The Welfare Cost of Inflation Risk*

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

September 2026

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. The same ingredients generate a positive inflation risk premium even when inflation and aggregate consumption are uncorrelated, consistent with recent U.S. data. We also show that the welfare cost is governed by two key moments: the size of nominal positions and the inflation risk premium. We quantify this channel in a calibrated heterogeneous-agent economy matching key macro and asset-pricing moments, and find that the welfare costs of inflation risk are between 0.18% and 0.3% of perpetual aggregate consumption, comparable in magnitude to the costs of trend inflation due to transaction frictions or price dispersion.

Keywords: unexpected inflation, inflation risk premium, incomplete markets, heterogeneous agents, welfare cost of inflation.
JEL codes: E31, G12, D52

*Caravello: tomasec@mit.edu. de la Barrera: mdelabarrera@iese.edu. We are especially grateful to Eduardo Dávila and Stavros Panageas for extensive feedback at various stages of this project. We also thank Ricardo J. Caballero, Juanpa Nicolini, Andrés Schneider, Luis M. Viceira, Christian K. Wolf, and participants at the BSE Summer Forum, SED, and EEA-ESEM Congress for useful comments and suggestions. De la Barrera gratefully acknowledges the support of the Ministry of Science, Innovation and Universities, Grant JDC2024-053113-I funded by MICIU/AEI/10.13039/501100011033 and ESF+. All remaining errors are our own.

1 Introduction

The recent resurgence of inflation has put inflation and its associated welfare costs at the center of the macroeconomic policy debate. 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 exposes both parties to risk, which reduces financial trade and welfare. We quantify the welfare cost of inflation through this mechanism, and find it to be of the same order of magnitude as 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 a maturity of 5 years. 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 simple computation raises several conceptual issues: it is unclear why this should be related to welfare; nominal positions in a closed economy 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 clarify the role of 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 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 as well as preference and technology parameters. We derive three key results in this setup.

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 generates positive inflation risk premium even when inflation and aggregate consumption are uncorrelated at all leads and lags, which is hard to generate in representative agent 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. Following Dávila and Schaab (2025), we compute the efficiency gains from a marginal decrease in inflation risk. It turns out that these gains only depend on the direct channel: general equilibrium effects cancel out in the aggregate. This is important because while the strength of this channel depends on the specific assumptions about preferences

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

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 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-the-envelope formula as a lower bound to the true welfare cost, since it ignores the borrowing side of the market. Since these conclusions are obtained in a stylized two-period setting, we build a quantitative model that retains the main assumptions at its core and is also able to match several key moments of the data.

Quantitative Model. Quantifying the welfare cost of inflation risk requires a framework that simultaneously incorporates heterogeneity, aggregate uncertainty, and incomplete markets, a combination that is computationally challenging in standard heterogeneous-agent models. 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. There is an underlying overlapping generations 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. 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.[3] Since traded nominal bonds have positive duration, inflation shocks generate capital gains or losses, and therefore expose its holders to risk. The assumed duration directly indexes the severity of the market incompleteness due to the lack of real bonds, 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 that of outstanding debt in the market. We use the remaining preference and technology parameters to jointly match: the processes for aggregate consumption growth and inflation; and standard asset pricing moments, such as short term rates and the equity premium. Importantly, the calibrated model generates values for the inflation risk premium and aggregate nominal positions that are broadly consistent with the data, despite neither moment being directly targeted in the calibration.

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

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 Lucas (1987) and compute each agent's willingness to pay to eliminate inflation risk, in terms of a proportional consumption reduction. This yields a distribution 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.

These compensating variations cannot be aggregated into a single aggregate measure of efficiency that can be compared 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. By expressing each individual's welfare change in common units, we can aggregate the efficiency component across agents and cohorts. Given that we have an infinite number of cohorts, summing compensations for all agents will generally diverge, so we focus on the welfare gains accruing from consumption changes for the first 200 years, although results are similar for other thresholds. Under our baseline calibration, we find the gains to be around 0.2% of perpetual aggregate consumption, which is economically meaningful, and close to the back of the envelope estimate.

We also analyze the sensitivity of this number to several calibration choices, such as changing the underlying measure of inflation, the duration of debt, or the preferences. This yields estimates between 0.18% and 0.3% of perpetual aggregate consumption, which is of the same order of magnitude as the baseline figure.

We compare these numbers to estimates of costs of *trend* inflation from two standard sources: transaction and price-setting frictions, which we review in Section 5.4. These are the two main costs of inflation in standard macroeconomic models, with a vast literature that estimates their magnitudes and determinants. The main takeaway is that, while the exact magnitude of these costs is uncertain, our estimate of the cost of inflation risk through our channel is of the same order of magnitude as 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 transaction, price-setting, and other distortions; for an early taxonomy of these channels, see 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 showing how nominal exposures and inflation risk premia discipline the magnitude of this welfare cost, and then providing 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).[4] 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 on a mechanism that can generate a 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

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

about the potential welfare losses stemming from inflation volatility. This channel is absent in a representative-agent setting without heterogeneous nominal exposures. An important related paper is Gomez Cram et al. (2025), which also emphasizes the distributional consequences of inflation risk in asset markets. Our present paper differs from theirs in that we focus only on private sector lending. We abstract from the government and from the sources of inflation risk. In contrast, Gomez Cram et al. (2025) focus explicitly on inflation risk coming through fiscal channels, and the redistribution between bondholders and taxpayers.

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 (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 then uses the model to quantify the welfare gains of eliminating inflation risk. Section 6 offers some final remarks.

2 Motivation: The Cost of Hedging Nominal Positions

To gauge the order of magnitude of this cost, we ask: how costly would it be to hedge all nominal positions in the economy against inflation risk? If inflation risk reduces welfare, the agents who bear it will be willing to pay to eliminate it, and this willingness to pay will be reflected in market prices. 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. For each nominal bond that is hedged, we give up the expected real return $i - E[\pi]$ (the nominal rate minus expected inflation) and receive instead the guaranteed real return r. The expected cost for each bond is thus $i - E[\pi] - r =$ Inflation Risk Premium. Multiplying this quantity by the stock of nominal positions yields the following estimate of the aggregate cost of inflation risk:

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

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. To obtain an order-of-magnitude estimate, we treat all nominal positions as a single aggregate with a common 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 trillion dollars for federal and state governments combined. With nominal GDP equal to 31.4 trillion dollars at the end of 2025, private-sector debt (households and businesses combined) amounts to roughly 139% of GDP, government debt to 121% of GDP, and total debt to approximately 260% of GDP.

