount factor and the law of motion of the state we can obtain $Q^{\$}(x,\tau)$ by solving the following PDE:

$$\begin{aligned} -Q^{\$}_{\tau}(x,\tau) + \mathcal{L}[Q^{\$}(x,\tau)] - (r(x)+\pi)Q^{\$}(x,\tau) - (\nabla Q^{\$})^T\Sigma_x(x)\kappa(x) = 0 \\ Q^{\$}(x;0) = 1. \end{aligned} \tag{4.10}$$

Here $Q^{\$}_{\tau}$ 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 $y^{\$}(x,\tau) = -\frac{1}{\tau}\log(Q^{\$}(x,\tau))$. Similarly, we can find the price of a *real* zero coupon bond by solving:

$$\begin{aligned} -Q_{\tau}(x,\tau) + \mathcal{L}[Q(x,\tau)] - r(x)Q(x,\tau) - (\nabla Q)^T\Sigma_x(x)\kappa(x) = 0 \\ Q(x;0) = 1. \end{aligned} \tag{4.11}$$

In order to compute term and inflation risk premia, we follow Haubrich et al. (2012). In particular, we create counterfactual zero risk premia nominal and real bond prices by solving equations (4.10)- (4.11) with $\kappa(x) = 0$ for all x. Denoting these prices by $Q^{\$,rf}(x,\tau)$ and $Q^{rf}(x,\tau)$ and their corresponding yields by $y^{\$,rf}, y^{rf}$. Then we compute the real and nominal term premia, and inflation risk-premia at maturity τ as:

$$\begin{aligned} RTP(x,\tau) &= y(x,\tau) - y^{rf}(x,\tau) \\ NTP(x,\tau) &= y^{\$}(x,\tau) - y^{\$,rf}(x,\tau) \\ IRP(x,\tau) &= NTP(x,\tau) - RTP(x,\tau). \end{aligned}$$

4.4 Calibration

We calibrate the model at the annual frequency. The parameters are shown in Table 1. Importantly, we do *not* target the inflation risk premium it in the calibration: we use it as an untargeted moment. We do this because the key source of the welfare cost is the size of market imperfection. Since the above model has zero correlation between inflation and

	Description	Value
	Inflation Process	
λ_π	Persistence of inflation	0.390
$\bar{\pi}$	Inflation target	0.026
σ_π	Volatility	0.013
	Output Process	
μ	Output growth rate	0.02
σ	Volatility of output	0.04
ω	Dividend share	0.15

Param.	Description	Value
	Agents	
ρ	Discount rate	0.001
φ	Death rate	0.02
$\bar{\alpha}$	Share of type-A investors	0.025
γ^A	Risk aversion of type-A agent	0.9
ψ^A	EIS of type-A agent	1.2
γ^B	Risk aversion of type-B agent	10
ψ^B	EIS of type-B agent	0.035
	Assets	
δ_q	Trees depreciation	0.03
δ_b	Bond coupon decay	0.2

Table 1: Calibrated parameters.

output shocks, all inflation risk premia must come from our mechanism. If in reality part of the inflation risk premia came from other sources but we fully matched it via market incompleteness, we would overestimate the size of the friction and thus the welfare costs. Our approach is to calibrate externally the size of the friction, and then, as a validation moment, ask how much of the observed inflation risk premium we can explain.

We use Gârleanu and Panageas (2015) specification of the output process, with an annual growth rate of 2%, and $\sigma = 0.04$ so that time-integrated data from our model can roughly reproduce the first two moments of annual consumption growth. An instantaneous volatility of $\sigma = 0.04$ corresponds to a volatility of 0.033 for model-implied, time-integrated, yearly consumption data (Campbell and Cochrane (1999)). We set ω to 0.15. This generates a (traded) wealth-to-GDP ratio of around 4.5 on average, which is roughly consistent with what is observed in the data. This calibration of ω also implies that the size of nominal positions is also left as an untargeted moment.

For inflation, we estimate an AR(1) process for inflation in the 1985-2025 sample, and map the parameters to continuous time. Figure 5.2 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 model-implied one, and it also has fatter tails. We note that the persistence of inflation is moderate, $\lambda_\pi = 0.39$, which means that the half life of the shock is around seven 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.

Preferences and $\bar{\alpha}$ are calibrated to roughly match the risk free rate and the equity risk

Moment	Baseline	$\sigma_\pi = 0$	Data
Real Rate	1.04%	1.16%	0.83%
Real Rate Volatility	0.47%	0.56%	1.33%
Equity Risk Premium	4.72%	5.01%	4.64%
Equity Volatility	22.27%	23.79%	16.05%
IRP 5-yr	0.14%	- %	0.17%
IRP 10-yr	0.19%	- %	0.45%
Credit / GDP	108.7%	125.7%	139.8%

Table 2: Model moments vs. data. The real rate is taken from FRED as the difference between the 1-year treasury yield (DTB1YR) and expected inflation (EXPINF1YR). The equity risk premium is from Martin (2017), the return volatility is from the Shiller dataset, and the inflation risk premia are from Haubrich et al. (2012). Except for the equity risk premium and the inflation risk premia, the time-frame is 1985–2025.

premium with moderate risk aversion, as in Gârleanu and Panageas (2015). 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 a low EIS for agent B is required to generate an empirically relevant equity risk premium. The depreciation rate for trees δ_q is chosen to match the level of real rates.

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 for the reasons discussed in Section 2.

5 Positive Implications

Having described the model and its calibration, the present section analyzes its positive implications. This is important, because in order to measure the welfare cost of some risk, we must make sure our model can reproduce the prices and quantities of risk present in the data (Alvarez and Jermann, 2004). This section shows that our model can, and thus we proceed to normative analysis in Section 6.

Standard Asset Pricing Moments. Table 2 shows several implied asset pricing moments on average and compares them to the data. Figure 5.1 reports those moments as a function of α (averaged across inflation values) and Figure 5.2 shows the stationary distributions for both states. As we can see in Table 2, the model roughly matches the average level of the real rate, as well as its volatility. It also produces a sizable equity risk premium, which is broadly in line with the data. The model is also able to produce non-trivial volatility in equity prices, which in fact is somewhat high relative to the data. The mechanism to match these classic asset pricing moments 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 negative 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 5.1 shows the workings of the model 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, around $\alpha = 0.3$), the equity risk premium is low and around 2% a year. However, when negative output shocks hit and the wealth share of risk-tolerant investors decreases, the equity risk premium increases sharply. This is due to two foces. On the one hand, for low α, asset prices mostly reflect the preferences of the risk averse agent. Second, individual risk-tolerant agents become highly levered for low α, which further amplifies stock market volatility at low α. This makes stock riskier and thus they command a large equity premium in those states. This behavior is consistent with Martin (2017), who documents that the equity premium is relatively low in normal times, but has very large spikes in times of stress.

Nominal Positions. The top right panel in Figure 5.1 shows the size of nominal positions, in this economy. We follow Schneider (2022) and denote this as “Credit". This is computed as Credit $= (1 - \alpha)\theta_B(\alpha, \pi)\tilde{q}(\alpha, \pi)$. As we can see, the size of nominal positions 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, Table 2 shows that the average Credit to GDP ratio under our calibration is around 109% of GDP, somewhat below the average for this sample, around 140% of GDP. Thus, our model generates nominal positions that are somewhat smaller (but in the same order of magnitude) of those observed in the data. Since larger nominal positions would make inflation risk more important, if anything this biases our cost estimate downwards.

Inflation Risk Premium. The bottom right panel in Figure 5.1 shows the inflation risk premium on a 10-year zero coupon bond. Relative to Gârleanu and Panageas (2015), the key novelty of our model is that it can generate positive inflation risk premium out of

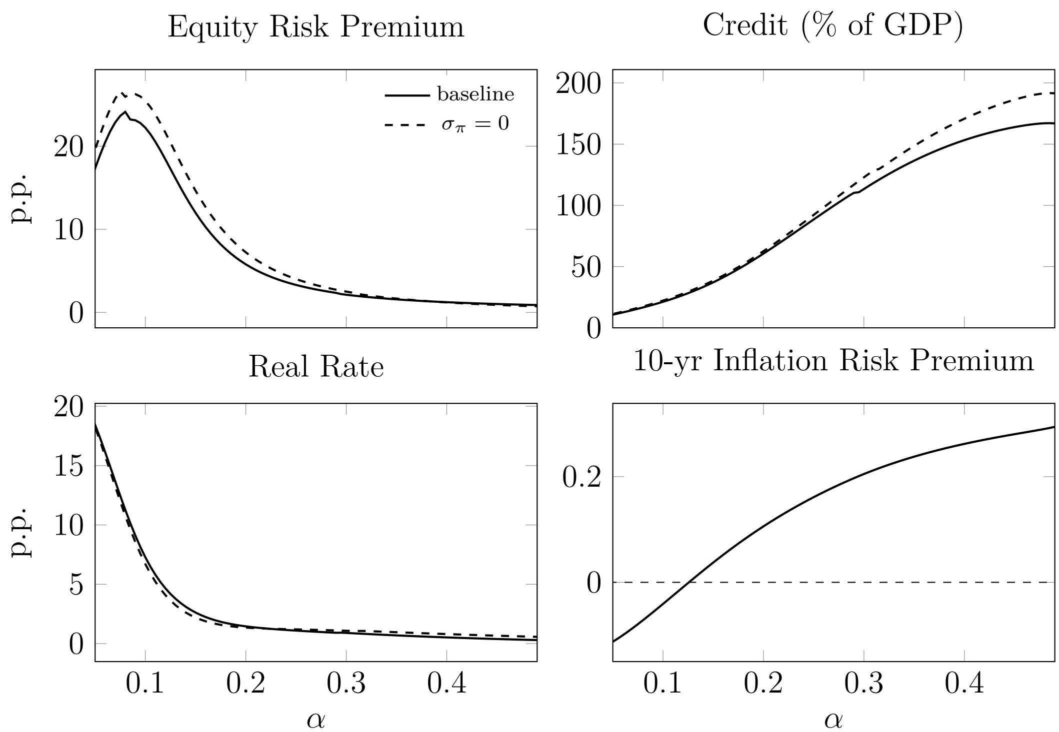

Figure 5.1: Selected asset-pricing and macroeconomic moments as a function of α. Solid line: baseline calibration. Dashed line: no inflation risk ($\sigma_\pi = 0$).