Translating these figures into a measure of inflation exposure requires taking a stance on two issues: i) the average maturity of this debt, and ii) the share of it that is effectively indexed and therefore insulated from inflation 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 set the maturity to 5 years. Measuring the degree of indexation is likewise challenging. However, virtually all private debt is issued in nominal terms, while only a small fraction of government debt is indexed. To remain conservative, we set nominal positions equal to 130% of GDP, roughly one-half of total outstanding debt. 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 considerably more difficult. 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) use 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 on 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

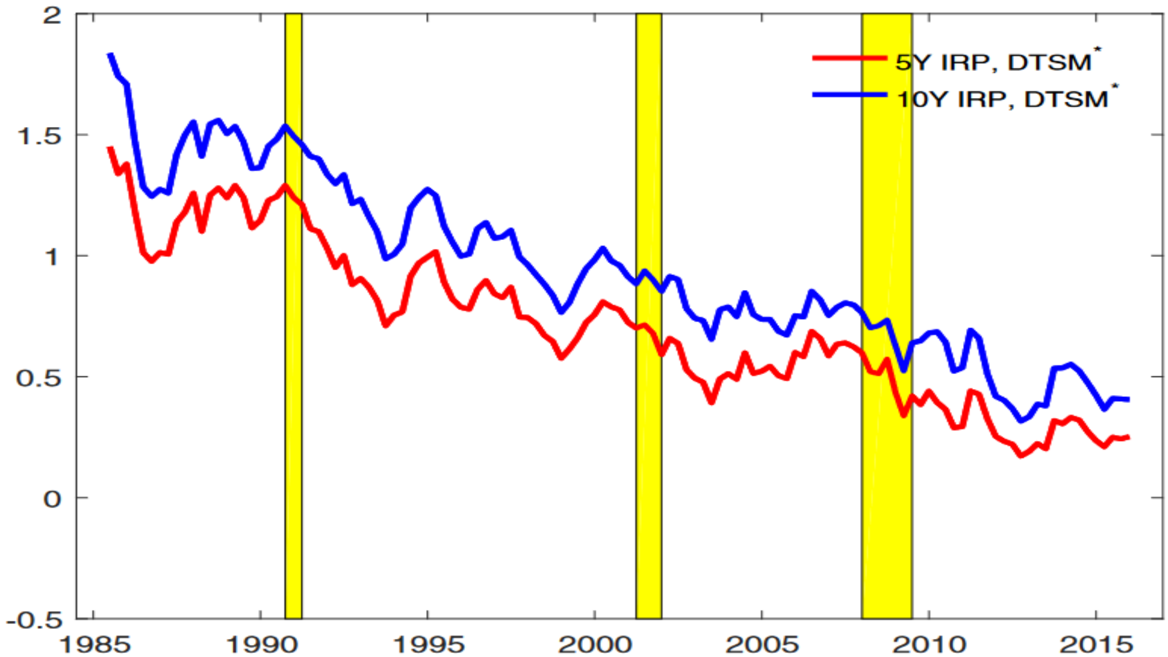

Figure 2.1: Inflation risk premium for 5 and 10 year bonds derived from the model estimated by Ajello et al. (2019).

model is estimated in the 1985Q1-2015Q4 sample and thus omits most of the large inflation episodes in the 70s and 80s. Although the inflation risk premium declines over the sample, it remains 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 cost of 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 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.[5] Our simple calculation indicates that, in principle,

[5] We review the literature on the quantification of typical welfare costs of inflation in Section 5.4.

the cost due to inflation risk can be of the same order of magnitude. However, reaching that conclusion at this point is clearly premature, since it is unclear whether (2.1) is a valid measure of the *welfare* cost of inflation risk in general equilibrium, or merely indicative of its order of magnitude.

Beyond improving this illustrative calculation, there are several conceptual issues with taking the above computation at face value. First, for each lender there is a borrower. One could equally measure the aggregate position from the borrowers' side and obtain the opposite sign, namely that inflation risk *increases* welfare. Second, it is unclear what counterfactual experiment the above exercise corresponds to. Ideally, one would like to compare 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 generally 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 addresses these conceptual issues 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 efficiency gain from reducing inflation risk is disciplined by nominal positions and the inflation risk premium.

3.1 Economic Environment

There are two time periods, $t = 0, 1$. Agents trade 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, and $\sigma_\pi, \sigma_y > 0$. 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 for 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 shares of the tree, ω_i. There are two agents $i = A, B$ with heterogeneous preferences and endowments. The optimization problem of agent i is:

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

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:

$$\mathbb{E}[u_i'(c_{i1})(\Pi^{-1} - \sigma_\pi \varepsilon_\pi)] = \lambda_i q_b^{\$} \tag{3.2}$$

$$\mathbb{E}[u_i'(c_{i1})(\bar{y} + \sigma_y \varepsilon_y)] = \lambda_i q_s \tag{3.3}$$

$$u'(c_{0i}) = \lambda_i \tag{3.4}$$

$$\mathbb{E}[u_i'(c_{i1})] \leq \lambda_i q_b \quad (\text{with } = \text{ if } b_i > 0). \tag{3.5}$$

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:

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

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

No Inflation Indexation. As we show below, a key assumption is the impossibility of borrowing in inflation indexed terms. This assumption is motivated by three considerations: 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.[6] The lack of indexation is often considered a puzzle (Fischer, 1975).[7] 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 generates empirically plausible inflation risk premia and nominal positions despite not targeting either moment in the calibration.

Third, we see the no indexation assumption as putting the proposed mechanism on an 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. Accordingly, inflation shocks can be interpreted as shocks to monetary policy. This interpretation highlights 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.[8]

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

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

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

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 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 establishes that both heterogeneity and incomplete indexation are necessary for inflation risk to affect welfare: if either is removed, then the equilibrium allocations do not depend on σ_π, and hence the welfare cost of inflation risk is zero.

Proposition 1. *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 way in which ε_π can affect allocations is if $|b_i^\$| > 0$ in equilibrium, so we need to show that, if either assumption is removed, 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:

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

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 is one agent whose consumption comoves positively with inflation, and the other agent has negative comovement. By the monotonicity of 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, i.e., there is some correlation between inflation and output, but inflation has some additional risk. In that case, by mixing trees and bonds, one can construct a portfolio that loads only on ε_z. Using this asset, one goes back to the same structure as before. 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, perhaps because these observed nominal exposures are hedged by nominal exposure in untraded assets, such as labor income. For that reason, we derive below a key distinctive implication of our setup: we can obtain a positive inflation risk premium even when aggregate output and inflation are uncorrelated.