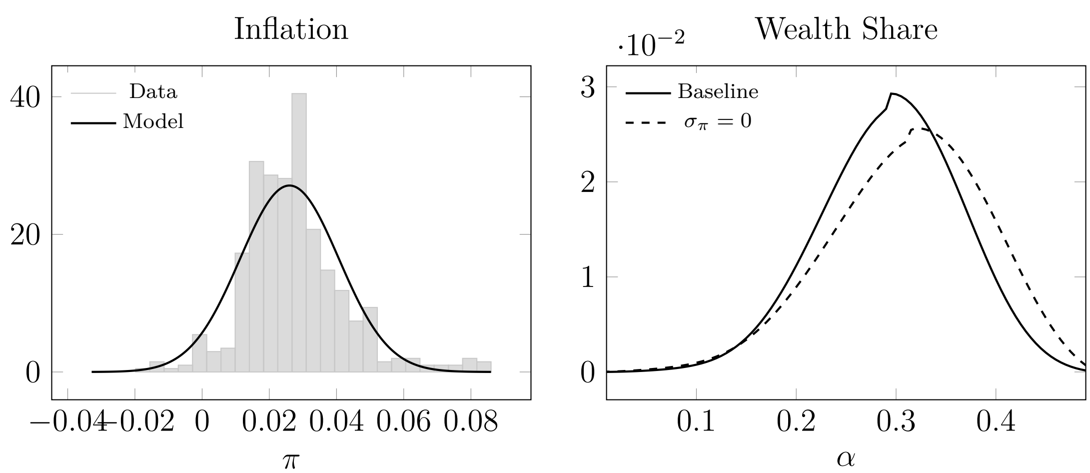

Figure 5.2: Stationary distributions. Left panel: inflation density in the data and in the estimated OU process. Right panel: stationary density of the wealth share α.

uncorrelated inflation and output shocks, due to imperfect inflation indexation. The basic mechanism is simple: since agents have different risk aversions, they would optimally choose different exposure to stocks. In equilibrium, this requires that they trade in the nominal bonds market. Then the intuition of Section 3 applies: since there is non-zero exposure to this risk in equilibrium, the inflation risk premium is tends to be positive. We can also see

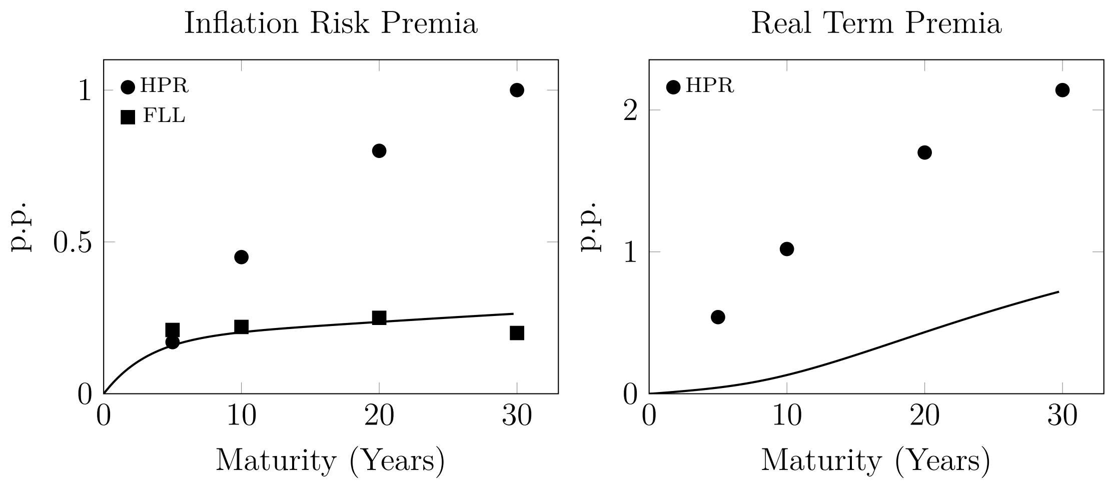

Figure 5.3: Term Structure of Nominal and Real Risk Premia. Circles: estimates from Haubrich et al. (2012). Squares: estimates from Fleckenstein et al. (2017).

that inflation risk-premium increases with α. The intuition is that, as α grows, there is more credit. This makes agents more exposed to inflation risk, and thus inflation risk-premium increases.

Note however that, for very low values of α, inflation risk premium may turn negative. The intuition is as follows: at low α, type B agents have most of their wealth in trees (since they have essentially no counterpart to trade with). Inflationary shocks boosts asset valuations by increasing α, and thus correlate with higher consumption growth through this channel. When α is low i) stock prices are very sensitive to α, and ii) nominal positions are small. Thus the effect on equity valuations dominates over the direct effect (by which inflationary shock liquate nominal assets) and thus a negative inflation risk premium emerges. This feature is consistent with Haubrich et al. (2012) and Fleckenstein et al. (2017), who find a positive risk premium on average, but it can occasionally turn negative. Note, however, that the states where this happens are very unlikely.

The average values across the α distribution are shown in the bottom part of Table 2. Our model is able to produce a sizable inflation risk-premium at a 5-year horizon, which is very much in line with the findings in the empirical studies discussed in Section 2. This is remarkable given that our calibration left inflation risk premium as an untargeted moment. In our view, this provides additional validation for the empirical relevance of the mechanism proposed in the paper.

Term Structure of Risk Premia. Figure 5.3 depicts the term structure for inflation and real term premium, averaging across all states with the stationary distribution. Inflation risk premium is essentially zero at low maturities, and its term structure is upward sloping.

The figure also shows the empirical estimates reported in Haubrich et al. (2012) (circles) and Fleckenstein et al. (2017) (squares). The model fits both estimates at 5-year maturities almost perfectly. For longer maturities, the model predicts a relatively flat term structure, in line with the estimates in Fleckenstein et al. (2017), but too low compare to the estimates in Haubrich et al. (2012).

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 Schneider (2022): in these models, negative output shocks reduce α and increase real rates, as can be seen in Figure 5.1. 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 Haubrich et al. (2012).

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 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 lines in Figures 5.1 and 5.2 shows what happens if we shut down inflation risk in the model, i.e if we set $\sigma_\pi = 0$ leaving all other structural parameters unchanged. 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 equity premium as a function of α shifts up. As we show in Figure 6.2 below, as inflation risk falls, borrowing becomes more attractive for type A agents. They use it to take larger positions in trees, which increases the drift in their wealth and shifts the stationary wealth distribution to the right. Since they are more levered, this also increases the equity risk premium.

Looking at the implied moments in Table 2, the impact is relatively small. Both the equity premium and volatility increase due to increased borrowing by type A agents. Credit to GDP increases by 17 percentage points. 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, but also the volatility of equity and the equity risk premium.

6 The Welfare Costs of Inflation Risk

Having shown that our model can successfully reproduce the key quantities and prices of the relevant risks, we now use it to quantify the welfare gains of eliminating inflation risk.

We proceed in two steps. In Section 6.1, we follow Robert E. Lucas (1987) and compute how much agents in the model would be willing to pay to eliminate inflation risk, in terms of a proportional consumption reduction. This will give us a histogram of welfare gains and losses for different agent types and cohorts. The downside is that these numbers cannot be aggregated to yield a single efficiency number, they are in different units and correspond to a global change. In Section 6.2, we follow the approach proposed by Dávila and Schaab (2025); Barcons et al. (2026) to compute an estimate of the aggregate efficiency gains from eliminating inflation risk. We compare this number with other costs of inflation in Section 7.

6.1 Gains Across Different Agents

In order to quantify the welfare consequences of inflation risk, we compare the welfare of agents in two different economies: the baseline economy, and one without inflation risk, where before time starts and any trade takes place, monetary policy sets makes inflation deterministic ($\sigma_\pi = 0$), and leave the rest of structural parameters unchanged. We assume that both economies start at the same level of output y_0, and initial state $x_0 = (\alpha_0, \pi_0)$. For notational purposes, denote all the objects in the equilibrium without inflation risk by adding a superscript $\sigma_\pi = 0$. For example, the consumption policy function of type A agents is denote as $V^{i,\sigma_\pi=0}(x)$, $\tilde{q}^{\sigma_\pi=0}(x)$ denotes the price of a tree at state x in that equilibrium, and so on.

Our formulation implies that all agents with the same preferences $i =\in \{A, B\}$ born at the same date $t \in [0, +\infty)$ have the same endowments, and thus have exactly the same welfare. Therefore we have two dimensions of heterogeneity (preferences and cohort of birth), and thus the type of the agent is defined by a pair $(i, t) \in \{A, B\} \times [0, +\infty)$. Our population consists of mass 1 of initially alive agents, which are treated as born at time $t = 0$ (which we assume have homogeneous endowments within each i), and then a number total number of agents φdt being born at each instant, out of which a share $\bar{\alpha}$ is risk-tolerant ($i = A$), and the rest is risk averse.

Denote the wealth at birth of a cohort born at time t of type i in the baseline equilibrium as $\bar{n}^{i,t}(x_t)$. For an agent of type (i, t). We measure the welfare consequences of eliminating inflation risk using the proportional compensating variation $\mathcal{C}^{i,t}(x_0)$ as in defined as the solution to the following equation:

$$\mathbb{E}_0[\mathcal{U}_t^{i,t} | x_0, n_t^{i,t} = \bar{n}^{i,t}(x_t)] = \mathbb{E}_0[\mathcal{U}_t^{i,t,\sigma_\pi=0} | x_0, n_t^{i,t} = (1 - \frac{\mathcal{C}^{i,t}(x_0)}{100})\bar{n}^{i,t,\sigma_\pi=0}(x_t)] \tag{6.1}$$

that is, for each type and conditional on the initial state x_0, $\mathcal{C}^{i,t}(x_0)$ is the value such that an agent is indifferent between being born in the economy with inflation risk, or being born in the economy without inflation risk, but $\mathcal{C}^{i,t}(x_0)$ percent less consumption in all states.[15] Given the homotheticity in preferences, this can also be interpreted as saying that the agent is willing to pay $\mathcal{C}^{i,t}(x_0)$ percent of her consumption in all states in order to get rid of inflation risk.