3.3 Positive Inflation Risk Premium with no Output-Inflation Correlation

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

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

so the inflation risk premium crucially depends on $Cov\left(u_i'(c_{i1}), R^\$ - R\right)$. For the 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 it optimal to switch all her savings to real bonds: she gets a higher expected return and lower risk. In order to see that such a 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 decreasing marginal utility and is risk averse, this agent prefers to lend in real, and 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 $\mathbb{E}[R^\$ - R] > 0$. $\square$

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 the inflation risk premium is positive. We note that the argument in the proof is more general than the current setup. The intuition of the result is simple:

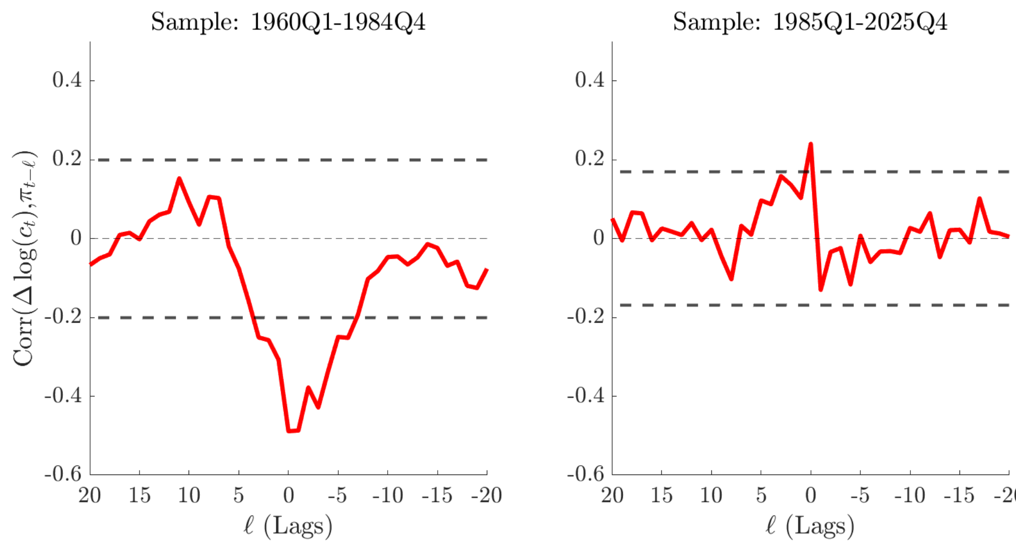

Figure 3.1: Cross-correlogram between real consumption growth and inflation for two different subsamples. Dashed lines are 95% confidence intervals.

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. For 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 hold in several extensions of the model, such as adding idiosyncratic income shocks, recursive preferences, allowing inflation to be correlated with output (in which case the result corresponds to the premium of an output-hedged position).

Empirical Evidence. While the assumption of zero correlation may sound 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 1985Q1-2025Q4 (right).[9] 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. This is consistent with the view that, in recent times, demand shocks account for a substantial share of aggregate consumption fluctuations, while their effect on inflation is limited (Angeletos et al., 2020).

[9]The data is obtained from FRED. Real consumption is obtained by dividing Personal Consumption Expenditures (PCE) by its corresponding price index (PCEPI). Inflation is measured using the Consumer Price Index (CPIAUCSL). Results are robust to using core inflation, and to excluding the Covid-19 period, see Appendix A.

Under the recent correlation structure, traditional consumption-based representative agent models would tend to generate 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 has been documented in studies using post-1980s data (Fleckenstein et al., 2017; Ajello et al., 2019). In short, we see these empirical patterns as evidence in favor of the two key assumptions on which our analysis hinges.

3.4 Inflation Risk Premium, Nominal Positions, and Welfare

Having explained what the key ingredients are, 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 consumption 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 on 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 order to get the welfare consequences of this in terms of $t = 0$ goods for agent i, we use the basic formula in Dávila and Schaab (2025) to obtain:

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

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

Since the changes in welfare are measured in common units, we can sum them to obtain a measure of the change in welfare due to efficiency following Dávila and Schaab

(2025). Letting Ξ_E be this change in aggregate welfare due to efficiency, we have:

$$\begin{aligned}
\Xi_E &= \sum_{i=A,B} \frac{dc_{i0}}{d\sigma_\pi} + \frac{\mathbb{E}[u_i'(c_{1i}) \times \left(-\varepsilon_\pi b_i^\$ + \frac{da_i}{d\sigma_\pi}(\bar{y} + \sigma_y \varepsilon_y) + \frac{db_i^\$}{d\sigma_\pi}(\Pi^{-1} - \sigma_\pi \varepsilon_\pi)\right)]}{u'(c_{i0})} \\
&= \sum_{i=A,B} \left(b_i^\$ \frac{\mathbb{E}[u_i'(c_{i1})(-\varepsilon_\pi)]}{u'(c_{i0})} \right) + \underbrace{\sum_{i=A,B} \left(\frac{dc_{i0}}{d\sigma_\pi} + \frac{da_i}{d\sigma_\pi} q_s + \frac{db_i^\$}{d\sigma_\pi} q_b^\$ \right)}_{=0} \\
&= \sum_{i=A,B} \left(b_i^\$ \frac{\mathbb{E}[u_i'(c_{i1})(-\varepsilon_\pi)]}{u'(c_{i0})} \right) \qquad (3.9)
\end{aligned}$$

where in the second line we used that $\frac{da_i}{d\sigma_\pi}$ and $\frac{db_i^\$}{d\sigma_\pi}$ do not depend on the realized states, and from the FOCs (3.2)-(3.4) for both agents we can replace the ratio of marginal utilities by prices. Then, due to market clearing, we have that $\sum_{i=A,B} \frac{dc_{i0}}{d\sigma_\pi} = \sum_{i=A,B} \frac{da_i}{d\sigma_\pi} = \sum_{i=A,B} \frac{db_i^\$}{d\sigma_\pi} = 0$, since neither aggregate output at time 0 nor the aggregate quantity of assets changes with the perturbation.

Equation (3.9) shows that the aggregate welfare effects depend *only* on the direct effect: indirect effects via re-optimization and general equilibrium forces cancel out in the aggregate, because agents already have equal marginal valuations between assets and time 0 good at the initial equilibrium. This is important because the strength of the direct channel only depends on the size of nominal positions, *regardless of why agents trade.*[10] Once $|b_i^\$| > 0$, reducing inflation risk directly reduces consumption risk, regardless of 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. Our calculation shows that, in this simple two period model, this only affects the distribution of welfare gains across agents, but not the total efficiency gains measured by Ξ_E.

In what follows, we show that a computation similar in spirit to Section 2 provides a lower bound for the welfare effects through this direct channel. Assume that this economy features non-zero nominal trade. This implies that there is a lender (borrower) with $b_i^\$ > 0$ ($b_i^\$ < 0$). The budget constraint implies 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). Thus, $\mathbb{E}[u_i'(c_{i1})(-\varepsilon_\pi)]b_i^\$/u'(c_{i0})$ is negative for both agents: a reduction in σ_π increases welfare for both agents through the direct channel. This may seem at odds with the standard intuition whereby inflation shocks

[10] In a sense, this is an ex-ante version of the exercises done in Doepke and Schneider (2006b), where nominal positions are taken as given.

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" yields:

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

This formula shows that, fixing the asset allocation and focusing on the unconstrained agent, we can evaluate the effects on 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, $b_i^\$$, appropriately discounted by time using 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 that $q_b \approx 1$ we can estimate the welfare gains (due to the direct channel) for the pricer of completely eliminating inflation risk as:

$$\underbrace{\text{Nominal Positions}}_{=|b_i^\$|}\times\underbrace{\text{Inflation Risk Premium}}_{=\sigma_\pi\kappa_\pi} \tag{3.10}$$

which is exactly the formula we used in Section 2. Note, however, that this only captures the gains for one of the agents. Using (3.9) and the fact that $b_i^\$$ and $\mathbb{E}[u_i'(c_{i1})(\varepsilon_\pi)]$ always have the same sign, we can bound the efficiency gains as:

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

Thus, through the lens of this model, 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.

3.5 Takeaways

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. We argue that this is a reasonable characterization of recent US data. Finally, we showed that the welfare cost of inflation risk can be approximated by the product of the inflation risk premium and the absolute size of nominal positions, and that this approximation constitutes a lower bound on the true efficiency gains from eliminating inflation risk. This justifies the back of the envelope calculation in Section 2 using a simple two period model. However, a key shortcoming is that this model is highly stylized: once we move to an infinite period, dynamic model, the indirect effects will not necessarily cancel out anymore. Thus, in order to quantify the welfare costs of inflation risk in a richer setting, 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 in 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 simply comparing the baseline economy to one where 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, who immediately consume it.[11] 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 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 infeasible, so we use tree depreciation as a tractable way of capturing the same force. See Panageas (2020) for more details.

[11] 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.

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:

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

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 obtains 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 bonds. 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 (Hatchondo and Martinez, 2009; Lorenzoni and Werning, 2019). 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 $\frac{P_t}{P_s} e^{-\delta_b(t-s)}$ 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 their 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 due to inflation risk in this economy is directly tied to the assumed duration of 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 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 we refer to this as trees, and treat it as a single asset. Thus, owning a tree yields $\int_{s=-\infty}^{t} y_{t,s} ds = \omega y_t$ 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 it depreciates at rate δ_q.

The expected return of a tree is $\mu_{q,t}$, and $\sigma_{q,t}^T = [\sigma_{q,t}^y, \sigma_{q,t}^\pi]$ is how much returns are affected by shocks to output and inflation. For notational simplicity, we write returns in matrix form

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

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 to 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.[12] 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 actuarially 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.

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

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

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

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 us 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 asset demands within each group. For example, consumption of type A agents 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 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:

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

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

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

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's 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 have 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:

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

Following Gârleanu and Panageas (2015), the corresponding SDF has drift $-r^i(x)$ and price of the risk $\kappa^i(x)$.[13] 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 nominal bonds because these are freely traded in equilibrium. For pricing all other assets, we use the SDF of the 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 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 $Q^\$(x,\tau)$

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

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. Denote 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 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 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 the specification of the output process in Gârleanu and Panageas (2015), with an annual growth rate of 2%, and $\sigma=0.04$ so that time-integrated data from our model can

	Description	Value
	Inflation Process	
λ_π	Persistence of inflation	1.347
$\bar{\pi}$	Inflation target	0.027
σ_π	Volatility	0.025
	Output Process	
μ	Output growth rate	0.02
σ	Volatility of output	0.04
ω	Dividend share	0.15

	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	Tree depreciation	0.03
δ_b	Bond duration	0.2

Table 1: Calibrated parameters.

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 left as an untargeted moment.

For inflation, we estimate an annual AR(1) process for CPI inflation in the 1985-2025 sample, and map the parameters to continuous time. Figure 4.2 shows the stationary distribution of inflation and compares 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, $\lambda_\pi = 1.347$, which means that the half life of the shock is around two 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.[14]

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

[14] In Section 5.3 we present robustness to targeting Core PCE inflation, and using a more persistent process with the same unconditional variance.

Moment	Baseline	$\sigma_\pi = 0$	Data
Real Rate	1.06%	1.10%	0.83%
Real Rate Volatility	0.52%	0.54%	1.33%
Equity Risk Premium	4.81%	4.86%	4.64%
Equity Volatility	22.89%	23.70%	16.05%
IRP 5-yr	0.11%	-	0.17%
IRP 10-yr	0.11%	-	0.45%
Credit / GDP	117.1%	129.2%	139.8%
Traded Wealth / GDP	452.5%	456.3%	448.0%

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.

As explained above, bond duration is the central parameter governing market incompleteness: longer duration increases agents' exposure to inflation risk and therefore amplifies the inflation risk premium generated by the model. With this in mind, long-term bonds are calibrated to have an average maturity of 5 years for the reasons discussed in Section 2.

4.5 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 5.

Standard Asset Pricing Moments. Table 2 shows several implied asset pricing moments on average and compares them to the data. Figure 4.1 reports those moments as a function of α (averaged across inflation values) and Figure 4.2 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

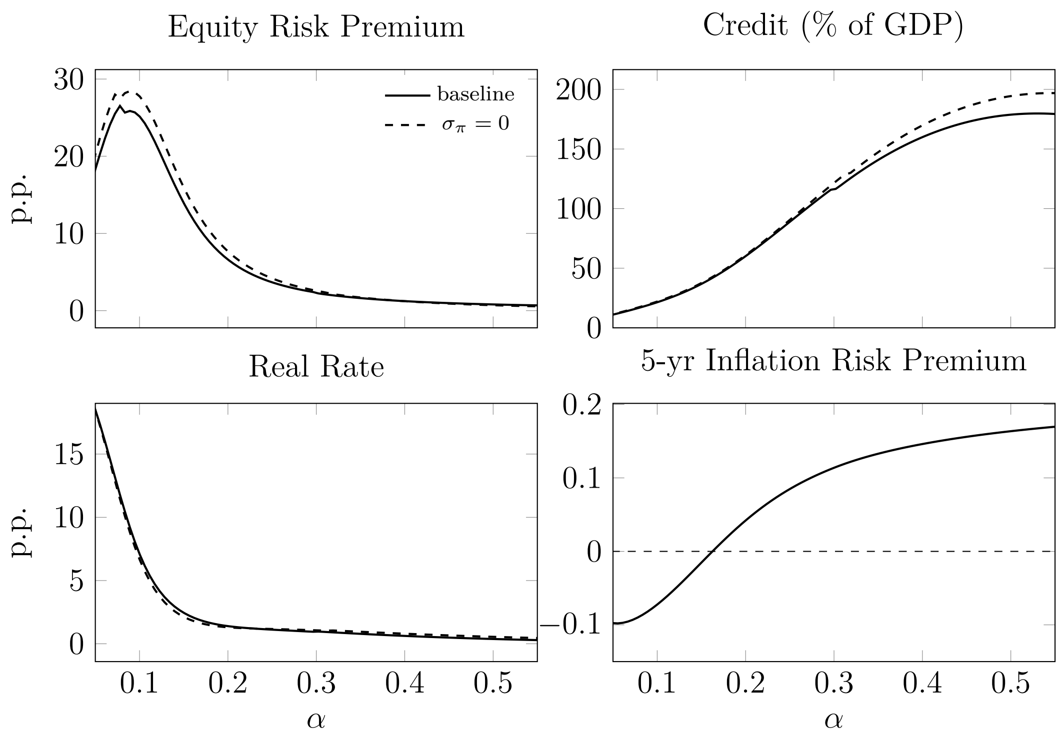

Figure 4.1: Selected asset-pricing and macroeconomic moments as a function of α. Solid line: baseline calibration. Dashed line: no inflation risk ($\sigma_\pi = 0$).

down). The particular combination of risk aversion and EIS chosen makes it so that changes in real rates and equity premium tend to reinforce each other, at least for a significant range of α. This generates significant volatility in price-dividend ratios, which in turn increases the equity premium.