In order to obtain a formula in terms of equilibrium functions, denote the number of trees that an agent of type (i,t) is born with as $s^{i,t}$.[16] The total wealth at birth is $y_t\bar{n}^{i,t}(x_t) = y_t s^{i,t}\tilde{q}(x_t)$, i.e the number of trees times its price. Thus, using (4.8) to evaluate utility at birth, we can compute the expected utility at time 0 of an agent that will be born at time t conditional on the initial state as:

$$U^{i,t}(x,y) = \mathbb{E}\left[\mathcal{U}^{i,t}dt|x_0 = x, y_0 = y\right] = \mathbb{E}\left[\frac{(s^{i,t}y_t\tilde{q}(x_t))^{1-\gamma_i}}{1-\gamma_i}(V^i(x_t))^{\frac{1-\gamma_i}{1-\psi_i}}|x_0 = x, y_0 = y\right]$$

Using homogeneity we can write:

$$\begin{aligned} U^{i,t}(x,y) &= \frac{(s^{i,t}y)^{1-\gamma_i}}{1-\gamma_i}u^i(x,\tau) \\ u^i(x_t,\tau) &= \mathbb{E}\left[(\frac{y_{t+\tau}}{y_t}\tilde{q}(x_{t+\tau}))^{1-\gamma^i}(V^i(x_{t+\tau}))^{\frac{1-\gamma^i}{1-\psi^i}}|x_t\right] \end{aligned} \tag{6.2}$$

and Appendix D.2 shows that $u^i(x_t,\tau)$ can be found by solving a PDE. Given this function, and noting that both economies feature the same $\frac{(s^{i,t}y)^{1-\gamma_i}}{1-\gamma_i}$, we can compute $\mathcal{C}^{i,t}(x_0)$ as

$$\mathcal{C}^{i,t}(x) = \left[1 - \left(\frac{u^i(x,t)}{u^{i,\sigma_\pi=0}(x,t)}\right)^{1/(1-\gamma_i)}\right] \times 100. \tag{6.3}$$

Results. Figure 6.1 plots $\mathcal{C}^{i,t}(x_0)$ for $i = A$ (blue) and B (red), with different cohorts on the horizontal axis, and for different initial states.[17] The welfare effects of eliminating inflation risk for these agents are sizable: for initially alive cohorts, type A (B) agents value

[15]Note that this is using that the correct way of valuing the utility at time 0 of an agent that is going to born at $t > 0$ is by taking the expected value of its $\mathcal{U}^{i,t}$ at birth. Appendix D shows that this is correct under our assumptions about the form of the aggregator.

[16]The discussion in Section 4 implies that $s^{A,0} = \frac{\alpha_0}{\bar{\alpha}}$ and $s^{B,0} = \frac{1-\alpha_0}{1-\bar{\alpha}}$ for initially alive agents, and $s^{i,t} = \frac{\delta_q}{\varphi}$ for $i = A, B$ with $t > 0$.

[17]The figure fixes $\pi_0 = 2\%$ and only varies α_0. Varying π_0 has insignificant effects on the required compensations.

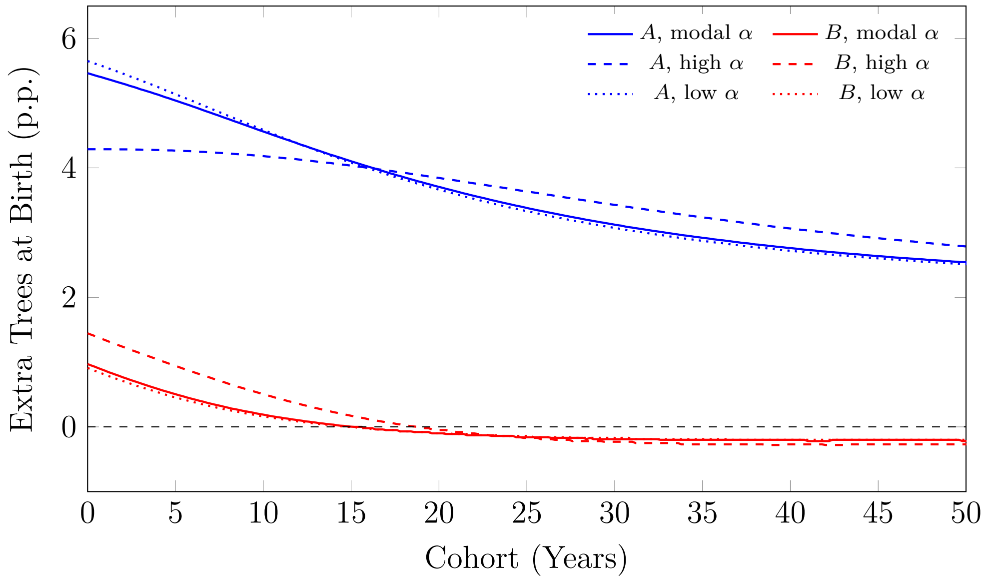

Figure 6.1: Willingness to pay to fully eliminate inflation risk, in units of proportional consumption reduction in all states, as a function of the cohort of birth.

eliminating inflation risk as much as having 5.5% (1%) higher consumption in all states. The Figure also shows that the required compensation varies significantly by both preferences and cohort. First, risk-tolerant agents require much higher compensation than risk averse agents. 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 that, in general equilibrium, investment opportunities and optimal portfolios change when $\sigma_\pi = 0$.

In particular, Figure 6.2 shows the optimal portfolios, consumption policies, as well as the implied drift and volatility of the consumption process, and welfare at birth, as a function of α in both economies.[18] Focus first on type A agents. Given that leveraging is cheaper in the economy without inflation risk (since they do not need to pay the inflation risk premium) and expected returns are higher, type A agents invest more in trees and consume less. This is intuitive: the expected return on trees is higher, so they have incentives to shift their portfolio towards it. Since these agents have high EIS, it is optimal for them to consume less of their wealth, since their savings have higher returns. These optimal decisions have direct implications for the consumption process for type A agents: given that expected returns are higher and they save more, their consumption growth is higher. But since they load more on risky trees, their consumption is more volatile. Given that their elasticity of intertemporal substitution is high and their risk aversion is low, the force of higher drift dominates, and

[18] Appendix C.5 describes how to compute the drift and volatility of the consumption process from the solution to the equilibrium.

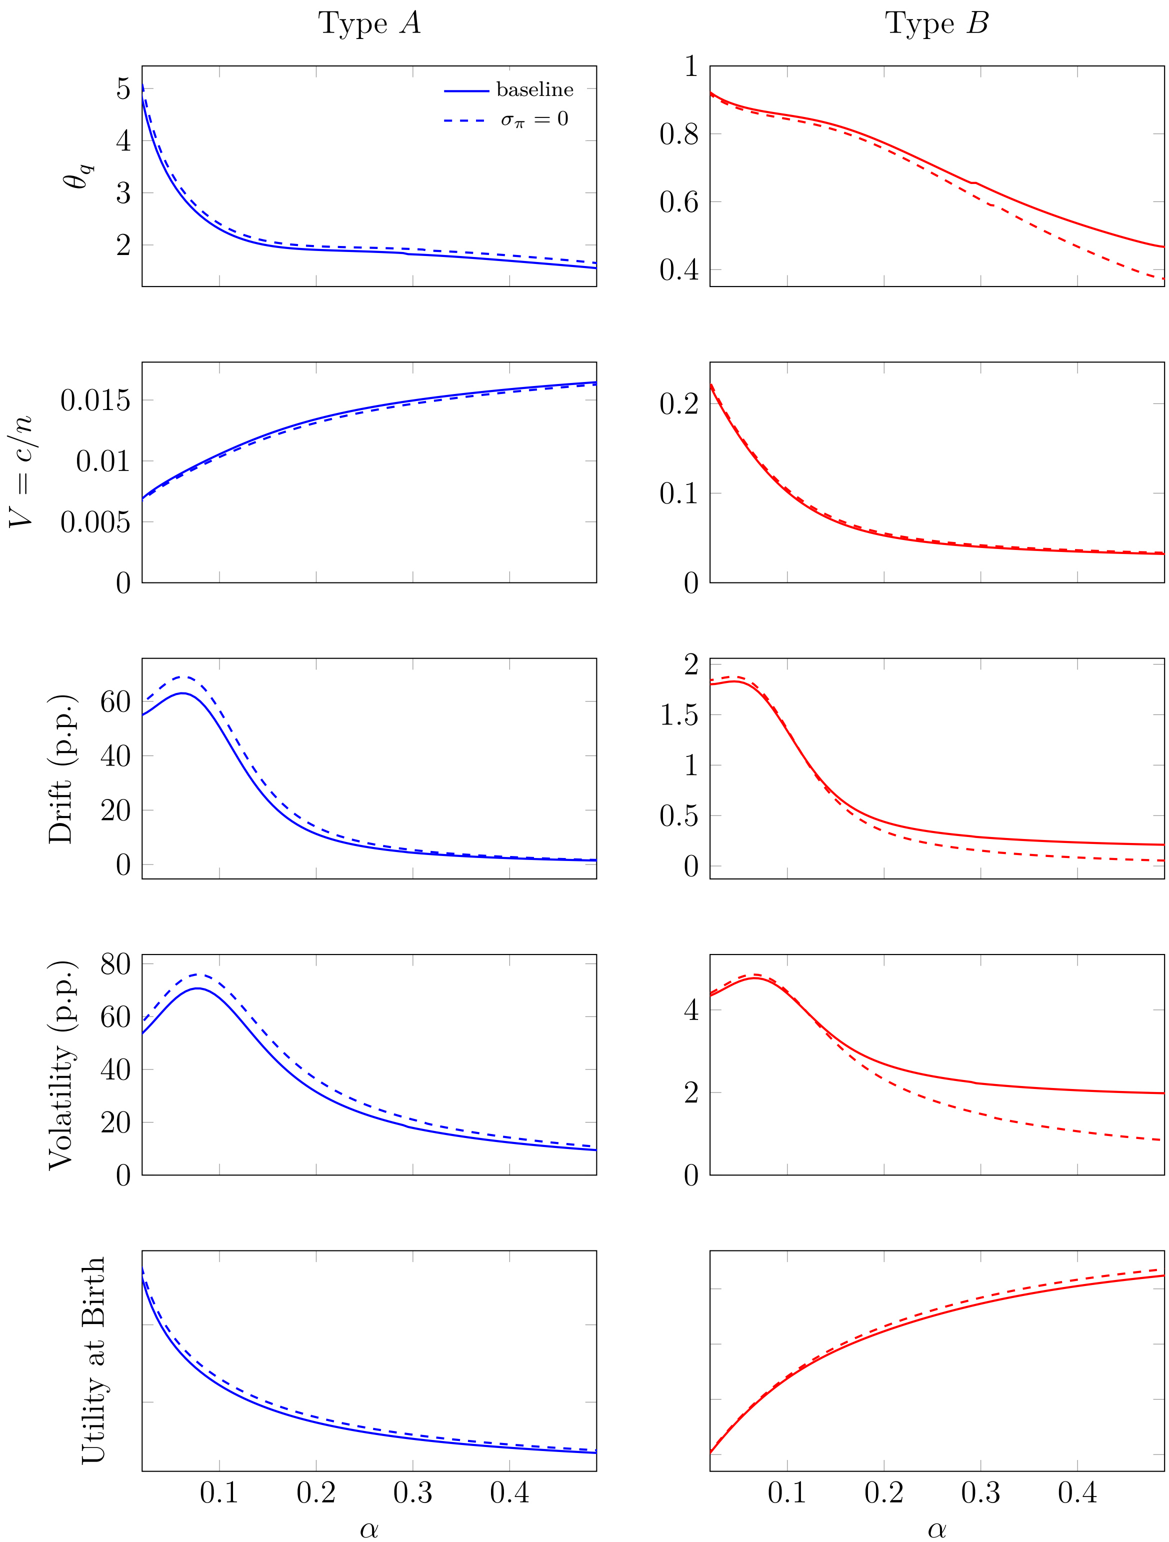

Figure 6.2: First Row: Portfolio shares for trees, θ_q^i. Second: consumption to wealth ratio, V^i. Third: Drift of consumption the process μ_c^i. Fourth: Instantaneous volatility of the consumption process, $\sqrt{(\sigma_c^i)^T(\sigma_c^i)}$. Last: Utility at birth $U^{i,0}(\alpha) = \frac{q(\alpha)^{1-\gamma_i}}{1-\gamma_i}\tilde{V}_i(\alpha)^{\frac{1-\gamma_i}{1-\psi_i}}$. In all cases, we fix α and average over π. Blue: Agent A. Red: agent B. Solid: baseline economy with inflation risk. Dashed: no inflation risk.