Figure 4.1 shows the workings of the model in more detail. The equilibrium equity risk premium 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 forces. On the one hand, for low α, asset prices mostly reflect the preferences of the risk averse agent. On the other hand, individual risk-tolerant agents become highly levered for low α, which further amplifies stock market volatility at low α. This makes stocks 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 4.1 shows the size of nominal positions in this economy. We follow Schneider (2022) and denote this as “Credit". This is computed

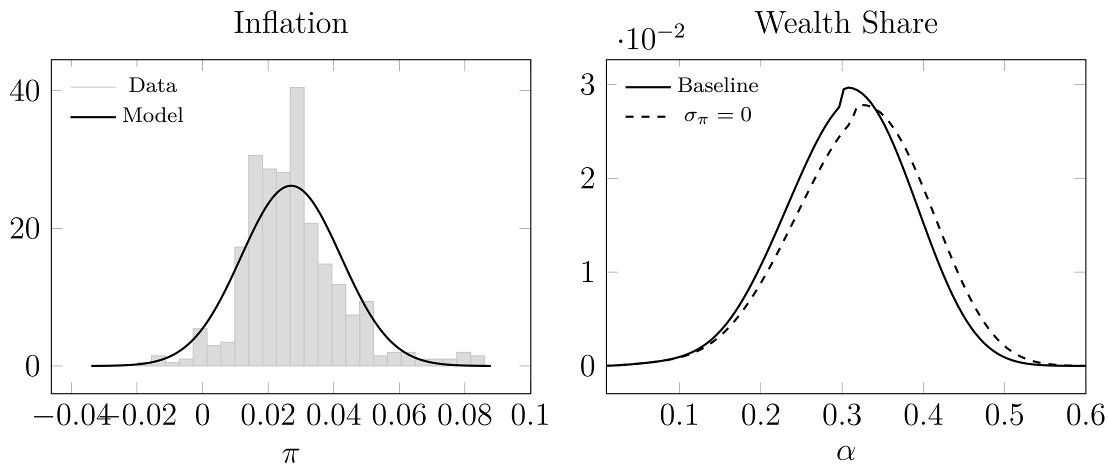

Figure 4.2: Stationary distributions. Left panel: inflation density in the data and in the estimated OU process. Right panel: stationary density of the wealth share α.

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 to 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 agent, 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 117% of GDP, somewhat below the average for this sample, around 140% of GDP. 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 4.1 shows the inflation risk premium on a 5-year zero coupon bond. Relative to Gârleanu and Panageas (2015) and Schneider (2022), the key novelty of our model is that it can generate a positive inflation risk premium out of uncorrelated inflation and output shocks, due to imperfect inflation indexation. The basic mechanism is simple: since agents have different preferences, 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 tends to be positive. We can also see 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, the 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 boost

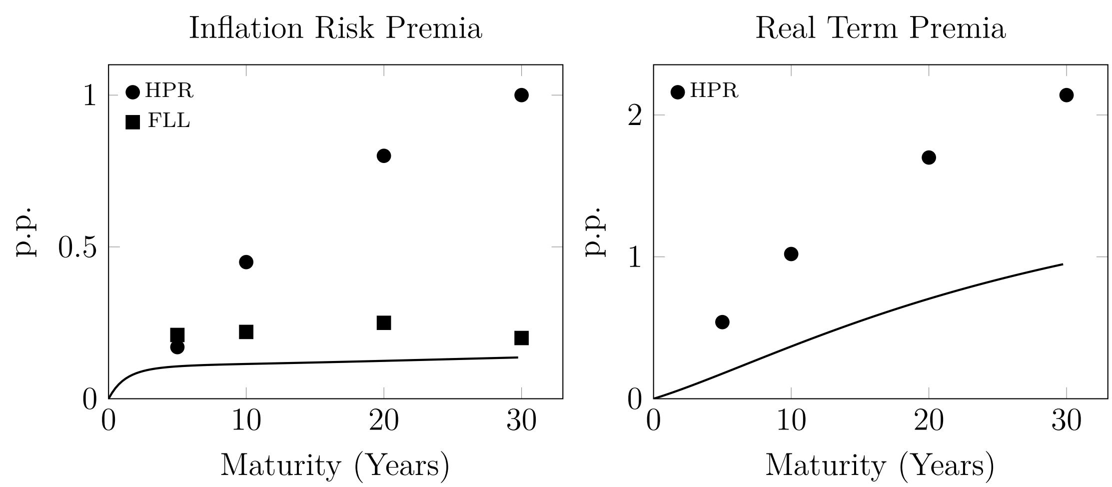

Figure 4.3: Term Structure of Inflation and Real Risk Premia. Circles: estimates from Haubrich et al. (2012). Squares: estimates from Fleckenstein et al. (2017).

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, which 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, although it is somewhat smaller than the one found in the empirical studies discussed in Section 2. This shows that our channel can explain an important part of the observed inflation risk premium at shorter maturities.

Term Structure of Risk Premia. Figure 4.3 depicts the term structure of inflation and real term premia, averaging across all states with the stationary distribution. The 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 predicts a 5-year inflation risk premium that is somewhat smaller than those estimates. For longer maturities, the model predicts a relatively flat term structure, in line with the estimates in Fleckenstein et al. (2017), but too low compared to the estimates in Haubrich et al. (2012). This is not surprising, given that our inflation process is not very persistent, and thus cannot capture other sources of long-term inflation risk, such as changes in the inflation target.

Regarding the real term premium, it 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 4.1. 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 the inflation and term premia, which are consistent with the data. While the model understates the level of the premia relative to some empirical estimates, it captures their key qualitative features. 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 4.1 and 4.2 show 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, especially for high values of α. The equity premium as a function of α shifts up. As we show in Figure 5.2 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. 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.

5 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 5.1, we follow 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 5.2, 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 5.4.

5.1 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 makes inflation deterministic ($\sigma_\pi = 0$), and leaves the rest of the 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 denoted 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 total number of agents φdt 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 $\mathcal{C}^{i,t}(x_0)$ defined as the solution to the following equation:

$$\mathbb{E}_0[\mathcal{U}_t^{i,t}|x_0, n_t^{i,t} = \bar{n}^{i,t}(x_t)] = \mathbb{E}_0[\mathcal{U}_t^{i,t,\sigma_\pi=0}|x_0, n_t^{i,t} = (1 - \frac{\mathcal{C}^{i,t}(x_0)}{100})\bar{n}^{i,t,\sigma_\pi=0}(x_t)], \tag{5.1}$$

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

[15]Note that this is using that the correct way of valuing the utility at time 0 of an agent that is going to be 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.

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 their 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:

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

Using homogeneity we can write:

$$\begin{aligned} U^{i,t}(x,y) &= \frac{(s^{i,t}y)^{1-\gamma_i}}{1-\gamma_i}u^i(x,\tau) \\ u^i(x_t,\tau) &= \mathbb{E}\left[(\frac{y_{t+\tau}}{y_t}\tilde{q}(x_{t+\tau}))^{1-\gamma^i}(V^i(x_{t+\tau}))^{\frac{1-\gamma^i}{1-\psi^i}}|x_t\right] \end{aligned} \tag{5.2}$$

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

$$\mathcal{C}^{i,t}(x) = \left[1 - \left(\frac{u^i(x,t)}{u^{i,\sigma_\pi=0}(x,t)}\right)^{1/(1-\gamma_i)}\right] \times 100. \tag{5.3}$$

Results. Figure 5.1 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 are willing to pay 3.7% (0.7%) of consumption in all states to eliminate inflation risk. 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. At 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 5.2 shows the optimal portfolios, consumption policies, as well as

[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.

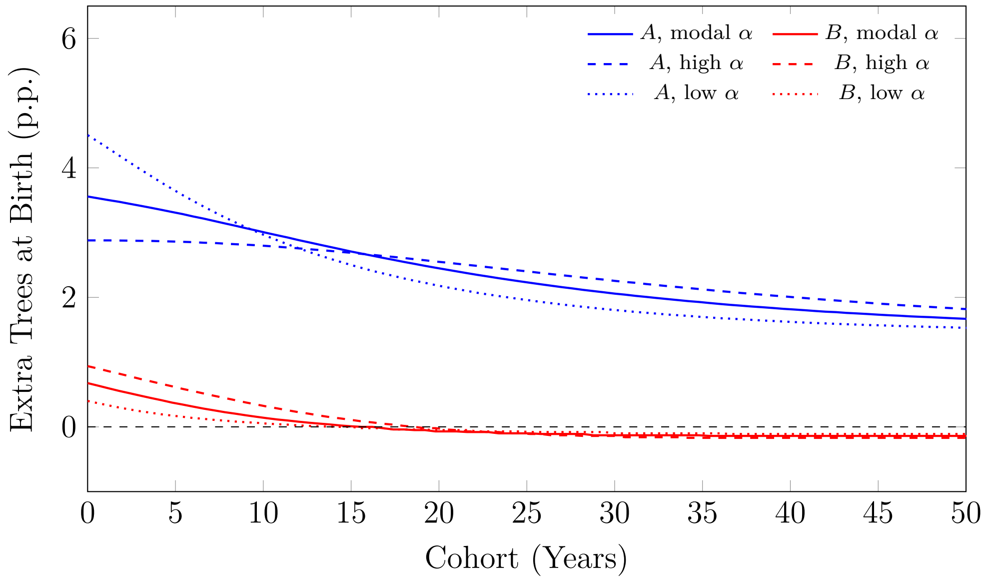

Figure 5.1: Willingness to pay to fully eliminate inflation risk, in units of proportional consumption reduction in all states, as a function of the cohort of birth.

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, it is also more volatile. Given that their elasticity of intertemporal substitution is high and their risk aversion is low, the force of higher drift dominates, and their welfare increases significantly upon elimination of inflation risk. This is shown in the last panel, which shows 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 hold a smaller share of their portfolio in trees in the 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