their welfare increases significantly upon elimination of inflation risk. This is shown in the last panel, where we can see that utility at birth as a function of α shifts up. This curve is decreasing in α because investment opportunities are more favorable to individual type A agents when they are relatively scarce in terms of wealth.

Turning to type B agents, they buy significantly less trees in economy without inflation risk. Intuitively, now that bonds are not exposed to inflation risk, risk-averse agents optimally shift their portfolio towards bonds. Since real rates are higher and they have low EIS, they consume a larger fraction of their wealth. Turning to the implications for their consumption process, we see that for most values of α, they enjoy lower consumption volatility, but also a lower drift. This is reversed for very low α, where both the drift and volatility are higher. The reason is that, at low α, type B agents are essentially forced to hold most of their wealth in trees, which have higher expected returns but are more volatile. Thus, their consumption process reflects that. We can also see that their utility at birth shifts up, although the effect is more noticeable for intermediate to high α. This curve is increasing for the same reason as before: when type A agents hold more wealth, prices are more favorable to individual type B agents, and thus their welfare at birth increases.

Another salient feature of Figure 6.1 is that welfare gains are concentrated among initially alive cohorts: the gain is decreasing in t. This is mainly for two reasons. First, the stationary distribution without inflation risk features higher α. Thus, if we start at the modal α of the baseline distribution, there is an initial transition towards the higher α associated with the economy without inflation risk. Since individual type A agents prefer lower α, initially born agents of type A can take advantage of the initially lower α and reap higher welfare gains. This also explains why the agents for type A agents are larger when the initial alpha is lower (dotted line), and vice versa for high initial α (dashed). Given that type B agents prefer high α, the opposite pattern emerges when comparing initial values: gains are larger when the starting α is high. However, the term-structure of type B agents still shows a decreasing pattern. This has to do with the second reason, which is the weights given to different events in the future. Note that, in (6.2), the value of the function $(\tilde{q})^{1-\gamma}V^{\frac{1-\gamma}{1-\psi}}$ is weighted by a $(\frac{y_{t+\tau}}{y_t})^{1-\gamma}$ term. For the risk averse agent $\gamma > 1$, this is effectively weighting more the states where $y_{t+\tau}$ is low, which tend to be correlated with low α. Furthermore, very low α states are (marginally) more likely to occur in the economy with no inflation risk, as shown by comparing the left tail of the stationary distributions in Figure 5.2. This effect is specially strong for type B agents (since γ is larger), and it gets stronger for later cohorts: the variance of $(\frac{y_{t+\tau}}{y_t})^{1-\gamma}$ grows over time. This effect is so strong that it can lead to small *negative* compensations required for later cohorts.[19] Intuitively, eliminating inflation risk marginally increases the change of large market crashes due to increases in leverage, and for risk averse agents that will born many years into the future, these states dominate in terms of their welfare evaluation. For type A agents, since $\gamma < 1$, they weight events with higher aggregate output (and thus higher α) more. Since their utility at birth is decreasing in α, this weighting also reduces their welfare as we consider cohorts to be born further away in the future.

In summary, eliminating inflation risk increases welfare for all initially alive agents, as well as for earlier cohorts. Welfare gains are much larger for type A agents, and to earlier cohorts. The magnitude of the effects is large at the individual level. However, given the nature of the exercise, we cannot aggregate these welfare gains into an aggregate efficiency number, which we can compare to the losses due to other standard costs of inflation. We turn to the aggregation issue next.

6.2 Towards an Aggregate Welfare Gain

Measuring aggregate welfare gains with heterogeneous agents is notoriously hard. In order to make progress, we use the recent approach proposed by Dávila and Schaab (2025); Barcons et al. (2026). In particular, consider an utilitarian social welfare function (SWF) $\mathcal{W}$, of the form:

$$\mathcal{W}(x,y) = \lim_{T\to+\infty} \int \int_0^T w^{i,t} U^{i,t}(x,y) dt di \tag{6.4}$$

where the welfare weights $w^{i,t}$ are such that $\mathcal{W}(x,y) < +\infty$ as $T \to \infty$. As it is well understood, the problem with assuming a specific SWF is that the answer to any aggregate welfare question will depend on the welfare weights used.

In this setup, Barcons et al. (2026) show that, if one is willing to consider marginal perturbations to the allocation of interest and choose a welfare numeraire (more on this below), then one can decompose the marginal change in the SWF ($\frac{\partial \mathcal{W}}{\partial \sigma_\pi}$) in two terms, efficiency and redistribution, as follows:

[19] One way to see the strength of this effect to mechanically compute $\tilde{\mathcal{C}}^{i,t}(x) = 1 - \left(\frac{\mathbb{E}[u^i(x_t,0)|x_0 = x]}{\mathbb{E}[u^{i,\sigma_\pi=0}(x_t,0)|x_0 = x]}\right)^{1/(1-\gamma_i)}$. and consider the limit where $t \to +\infty$. In this case, we are essentially computing the expected value of $u^i(x_t,0)$ with and without inflation risk, using the corresponding stationary distributions. This calculation differs from (6.3) in that it ignores the $(\frac{y_{t+\tau}}{y_t})^{1-\gamma}$ term. In this case, both agents have positive long-term gains.

$$\frac{d\mathcal{W}(x,y)}{d\sigma_\pi} = \Xi^E(x,y) + \Xi^{RD}(x,y), \quad \Xi^E(x,y) = \lim_{T\to+\infty} \int \int_0^T \frac{dU^{i,t}(x,y)/d\sigma_\pi}{\lambda^{i,t}(x,y)} dtdi. \tag{6.5}$$

Where $\lambda^{i,t}(x,y)$ captures the value in utils of the welfare numeraire. In particular, the welfare numeraire is a bundle of goods at each date and state, which must be chosen such that all agents value it. We follow Barcons et al. (2026) and let the welfare numeraire be a bundle that contains aggregate output at each date and state. Therefore:

$$\lambda^{i,t}(x,y) = \mathbb{E}\left[\int_0^\infty M^{i,t}_{t+h} y_{t+h} dh | x_0 = x, y_0 = y\right] \tag{6.6}$$

where $M^{i,t}_{t+h} = \mathbf{1}_{h\geq 0} \exp(\int_0^h f_{\mathcal{U}}(c^{i,t}_{t+s}, \mathcal{U}^{i,t}_{t+s})ds) f_{\mathcal{U}}(c^{i,t}_{t+h}, \mathcal{U}^{i,t}_{t+h})$ is the marginal utility for an agent of type i,t of additional consumption at date $t+h$ in a given state of the world. Once we divide the change in utility by $\lambda^{i,t}(x,y)$ for all agents, then all the $\frac{dU^{i,t}/d\sigma_\pi}{\lambda^{i,t}}$ are expressed in the same units (welfare numeraire over inflation volatility), and thus we can sum them across agents. The interpretation is also straight forward: if $\Xi^E = -2$, then reducing inflation volatility by 0.01 is equivalent to increasing aggregate output in all periods and states by $2 \times 0.01 = 2\%$ and distributing that uniformly across agents. Crucially, the efficiency term *does not* depend on the welfare weights, so we focus on it as our measure of aggregate welfare gains.

A problem of computing Ξ_E is that the integral is not discounted, so it will generally diverge. This makes sense: we have an infinite number of agents, so we considered a perturbation that increased welfare of all agents cohorts by 0.01, then $\Xi_E = \lim_{T\to+\infty} 0.01 \times T = +\infty$. For this reason, we focus on two objects. First, we compute the cumulative sums of the efficiency terms for the first T cohorts. Noting that a share of aggregate output $1-\omega$ goes to hand-to-mouth agents that are not affected by the perturbation, only type A and B agents are affected, and thus the cumulative sum up to cohort T is:

$$\Xi^E(x,y,T) = \sum_{i=A,B} \bar{\alpha}_i \left[\underbrace{\frac{dU^{i,0}(x,y)/d\sigma_\pi}{\lambda^{i,0}(x,y)}}_{\text{Initially Alive}} + \underbrace{\varphi \int_0^T \frac{dU^{i,t}(x,y)/d\sigma_\pi}{\lambda^{i,t}(x,y)} dt}_{\text{Future Cohorts}}\right] \tag{6.7}$$

where $\bar{\alpha}_A = \bar{\alpha}$ and $\bar{\alpha}_B = 1-\bar{\alpha}$ are the shares of agents born of each type. For any given fixed T, this yields the efficiency gains if we only focused on the initially alive agents, and agents born until period T. The second measure focuses on the long run, in particular, we normalize the total Ξ_E by the number of agents, that is $\bar{\Xi}^E(x,y) = \lim_{T\to+\infty} \frac{\Xi^E(x,y,T)}{1+\varphi T}$. Note that if $\frac{dU^{i,t}(x,y)/d\sigma_\pi}{\lambda^{i,t}(x,y)}$ converges as $t \to +\infty$, then

$$\bar{\Xi}^E(x,y) = \sum_{i=A,B} \bar{\alpha}_i \frac{dU^{i,\infty}(x,y)/d\sigma_\pi}{\lambda^{i,\infty}(x,y)} \tag{6.8}$$

where $\frac{dU^{i,\infty}/d\sigma_\pi}{\lambda^{i,\infty}(x,y)} = \lim_{t\to+\infty} \frac{dU^{i,t}/d\sigma_\pi}{\lambda^{i,t}(x,y)}$. We refer to this measure as the long run welfare change. Note that this measure is in "per agent" terms. Thus, if $\bar{\Xi}^E(x,y) = -1$, then reducing inflation risk in 0.01 will increase welfare of *each* agent as if we gave those agents 1% extra units of aggregate output.