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

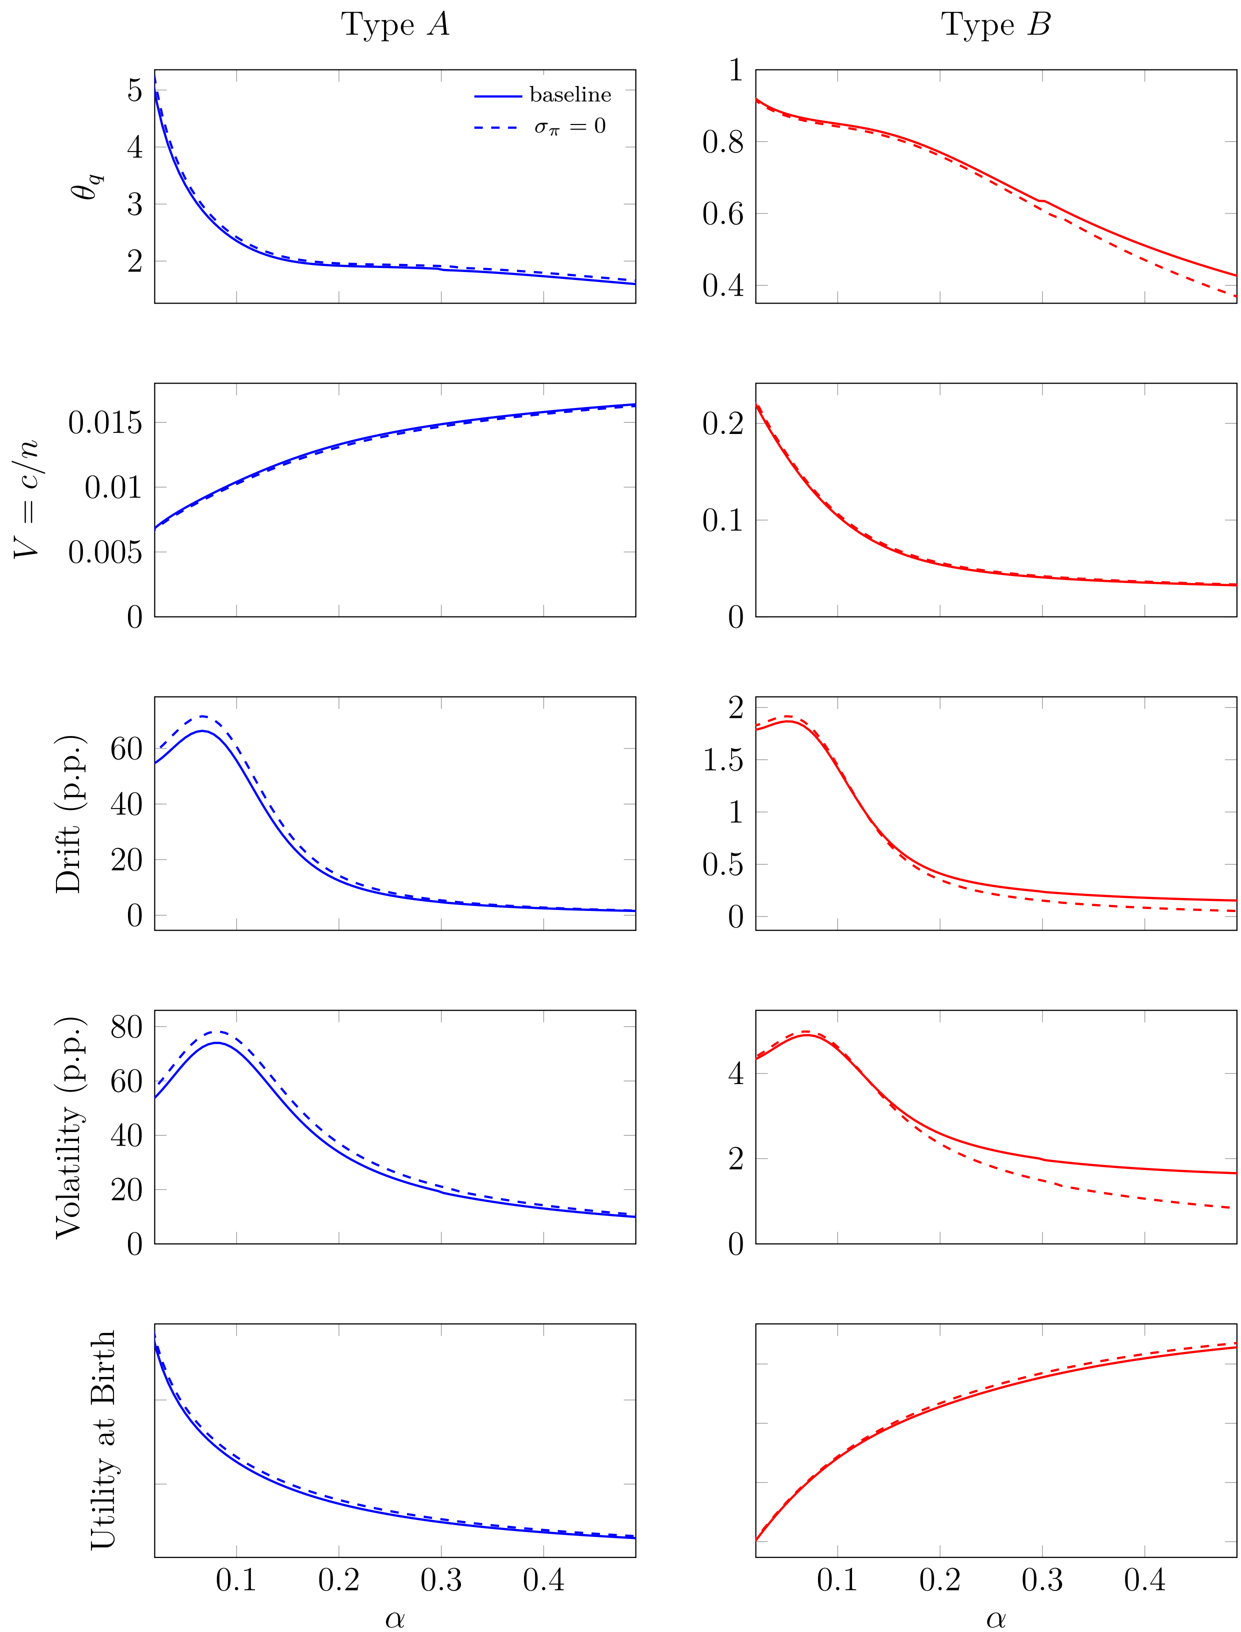

Figure 5.2: First Row: Portfolio shares for trees, θ_q^i. Second: consumption to wealth ratio, V^i. Third: Drift of the consumption process μ_c^i. Fourth: Instantaneous volatility of the consumption process, $\sqrt{(\sigma_c^i)^T(\sigma_c^i)}$. Last: Utility at birth $U^{i,0}(\alpha) = \frac{\tilde{q}(\alpha)^{1-\gamma_i}}{1-\gamma_i} V_i(\alpha)^{\frac{1-\gamma_i}{1-\psi_i}}$. 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.

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 5.1 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 get higher welfare gains. This also explains why the gains for type A agents are larger when the initial α 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 (5.2), 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 4.2. This effect is especially 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 chance of large market crashes due to increases in leverage, and for risk averse agents that will be born many years into the future, these states dominate in

[19]One way to see the strength of this effect to mechanically compute $\tilde{\mathcal{C}}^{i,t}(x) = 1 - \left(\mathbb{E}[u^i(x_t, 0)|x_0 = x]/\mathbb{E}[u^{i,\sigma_\pi=0}(x_t, 0)|x_0 = x]\right)^{1/(1-\gamma_i)}$ 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 (5.3) in that it ignores the $(\frac{y_{t+\tau}}{y_t})^{1-\gamma}$ term. In this case, both agents have positive long-term gains.

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.

The main quantitative result is that eliminating inflation risk generates economically large welfare gains for currently alive agents, with gains concentrated among risk-tolerant investors and earlier cohorts. 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.

5.2 Aggregate Welfare Gains

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 a utilitarian social welfare function (SWF) $\mathcal{W}$, of the form:

$$\mathcal{W}(x,y) = \lim_{\bar{T}\to+\infty} \int \int_0^{\bar{T}} w^{i,t} U^{i,t}(x,y) dt di \tag{5.4}$$

where the welfare weights $w^{i,t}$ are such that $\mathcal{W}(x,y) < +\infty$ as $\bar{T} \to \infty$.[20] 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:

$$\frac{d\mathcal{W}(x,y)}{d\sigma_\pi} = \Xi^E(x,y) + \Xi^{RD}(x,y), \quad \Xi^E(x,y) = \lim_{\bar{T}\to+\infty} \int \int_0^{\bar{T}} \frac{dU^{i,t}(x,y)/d\sigma_\pi}{\lambda^{i,t}(x,y)} dt di. \tag{5.5}$$

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:

$$\lambda^{i,t}(x,y) = \mathbb{E}\left[\int_0^\infty M^{i,t}_{t+h} y_{t+h} dh | x_0 = x, y_0 = y\right] \tag{5.6}$$

[20] We leave implicit in the notation that the initially alive cohort corresponds to a mass of agents.

where $M^{i,t}_{t+h} = \mathbf{1}_{h\geq 0}\exp(\int_0^h f_{\mathcal{U}}(c^{i,t}_{t+s}, \mathcal{U}^{i,t}_{t+s})ds) f_c(c^{i,t}_{t+h}, \mathcal{U}^{i,t}_{t+h})$ 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 straightforward: 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 use 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 if we considered a perturbation that increased welfare of all agents cohorts by 0.01, then $\Xi_E = \lim_{T\to+\infty} 0.01 \times \bar{T} = +\infty$. For this reason, we focus on a truncated version of Ξ_E, using the term structure decomposition introduced in Dávila and Schaab (2025).