Computing the Key Terms. From the above discussion, we need to compute $dU^{i,t}/d\sigma_\pi$ and $\lambda^{i,t}(x,y)$ for each i,t. We approximate $dU^{i,t}/d\sigma_\pi$ using finite differences: we solve the equilibrium for close values of σ_π, then compute the u function defined in (6.2), and thus approximate:

$$dU^{i,t}/d\sigma_\pi \approx \frac{(s^{i,t}y)^{1-\gamma_i}}{1-\gamma_i} \frac{\Delta u^i(x,\tau)}{\Delta\sigma_\pi} \tag{6.9}$$

We can compute $\lambda(x,y)$ as:

$$\begin{aligned}\lambda^{i,t}(x,y) &= \mathbb{E}\left[\int_t^\infty M_s^i y_s^i ds | x_0 = x, y_0 = y\right] \\ &= (s^{i,t})^{-\gamma_i} y^{1-\gamma_i} \ell^i(x_t,\tau)\end{aligned}$$

where $\ell^i(x_t,\tau) = \mathbb{E}\left[(\frac{y_t}{y_0})^{1-\gamma_i}(\tilde{q}(x_t))^{-\gamma_i}(V(x_t))^{\frac{1-\gamma_i}{1-\psi_i}} q^{a,i}(x_t) | x_0 = x, y_0 = y\right]$ and $q^{a,i}(x_t)$ is the value of a claim on aggregate output for an agent of type i born today if $y_0 = 1$, i.e $q^{a,i}(x) = \mathbb{E}\left[\int_t^\infty \frac{M_s^i}{M_t^i}\frac{y_s}{y_t} ds | x_t = x\right]$. We also used that marginal utility at birth is given by $M_t^i = (y_t \bar{n}^{i,t}(x_t))^{-\gamma_i} V^{\frac{1-\gamma_i}{1-\psi_i}}$. Appendix D.3 shows how to compute $\ell^i(x_t,\tau)$ by solving a PDE. With these two, the welfare change of a reducing inflation risk for an agent of type i,t is:

$$\frac{dU^{i,t}/d\sigma_\pi}{\lambda^{i,t}}(x) = \frac{s^{i,t}}{1-\gamma_i} \frac{\frac{\Delta u^i(x,\tau)}{\Delta\sigma_\pi}}{\ell^i(x_t,t)} \tag{6.10}$$

which is independent of the initial y_t by homotheticity. Once we have those terms, we can sum them as dictated by equation (6.7) to obtain the efficiency gains for the first set of cohorts, and then compute the long run value by using (6.8). We then approximate the gain of completely eliminating inflation risk by multiplying these values by $(-\sigma_\pi)$.

Results. Table 3 presents the welfare gains for the first 50 cohorts, and then in the long run (i.e for a cohort to be born far out in the future). If we focus on the aggregate efficiency term for the first 50 cohorts, eliminating inflation risk is equivalent to a 0.45% perpetual

	At modal $(\alpha, \bar{\pi})$	Integrated over (α, π)
First 50 cohorts	0.451	0.441
A	12.886	12.579
B	0.132	0.130
Initially alive	0.425	0.412
A	11.979	11.677
B	0.128	0.123
Rest (future cohorts)	0.026	0.029
A	0.907	0.902
B	0.004	0.007
Long-run $(t \to \infty)$	-0.033	-0.033
A	0.572	0.572
B	-0.048	-0.048

Table 3: Aggregate efficiency gains $\Xi^E(x, T)$ for $T = 50$ years, and long run gains $\bar{\Xi}^E(x)$, multiplied by the total inflation risk $(-\sigma_\pi)$. This gives an approximation of the aggregate welfare gain of eliminating inflation risk. Units: percentage points of perpetual aggregate output. The A and B rows indicate the average gains corresponding to each agent of that type, such that the aggregate gains are computed as $\bar{\alpha}$(Gains from A) + $(1 - \bar{\alpha})$(Gains from B). The left column evaluates the gains at the mode of the distribution, whereas the right column averages across initial conditions using the stationary distribution of the baseline economy.

increase in aggregate GDP, which is a non-trivial magnitude, especially compared to other typical costs of inflation. The table also decomposes those gains across different agents: the initially alive generation, and cohorts to be born during the first 50 years, as well as by preference type. As we can see, most of the gains are driven by type A agents, and the initially alive generation, which is consistent with Figure 6.1.

The magnitude of the split between initially alive type A agents and later born ones deserve some discussion. In particular, while in cohorts born in $t > 0$ both types of agents have the same wealth, this is not true for the initially alive agents. In our baseline calibration, type A agents represent 2.5% of the population, but own around 30% of the wealth. Since the eliminating inflation risk benefits each individual type A agent greatly, a relatively large amount of aggregate output is required to compensate them. In contrast, once we consider unborn cohorts, type A agents only represent 2.5% of the wealth, so compensating them is much cheaper.[20]

[20]We can see this from equations (6.7) and (6.10). Using that given α each type A agent starts with $s^{i,0} = \frac{\alpha_0}{\bar{\alpha}}$ trees, the efficiency component that comes from the initially alive cohort is $\alpha \frac{\frac{\Delta u^A(x,\tau)}{\Delta \sigma_\pi}}{(1-\gamma_A)\ell^A(x_0,0)} + (1-\alpha)\frac{\frac{\Delta u^B(x,\tau)}{\Delta \sigma_\pi}}{(1-\gamma_B)\ell^B(x_0,0)}$, while if we consider a cohort born at time $t = dt$, we have a sum of two similar terms, but where the weights are given by $\bar{\alpha}$.

Table 3 also shows that eliminating inflation risk produces small long run *losses*. Mechanically, type B agents to be born many years in the future lose due to the reasons explained above. Since they are the overwhelming majority of agents being born (and agents of both types are born with the same initial wealth), the gains of type A agents are not enough to compensate them, and the long run efficiency gains are negative. This may sound puzzling, since under $\sigma_\pi = 0$, both agents face dynamically complete markets since the moment of birth. However, due to the Overlapping Generations structure of the model, markets are still incomplete, because agents cannot trade before they are born. If they could, type A agents could ex-ante insure type B agents against the risk of being born in a state with low y_t, and get higher expected y_t in return. As it is well understood, once we have incomplete markets, adding another distortion ($\sigma_\pi > 0$) can increase welfare. However, the losses are quite small: it would take around 700 years for these losses to overcome the gain of the initially alive agents, and they rely on the way in which we model welfare at time 0 of agents that will be born hundreds of years in the future. Furthermore, Gârleanu and Panageas (2015) introduce an Overlapping Generations structure mainly to ensure a stationary wealth distribution. Thus, for practical purposes, we will consider these to be essentially zero.

6.3 Discussion

Comparison with Back of the Envelope Calculation. It is interesting to note that the estimated figure, of almost 0.45% of GDP, is around twice as large as the back of the envelope estimate in Section 2, justified by focusing on the direct effect in Section 3. Intuitively, once agents are able to re optimize, eliminating inflation risk has additional benefits for welfare, because it improves risk-sharing among living agents: now risk-averse agents are less exposed to aggregate output risk, which increases their welfare.[21]

Sources of Indirect Gains. As explained in Section 3, the size of the gains that come from general equilibrium forces are inherently tied to the particular ingredients we use. In our model, the main driver of nominal trade is sharing aggregate output risk. In reality, there are several other potential reasons why agents will hold nominal positions, such as self-insuring against idiosyncratic shocks, saving due to life cycle motives, or buying some indivisible asset like a house or a car. Since inflation risk will also harm financial trade in those cases, one would conjecture that the broad mechanism will still apply: once inflation risk is removed, financial trade will increase, and that would lead to higher welfare above

[21]However, as explained above, it also marginally increases the likelihood of large stock market crashes, which are important when evaluating the ex-ante welfare of agents to be born many years in the future.

and beyond the direct effect of reducing risk on existing positions. A deeper dive on the implications of alternative reasons for trade is left for future research.

7 Comparison with Standard Sources of Inflation Costs

In this section, we compare the estimated cost to the main sources of inflation costs in standard macroeconomic models: transaction and pricing frictions. Note that the typical experiment in this literature is to increase trend inflation from 0 to 5 or 10% per year. Given the evolution of inflation in developed economies in the last 40 years, where inflation was relatively stable around 2% for most of the time, this exercise overestimates the costs in sample, since it extrapolates trend inflation to values seldom observed in sample. In spite of this, we find that the each of these costs of trend inflation are estimated on the same magnitude of our cost of inflation risk.

Transaction Frictions. The classical cost of trend inflation relates to its distortion effects in money holdings (Bailey, 1956; Friedman, 1969): if agents must hold money to transact, but the real return on this asset is lower due to inflation, then trend inflation introduces a distortion and thus lowers welfare.[22] Although there is disagreement on its exact magnitude, the estimated costs through this channel can be substantial. In particular, Lucas (2000) finds that reducing inflation from 10% to 0% is equivalent to a gain of 1% in real income. In contrast, Ireland (2009) finds the gains to be much smaller, of less than 0.25%, and moving trend inflation from 0 to 2% is essentially costless. As explained in Benati and Nicolini (2025), the key difference is the behavior of money demand at low inflation. After systematically reviewing the evidence, this paper presents a benchmark number of 0.35% of GDP for a 5% trend inflation. This is similar in magnitude to our baseline estimates for the welfare cost of inflation risk.

Price-Setting Frictions. The standard source of costs of inflation in New Keynesian models relates to pricing frictions. These costs come from two sources: i) increased price dispersion, which distorts production because firms with the same productivity may charge

[22] A related paper within the transaction frictions literature is Allais et al. (2020): 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 two channels: it increases the risk of lending in nominal terms, and it also it distorts risk-sharing of aggregate shocks.

different prices, which leads to an inefficient allocation of resources across firms, and ii) the costs (if any) of changing prices.

In this case, it is important to distinguish the cost of pricing frictions in steady state from how this cost increases with inflation. In models with idiosyncratic shocks (which are the norm in modern sticky-price models), even in the absence of trend inflation, there is a cost of price dispersion, which comes from the fact that firms do not adjust to their idiosyncratic shocks perfectly every period. We focus on the change in this cost due to increasing trend inflation.[23]

Important contributions include Burstein and Hellwig (2008); Nakamura et al. (2018). They show that the answer to this question crucially depends on the pricing frictions assumed. Under Calvo (1983), price dispersion increases steeply with trend inflation. Nakamura et al. (2018) report that increasing trend inflation from 0 to 12% in these models is roughly 10% of GDP, which is implausibly large. They then show that, in a calibrated menu cost model, price dispersion is essentially insensitive to inflation in the ranges considered. In a recent paper, Cavallo et al. (2023) use price microdata to calibrate a random menu cost model, which can be seen as an intermediate between the polar cases of Calvo and pure Menu Costs. In their model, increasing trend inflation from 0 to 10% has a welfare cost of around 0.4% of GDP in steady state. This is in the same order of magnitude of our estimated cost of inflation risk.

8 Final Remarks

We show that inflation volatility is socially costly because it exposes agents to risk due to existing nominal positions, and distorts financial trade. This cost is related to two moments of the data: the size of nominal positions and the inflation risk premium. We quantify it in a general equilibrium model, and find to be of a similar order of magnitude to standard costs of trend inflation. We see our paper only as a first approximation to the true magnitude of this cost. Future research can improve our estimates among several dimensions.

On measurement, a more disaggregated analysis of nominal positions of different agents at different maturities is a natural next step, since as we showed, these are key determinants of the magnitude of this cost. Our paper also shows that the magnitude and sources of inflation risk premium, which has been studied primarily in the finance literature, are also relevant for the macroeconomics literature on welfare costs of inflation. This makes additional empirical evidence on it more valuable.

[23] In our model, inflation risk is purely driven by monetary policy, so it is internally consistent to compare this cost to the full cost of eliminating inflation risk.

We also recognize that the stylized nature of our model could be important for the quantification, especially for the part of the welfare gains coming through general equilibrium channels. Future work may add some other features that potentially drive borrowing and lending in reality, such as life-cycle motives or buying large, indivisible assets. In this case, the key challenge is how to integrate these features in a model that matches key asset pricing moments, as the one we presented here.

Finally, we modeled the absence of inflation indexation as an exogenous friction, and showed that welfare consequences of it are non-trivial. A natural question is whether the magnitude of the welfare costs depend on the exact frictions that generate it. By analogy to the literature on other costs of inflation, we think the answer is far from trivial. In the case of transaction frictions, the area below the money demand curve provides a good approximation of their welfare cost, regardless of the underlying microfoundations for it (Lucas, 2000; Alvarez et al., 2019).[24] In stark contrast, a central finding of the literature on price-setting frictions is that the underlying microfoundations are a crucial determinant of the the welfare costs of inflation through this channel Burstein and Hellwig (2008); Nakamura et al. (2018).

References

Ajello, A., L. Benzoni, and O. Chyruk (2019): "Core and 'Crust': Consumer Prices and the Term Structure of Interest Rates," *The Review of Financial Studies*, 33, 3719–3765.

Allais, O., Y. Algan, E. Challe, and X. Ragot (2020): "The Welfare Cost of Inflation Risk under Imperfect Insurance," *Annals of Economics and Statistics*, 1–20.

Alvarez, F. and U. J. Jermann (2004): "Using Asset Prices to Measure the Welfare Cost of Business Cycles," *American Economic Review*, 94, 1223–1256.

Alvarez, F., F. Lippi, and R. Robatto (2019): "Cost of inflation in inventory theoretical models," *Review of Economic Dynamics*, 32, 206–226.

Angeletos, G.-M., F. Collard, and H. Dellas (2020): "Business-Cycle Anatomy," *American Economic Review*, 110, 3030–3070.

[24] Lucas (2000) considers a money-in-utility model, and a shopping time model, and Alvarez et al. (2019) extend this to inventory-theoretical models.

Auclert, A. (2019): "Monetary Policy and the Redistribution Channel," *American Economic Review*, 109, 2333–2370.

Auclert, A., M. Rognlie, and L. Straub (2020): "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," NBER Working Paper 26647, National Bureau of Economic Research.

Bailey, M. J. (1956): "The Welfare Cost of Inflationary Finance," *Journal of Political Economy*, 64, 93–110.

Bansal, R. and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *The Journal of Finance*, 59, 1481–1509.

Baqaee, D. R. and A. T. Burstein (2025): "Welfare and Output with Income Effects and Taste Shocks," *American Economic Review*, 115, 1–46.

Barcons, S., E. Dávila, and A. Schaab (2026): "Intergenerational Welfare Assessments," NBER Working Paper 34616, National Bureau of Economic Research.

Baumol, W. J. (1952): "The Transactions Demand for Cash: An Inventory Theoretic Approach," *Quarterly Journal of Economics*, 66, 545–556.

Benati, L. and J. P. Nicolini (2025): "The Welfare Costs of Inflation Reconsidered," Staff Report 675, Federal Reserve Bank of Minneapolis.

Bianchi, F., R. Faccini, and L. Melosi (2023): "A Fiscal Theory of Persistent Inflation," *The Quarterly Journal of Economics*, 138, 2127–2179.

Brunnermeier, M. K. and Y. Sannikov (2014): "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 104, 379–421.

Burstein, A. and C. Hellwig (2008): "Welfare Costs of Inflation in a Menu Cost Model," *American Economic Review*, 98, 438–43.

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

Calvo, G. A. (1983): "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12, 383–398.

Campbell, J. Y. and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.

Campbell, J. Y., C. Pflueger, and L. M. Viceira (2020): "Macroeconomic Drivers of Bond and Equity Risks," *Journal of Political Economy*, 128, 3148–3185.

Cavallo, A., F. Lippi, and K. Miyahara (2023): "Inflation and Misallocation in New Keynesian Models," in *Proceedings of the ECB Forum on Central Banking 2023*, Sintra, Portugal: European Central Bank, first draft: May 2023.

Cieslak, A. and C. Pflueger (2023): "Inflation and Asset Returns," *Annual Review of Financial Economics*, 15, 433–448.

Dávila, E. and A. Schaab (2025): "Welfare Assessments with Heterogeneous Individuals," *Journal of Political Economy*, 133, 2918–2961.

Doepke, M. and M. Schneider (2006a): "Aggregate Implications of Wealth Redistribution: The Case of Inflation," *Journal of the European Economic Association*, 4, 493–502.

——— (2006b): "Inflation and the Redistribution of Nominal Wealth," *Journal of Political Economy*, 114.

——— (2017): "Money as a Unit of Account," *Econometrica*, 85, 1537–1574.

Duffie, D. and L. G. Epstein (1992): "Stochastic Differential Utility," *Econometrica*, 60, 353–394.

Fang, X., Y. Liu, and N. Roussanov (2025): "Getting to the Core: Inflation Risks Within and Across Asset Classes," *The Review of Financial Studies*, hhaf050, advance article, published 16 July 2025.

Fischer, S. (1975): "The Demand for Index Bonds," *Journal of Political Economy*, 83, 509–534.

Fisher, S. and F. Modigliani (1978): "Towards an Understanding of the Real Effects and Costs of Inflation," *Weltwirtschaftliches Archiv*, 114, 810–833.

Fleckenstein, M., F. A. Longstaff, and H. Lustig (2017): "Deflation Risk," *The Review of Financial Studies*, 30, 2719–2760.

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

——— (1977): "Nobel Lecture: Inflation and Unemployment," *Journal of Political Economy*, 85, 451–472.

Gabaix, X. (2025): “A Theory of Complexity Aversion,” Working paper.

Gârleanu, N. and S. Panageas (2015): “Young, Old, Conservative, and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing,” *Journal of Political Economy*, 123, 670–685.

Golosov, M. and R. E. Lucas (2007): “Menu Costs and Phillips Curves,” *Journal of Political Economy*, 115, 171–199.

Haubrich, J., G. Pennacchi, and P. Ritchken (2012): “Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps,” *The Review of Financial Studies*, 25, 1588–1629.

He, Z. and A. Krishnamurthy (2013): “Intermediary asset pricing,” *American Economic Review*, 103, 732–770.

Ireland, P. N. (2009): “On the Welfare Cost of Inflation and the Recent Behavior of Money Demand,” *American Economic Review*, 99, 1040–1052.

Lucas, R. E. (1978): “Asset Prices in an Exchange Economy,” *Econometrica*, 46, 1429–1445.

——— (2000): “Inflation and Welfare,” *Econometrica*, 68, 247–274.

Martin, I. (2017): “What Is the Expected Return on the Market?” *The Quarterly Journal of Economics*, 132, 367–433.

Nakamura, E., J. Steinsson, P. Sun, and D. Villar (2018): “The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation,” *Quarterly Journal of Economics*, 133, 1933–1980.

Panageas, S. (2020): “The implications of heterogeneity and inequality for asset pricing,” *Foundations and Trends in Finance*, 12, 199–275.

Pflueger, C. (2025): “Back to the 1980s or Not? The Drivers of Inflation and Real Risks in Treasury Bonds,” *Journal of Financial Economics*, 167, 104027.

Piazzesi, M. and M. Schneider (2006): “Equilibrium Yield Curves,” *NBER Macroeconomics Annual*, 21, 389–472.

Robert E. Lucas, J. (1987): *Models of Business Cycles*, Oxford: Basil Blackwell.

Rotemberg, J. J. (1982): "Sticky Prices in the United States," *Journal of Political Economy*, 90, 1187–1211.

Rudebusch, G. D. and E. T. Swanson (2012): "The Bond Premium in a DSGE Model with Inflation Targeting," *American Economic Review*, 102, 3111–3144.

Schneider, A. (2022): "Risk-Sharing and the Term Structure of Interest Rates," *Journal of Finance*, 77, 2331–2374.

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

Schröder, M. and C. Skiadas (1999): "Optimal Consumption and Portfolio Selection with Stochastic Differential Utility," *Journal of Economic Theory*, 89, 68–126.

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

Taylor, J. B. (1980): "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy*, 88, 1–23.

Tobin, J. (1956): "The Interest-Elasticity of Transactions Demand for Cash," *Review of Economics and Statistics*, 38, 241–247.

Wachter, J. A. (2006): "A consumption-based model of the term structure of interest rates," *Journal of Financial Economics*, 79, 365–399.

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

A Empirical Evidence

Figure A.1 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 et al., 2019; Fang et al., 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.

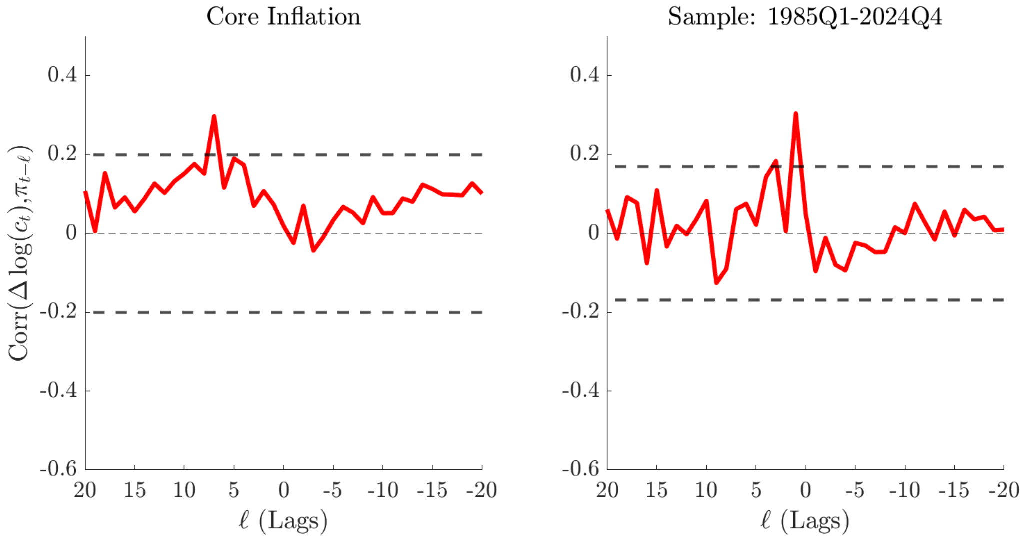

Figure A.1: Cross-correlogram between PCE real consumption growth and inflation. Left: use core CPI instead of headline CPI. Right: extend the sample to 2024Q4.Dashed lines are 95% confidence intervals.

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:

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

thus, for any a we can choose any P_0 by appropriately specifying a monetary policy rule of the form B.1. Similarly, for m_1 we have:

$$
\frac{M_1/a}{P_1} = \frac{1}{u_A'(c_{A,1})} + \frac{1}{u_A'(c_{B,1})} \tag{B.2}
$$

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 $\frac{\partial v_i}{\partial c_1}(c_{0,1}c_{i1})$ 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.[25] 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.

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

[25] 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.

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.

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

$$\begin{aligned}
0 = \max_{c^i, \theta^i} & \frac{1-\gamma^i}{1-1/\psi^i}\mathcal{U}^i \left[\left(\frac{c^i}{((1-\gamma^i)\mathcal{U}^i)^{\frac{1}{1-\gamma^i}}} \right)^{1-\frac{1}{\psi^i}} - (\rho + \varphi) \right] + \\
& + \mathcal{U}_n^i n \left(r + \left(\theta^i\right)^T (\mu - r) - \frac{c^i}{n} + \varphi \right) + \frac{1}{2}\mathcal{U}_{nn}^i n^2 \left(\theta^i\right)^T \Sigma\Sigma^T \theta^i \\
& + \left(\nabla_x \mathcal{U}^i\right)^T \mu_x + \frac{1}{2} Tr(\Sigma_x^T \mathcal{H}_x \mathcal{U}^i \Sigma_x) \\
& + n \left(\nabla_x \mathcal{U}_n^i\right)^T \Sigma_x \Sigma^T \theta^i \\
\text{s.t.} \quad & \mathbf{1}^T \times \theta^i \leq 1
\end{aligned} \tag{C.1}$$

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.

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

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

$$
c/n = V \tag{C.3}
$$

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.

$$
\sum \alpha^i (c/n)^i = \frac{1}{n}\sum c^i = \frac{\omega y}{q} = \frac{1}{\tilde{q}} \tag{C.4}
$$

The first order conditions with respect to porfolio shares read

$$
\mu_t - r - \gamma\Sigma\Sigma^T\theta + \frac{1-\gamma}{1-\psi}\left(\frac{(\nabla_x V)^T}{V}\Sigma_x\Sigma^T\right)^T = \gamma\begin{bmatrix}1\\1\end{bmatrix} \tag{C.5}
$$

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

$$d\alpha_t = \mu_{\alpha,t}dt + \sigma_{\alpha,t}^T dW_t \tag{C.6}$$

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

$$\frac{dn_t^A}{n_t^A} = (\theta_t^A)^T\left(\mu_t dt + \Sigma_t dW\right) - \frac{c_t^A}{n_t^A}dt + \delta_q\frac{\bar{\alpha}}{\alpha_t^A}dt \tag{C.7}$$

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

$$\frac{d(n_t^A + n_t^B)}{n_t^A + n_t^B} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}^T\left(\mu_t dt + \Sigma_t dW\right) - \alpha_t^T(c/n)_t dt + \delta dt. \tag{C.8}$$

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

$$\mu_{\alpha,t} = \alpha_t^A(1-\alpha_t^A)\left(\theta_q^T\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}\left(\mu - \Sigma_t(\Sigma_t)^T\begin{bmatrix} 1 \\ 0 \end{bmatrix}\right) - [1 - 1]\,(c/n)_t\right) + \delta_q\left(\bar{\alpha} - \alpha_t^A\right) \tag{C.9}$$

$$\sigma_{\alpha,t} = \alpha_t^A(1-\alpha_t^A)\theta_q^T\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}\Sigma_t \tag{C.10}$$

C.3 The law of motion of tree prices

Equipped 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 $dM = 0$.

$$\frac{d\tilde{q}}{\tilde{q}} = \left(\frac{\nabla \tilde{q}^T}{\tilde{q}} \mu_x + \frac{1}{2} \frac{Tr(\Sigma_x^T \mathcal{H} \tilde{q} \Sigma_x)}{\tilde{q}} \right) dt + \frac{\nabla \tilde{q}^T}{\tilde{q}} \Sigma_x dW \tag{C.11}$$

$$\frac{dq}{q} = \frac{d(\tilde{q}y)}{\tilde{q}y} = \left(\mu_{\tilde{q}} + \mu_y + \frac{\nabla \tilde{q}^T}{\tilde{q}} \Sigma_x \Sigma_z \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) dt + \frac{\nabla \tilde{q}^T}{\tilde{q}} \Sigma_x dW + \sigma_y dW_y. \tag{C.12}$$

Similarly for bond prices

$$\frac{dq^b}{q^b} = \left(\frac{\nabla q^{bT}}{q^b} \mu_x + \frac{1}{2} \frac{Tr(\Sigma_x^T \mathcal{H} q^b \Sigma_x)}{q^b} \right) dt + \frac{\nabla q^{bT}}{q^b} \Sigma_x dW. \tag{C.13}$$

From the definition of returns, we have that

$$\begin{aligned}
\mu_q &= \frac{\omega}{\tilde{q}} - \delta_q + \frac{\nabla \tilde{q}^T}{\tilde{q}} \mu_x + \frac{1}{2} \frac{Tr(\Sigma_x^T \mathcal{H} \tilde{q} \Sigma_x)}{\tilde{q}} + \mu_y + \frac{\nabla \tilde{q}^T}{\tilde{q}} \Sigma_x \Sigma_z \begin{bmatrix} 1 \\ 0 \end{bmatrix} \\
\sigma_q^T &= \frac{\nabla \tilde{q}^T}{\tilde{q}} \Sigma_x + \begin{bmatrix} 1 & 0 \end{bmatrix} \Sigma_z \\
\mu_b &= \frac{\delta_b}{q^b} + \left(\frac{\nabla q^{bT}}{q^b} \mu_x + \frac{1}{2} \frac{Tr(\Sigma_x^T \mathcal{H} q^b \Sigma_x)}{q^b} \right) - \delta_b - \pi \\
\sigma_b^T &= \frac{\nabla q^{bT}}{q^b} \Sigma_x
\end{aligned} \tag{C.14}$$

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

C.4 The Stochastic Discount Factor

From Schröder and Skiadas (1999), marginal utility in levels is given by:

$$M_t^i = \exp(\int_0^t f_{\mathcal{U}}^i ds) f_c \tag{C.15}$$

so applying Ito's lemma we obtain:

$$\frac{dM_t^i}{M_t^i} = f_{\mathcal{U}}^i dt + \frac{df_c}{f_c} \tag{C.16}$$

In order to compute $\frac{df_c}{f_c}$, it is useful to rewrite both partial derivatives as a function of V and n. In particular, note that:

$$f_c(c,\mathcal{U}) = c^{-1/\psi}[(1-\gamma)\mathcal{U}]^{\frac{1/\psi-\gamma}{1-\gamma}} \tag{C.17}$$

$$f_{\mathcal{U}}(c,\mathcal{U}) = \frac{1/\psi-\gamma}{1-1/\psi}\left[\frac{c}{((1-\gamma)\mathcal{U})^{1/(1-\gamma)}}\right]^{1-1/\psi} - (\rho+\varphi)\frac{1-\gamma}{1-1/\psi} \tag{C.18}$$

We can rewrite it in terms of V and n using C.3 and the functional form of the value function to get:

$$f_c = n^{-\gamma}V^{\frac{1-\gamma}{1-\psi}} \tag{C.19}$$

$$f_{\mathcal{U}} = \frac{1/\psi-\gamma}{1-1/\psi}(V-(\rho+\varphi)) - (\rho+\varphi) \tag{C.20}$$

note that $f_{\mathcal{U}}$ does not depend on n, only on the aggregate state. With that expression for f_c, we can apply Ito to obtain:

$$\frac{df_c}{f_c} = -\gamma\frac{dn}{n} + \frac{1-\gamma}{1-\psi}\frac{dV}{V} - \frac{1}{2}\gamma(-\gamma-1)(\frac{dn}{n})^2 + \frac{1}{2}\frac{1-\gamma}{1-\psi}(\frac{1-\gamma}{1-\psi}-1)(\frac{dV}{V})^2 - \gamma\frac{1-\gamma}{1-\psi}\frac{dn}{n}\frac{dV}{V} \tag{C.21}$$

where we have already derived the process for dn/n and dV/V as a function of equilibrium objects. Finally, note that this derivation computes the marginal utility process for each agent. This includes a discounting φ due to the probability of dying. As explained in the Appendix of Gârleanu and Panageas (2015), the SDF has the same volatility as the marginal utility process but the drift omits this φ term.

C.5 Consumption Process

In order to derive the consumption process for an agent born at time t, note that using its policy function, $c^{i,t}_{t+h} = V(x_{t+h})n^{i,t}_{t+h}$. Then $\frac{dc^{i,t}_{t+h}}{c^{i,t}_{t+h}} = \frac{dV}{V} + \frac{dn^{i,t}_{t+h}}{n^{i,t}_{t+h}} + (\frac{dn^{i,t}_{t+h}}{n^{i,t}_{t+h}})\frac{dV}{V}$. Let $\frac{dV}{V} = \mu_V(x)dt + \Sigma_V^T(x)dW$. Using the budget constraint we have that for all agents of type i:

$$\frac{dn^i}{n^i} = [r(x) + (\theta^i(x))^T(\mu(x)-r(x)) - V(x) + \varphi]dt + (\theta^i(x))^T\Sigma(x)dW$$

so consumption growth satisfies

$$\frac{dc^{i,t}_{t+h}}{c^{i,t}_{t+h}} = \mu^i_c(x)dt + (\Sigma^i_c(x))^T dW$$

where since $(\theta^i(x))^T\mathbf{1} = 1$

$$\begin{aligned}
\mu_c^i(x) &= (\theta^i(x))^T\mu(x) - V(x) + \varphi + \mu_V(x) + (\theta^i(x))^T\Sigma(x)\Sigma_V^T(x) \\
(\Sigma_c^i(x))^T &= (\theta^i(x))^T\Sigma(x) + \Sigma_V^T(x).
\end{aligned}$$

the objects $\mu, \Sigma, \theta^i,$, and V^i come directly from the solution to the equilibrium, and we can compute μ_V and Σ_V by applying Ito's lemma. Numerically, this implies computing the gradient and hessian of V, which is done using the same finite different procedures explained in C.6.

C.6 Numerical Solution

We solve the model with finite differences on a grid of 100 points for α and 25 for π, giving 2,500 nodes. At each node the unknowns are the two value functions V^A, V^B, the four entries of the return-volatility matrix Σ, and the long-bond price q_b, for 7 unknowns and 17,500 in total. The stock price $\tilde{q}$ is not an unknown: it follows from goods-market clearing, $\alpha c^A + (1-\alpha)c^B = \omega/\tilde{q}$.

Given a guess, we take gradients and Hessians with upwind differences, forward in a state when its drift is positive, backward when negative. For inflation this is immediate, since $\mu_\pi = \lambda_\pi(\bar{\pi} - \pi)$ is exogenous. For α the drift μ_α depends on the portfolio shares θ, which depend on $\nabla_\alpha V$ and hence on the difference direction. We solve θ both ways, evaluate μ_α under each, and blend the two one-sided gradients with a smooth weight that selects forward where $\mu_\alpha > 0$ and backward where $\mu_\alpha < 0$. The blend matters only in the narrow band where μ_α changes sign, where it removes the kink that would otherwise stall Newton's method.

With the drift known, Ito's lemma gives the return volatilities of $\tilde{q}$ and q_b, which we compare to the guessed Σ; we use the law of motion of the retuns to the long-term bond to pin down q^b, and together with the two HJB equations it yields 7 residuals per node, and a square system of 17,500 equations. We solve it by Newton's method, computing the Jacobian by automatic differentiation.

Finally, we obtain the stationary distribution $g(\alpha, \pi)$ as the zero-eigenvalue eigenvector of the discretized Kolmogorov-forward operator, and the asset-pricing and welfare objects of Appendix D by integrating their PDEs explicitly with step dt.

D Welfare Computations

D.1 Recursive Utility Before Birth

Consider the EZ preferences as specified in the Appendix of Gârleanu and Panageas (2015). Once we transform them as they do, we get their equation D.8:

$$\frac{1-\gamma}{\alpha}\left(\rho^{\frac{1-\gamma}{\alpha}}(1-\gamma)\mathcal{U}_t\right)^{\frac{\alpha}{1-\gamma}} = (\frac{1-\gamma}{\alpha})(1-e^{-\rho\Delta})c_t^{\alpha} + e^{-(\rho+\varphi)\Delta}\frac{1-\gamma}{\alpha}\left(\rho^{\frac{1-\gamma}{\alpha}}(1-\gamma)\mathbb{E}_t[\mathcal{U}_{t+\Delta}]\right)^{\frac{\alpha}{1-\gamma}}, \tag{D.1}$$

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 et al. (2026). In that case, the recursion simplifies to $\mathcal{U}_t = \mathbb{E}_t[\mathcal{U}_{t+\Delta}]$, so if we iterate backwards from the time of birth t to zero, we obtain that $\mathcal{U}_0 = \mathbb{E}_0[\mathcal{U}_t]$ so taking the expectation is justified even when preferences are recursive given our specification of the aggregator and the certainty equivalence operator.

D.2 Computing Utility Before Birth

Consider an unborn agent, that knows y_0, x_0, her preference type, and that she will be born at time t. Her expected (at time 0) lifetime utility at birth is:

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

The key object is something like $\mathbb{E}\left[(y_t\tilde{q}(x_t))^{1-\gamma_i}(V^i(x_t))^{\frac{1-\gamma_i}{1-\psi_i}}|x_0, y_0\right]$. We show below how to do so using Feyman-Kac. First, define the function $U^i(y_t, x_t, \tau) = \mathbb{E}\left[\frac{(y_{t+\tau}\tilde{q}(x_{t+\tau}))^{1-\gamma^i}}{1-\gamma_i}(V^i(x_{t+\tau}))^{\frac{1-\gamma^i}{1-\psi^i}}|x_t, y_t\right]$ Note that this function can be written as:

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

where $u^i(x_t, \tau) = \mathbb{E}\left[(\frac{y_{t+\tau}}{y_t}\tilde{q}(x_{t+\tau}))^{1-\gamma^i}(V^i(x_{t+\tau}))^{\frac{1-\gamma^i}{1-\psi^i}}|x_t\right]$, 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. Suppress the i index to lighten the notation. 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:

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

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

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

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 appropiate equilibrium law of motions for the states. With this, we have

$$\begin{aligned}
U^{i,t}(x,y) &= (s^{i,t})^{1-\gamma_i}\mathbb{E}\left[\frac{(y_t\tilde{q}(x_t))^{1-\gamma_i}}{1-\gamma_i}(V^i(x_t))^{\frac{1-\gamma_i}{1-\psi_i}}|x_0,y_0\right] \\
&= (s^{i,t})^{1-\gamma_i}\frac{(y)^{1-\gamma_i}}{1-\gamma_i}u(x,t),
\end{aligned}$$

which is the expression in the main text.

D.3 Valuing Aggregate Output Claims

Pricing Aggregate Output Claims Computationally, it is easier to price $yq^{a,i}(x,\tau) = \mathbb{E}[\int_0^\tau \frac{M_t}{M_0}\frac{y_t}{y_0}|x_0 = x, y_0 = y]$, i.e the value of aggregate output for the next τ years. Using the typical pricing formula, $yq(x,\tau)$ satisfies:

$$0 = y_t dt + \frac{E_t[d(M_t^i y_t q_t^{a,i})]}{M_t^i}) \tag{D.2}$$

Using Ito's Lemma,

$$\frac{d(M_t^i y_t q_t^{a,i})}{M_t^i y_t q_t^{a,i}} = \frac{dM_t^i}{M_t^i} + \frac{dy_t}{y_t} + \frac{dq_t^{a,i}}{q_t^{a,i}} + \frac{dM_t^i}{M_t^i}\frac{dy_t}{y_t} + \frac{dM_t^i}{M_t^i}\frac{dq_t^{a,i}}{q_t^{a,i}} + \frac{dy_t}{y_t}\frac{dq_t^{a,i}}{q_t^{a,i}}$$

but using Ito to compute $dq_t^{a,i}$, we have:

$$\begin{aligned}
\mathbb{E}_t\left[\frac{d(M_t^i y_t q_t^{a,i})}{M_t^i y_t q_t^{a,i}}\right] = \Bigg(&-\frac{q_\tau^{a,i}}{q^{a,i}} - (r^i+\varphi) + \mu + \frac{(\mu_x)^T\nabla q^{a,i}}{q^{a,i}} + \frac{1}{2}\frac{Tr(\Sigma_x\Sigma_x^T\nabla_x^2 q^{a,i})}{q^{a,i}} \\
&-([\sigma,0]\times\kappa^i) - \frac{(\kappa^i)^T\nabla q^{a,i}}{q^{a,i}} + [\sigma,0]\times\frac{\nabla q^{a,i}}{q^{a,i}}\Bigg)dt
\end{aligned}$$

replacing in D.2 we obtain:

$$q_\tau^{a,i} = 1 + q^{a,i}(-(r^i+\varphi) - [\sigma,0]\times\kappa^i + \mu) + ((\mu_x)^T - (\kappa^i)^T + [\sigma,0])\nabla q^{a,i} + \frac{1}{2}Tr(\Sigma_x\Sigma_x^T\nabla_x^2 q^{a,i}) \tag{D.3}$$

and we can solve this PDE backwards from $q^{a,i}(x,0) = 0$ for all x. Numerically, we iterate for a long number of periods until $q^{a,i}(x,\tau)$ is close to $q^{a,i}(x,\tau+\Delta)$.

Converting to marginal utility at time 0 A similar trick as above can be applied to derive the PDE for $\ell(x,\tau)$. For notational simplicity, define $k^i(x) = (\tilde{q}(x))^{-\gamma_i}(V(x))^{\frac{1-\gamma_i}{1-\psi_i}} q^{a,i}(x)$. Let $\ell^i(x,\tau) = \mathbb{E}\left[\left(\frac{y_{t+\tau}}{y_t}\right)^{1-\gamma_i}(k^i(x_{t+\tau}))|x_t = x\right]$. Letting $L^i(x_t, y_t, \tau) = y_t^{1-\gamma_i}\ell^i(x_t,\tau)$ and suppressing the i index to lighten the notation we have

$$dL = d(y_t^{1-\gamma_i})\ell + d\ell y_t^{1-\gamma_i} + d(y_t^{1-\gamma_i})d\ell$$

but $d(y_t^{1-\gamma_i}) = (1-\gamma_i)y_t^{1-\gamma_i}((\mu - \gamma_i\frac{\sigma^2}{2})dt + \Sigma_y^T dW$ as computed above, and $d\ell = (-\ell_\tau + (\nabla_x\ell)^T\mu_x + \frac{1}{2}Tr(\Sigma_x\Sigma_x^T\nabla_x^2\ell))dt + (\nabla_x\ell)^T\Sigma_x dW$ from Ito, thus:

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

For the same reason as before, $\mathbb{E}_t[dL] = 0$, so we obtain the PDE:

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

which we can iterate backwards starting from $\ell(x,0) = k^i(x)$. Note that the form of the PDE is the same as before, only the initial condition changes. Once we have ℓ^i, we can compute:

$$\begin{aligned} \lambda^{i,t}(x,y) &= (s^{i,t})^{-\gamma_i}L^i(x_0, y_0, t) \\ &= (s^{i,t})^{-\gamma_i}y^{1-\gamma_i}\ell^i(x,t). \end{aligned}$$