In particular, note that, for any type-cohort (i,t), we can decompose the marginal change in welfare, $dU^{i,t}(x,y)/d\sigma_\pi$, as the sum of the changes in consumption at each date, as:

$$\frac{dU^{i,t}(x_0,y_0)/d\sigma_\pi}{\lambda^{i,t}(x_0,y_0)} = \int_0^\infty \chi^{i,t}(x_0,h)dh$$

$$\text{where } \chi^{i,t}(x_0,h) = \frac{\mathbb{E}_0\left[M^{i,t}(s^{t+h})\, dc^{i,t}_{t+h}(s^{t+h})/d\sigma_\pi\right]}{\lambda^{i,t}(x_0,y_0)}. \tag{5.7}$$

Here, $\chi^{i,t}(x_0,h)dh$ represents the marginal gain or loss in welfare for agents of type-cohort (i,t), due to changes in their consumption at time $t+h$, conditional on the state at zero being (x_0,y_0), in units of the welfare numeraire. We show in Appendix D that this term does not depend on the initial y_0 due to homotheticity. Integrating this for all dates, we get the total change in welfare for those agents. However, given those, we can also sum the welfare gains differently: we can aggregate them by date of accrual. In particular, fixing a time τ, we let $\mathcal{W}^i(x_0,\tau)$ denote the total welfare gains of all type i agents at date τ (regardless of when they were born), and $\mathcal{W}(x_0,\tau)$ to be the aggregate across both types, we can compute:

$$\mathcal{W}^i(x_0,\tau) = \chi^{i,0}(x_0,\tau) + \varphi \int_0^\tau \chi^{i,t}(x_0,\tau-t)\,dt, \tag{5.8}$$

$$\mathcal{W}(x_0,\tau) = \bar{\alpha}\mathcal{W}^A(x_0,\tau) + (1-\bar{\alpha})\mathcal{W}^B(x_0,\tau). \tag{5.9}$$

The second term is particularly important: it measures the welfare gains or losses coming from consumption at time τ, regardless of which agents are enjoying them. Given this

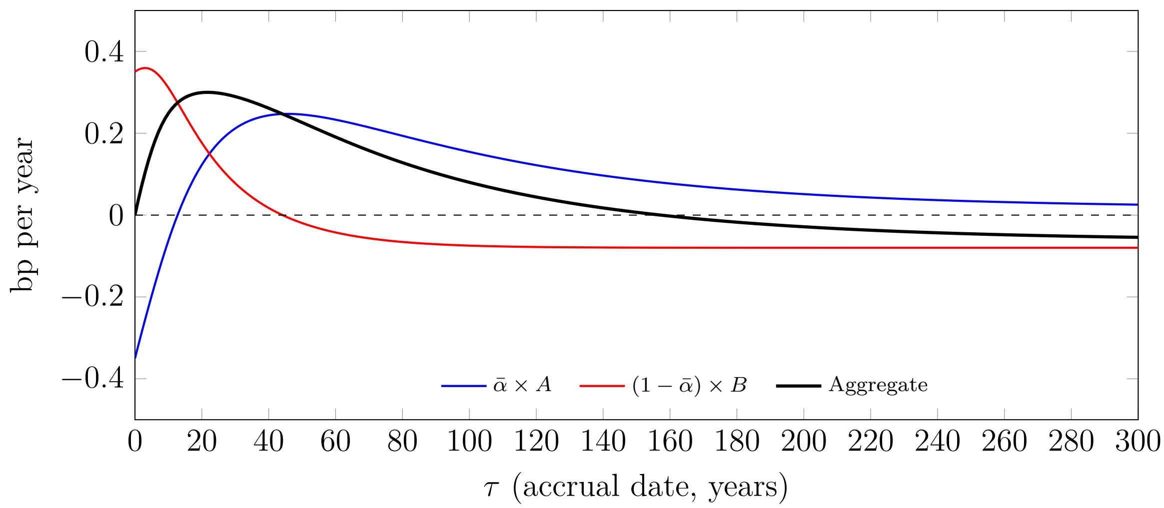

Figure 5.3: Welfare Term Structure. Black: term $-\sigma_\pi \mathcal{W}(x_0, \tau)$ from equation (5.9). Blue and Red: contributions of each type of agent to this term, $-\sigma_\pi \bar{\alpha} \mathcal{W}^A(x_0, \tau)$ and $-\sigma_\pi(1 - \bar{\alpha})\mathcal{W}^B(x_0, \tau)$, where $\mathcal{W}^i(x_0, \tau)$ is defined in (5.8).

definition, we can decompose the aggregate efficiency term by date of accrual. In particular, fixing some T and $\bar{T}$, we can decompose Ξ_E as:

$$\Xi_E = \int_0^T \mathcal{W}(x_0, \tau) d\tau + \int_T^\infty \mathcal{W}(x_0, \tau) d\tau.$$

We report $\mathcal{W}(x_0, \tau)$ as a function of τ for a fixed initial state, and also the cumulative gains up to different horizons. Appendix D.3 shows how to compute each of the terms from the recursive solution of the model. We report all results fixing x_0 at its modal values.

Results. Figure 5.3 reports $\mathcal{W}(x_0, \tau)$ as a function of τ, and it also shows the contribution of each type of agent to the total. At the beginning, the aggregate gain is very close to 0, with a positive contribution from type B agents and a negative one from type A. This is consistent with our previous discussion: when σ_π falls, type A agents start saving more, so their consumption drops; whereas type B agents consume more. As time goes by, the gains from type A turn positive and the gains of type B agents fall and eventually turn negative. As we discussed in Section 5.1, changes in time-0 welfare for type-B agents to be born far out in the future are mainly driven by minor changes in the probability of α reaching very low values, which is why it turns negative.

Table 3 presents the total efficiency gains accrued up to different horizons, as well as their sources across agents and cohorts. Based on this analysis, eliminating inflation risk (and focusing only on the gains accrued for the first 200 years) is approximately equivalent to a 0.20% perpetual increase in aggregate GDP, which is a non-trivial magnitude, especially

Horizon T (years)	50	100	150	200	250	300
Aggregate	0.125	0.198	0.217	0.211	0.192	0.167
A	1.966	6.030	8.368	9.703	10.537	11.123
B	0.078	0.048	0.008	-0.033	-0.073	-0.114
Initially alive	0.139	0.243	0.293	0.318	0.330	0.336
A	1.893	5.615	7.584	8.561	9.044	9.281
B	0.094	0.105	0.106	0.106	0.106	0.106
Unborn	-0.014	-0.045	-0.076	-0.107	-0.138	-0.169
A	0.074	0.415	0.784	1.142	1.493	1.841
B	-0.017	-0.057	-0.098	-0.139	-0.180	-0.220

Table 3: Cumulative efficiency gains accrued at different horizons, multiplied by the total reduction in 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).

compared to other typical costs of inflation. This means that the consumption effects (during the first 200 years) of eliminating inflation risk are approximately equivalent to an increase of 0.2% in aggregate output for all times and states.[21]

This aggregate number is also quite close to the simple approximation undertaken in Section 2. In terms of the sources across cohorts, the initially alive cohort is the main driver of gains. Cohorts born after $t = 0$ have a *negative* contribution, which is mainly driven by the welfare losses of type B agents discussed above, and to which we return below.

Another striking pattern is that eliminating inflation risk significantly benefits type-A agents. While these agents are only 2.5% of the population under the baseline calibration, compensating them for the first 200 years would require approximately $0.025 \times 9.7 \approx 0.24\%$ of perpetual aggregate consumption, which is more than the total compensation. The magnitude of the split between initially alive type A agents and later born ones also deserves 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 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

[21]Note that the magnitudes are different from those in Figure 5.1 because the units are different. In particular, these gains are measured in units of perpetual aggregate consumption, whereas the gains in Figure 5.1 are measured in terms of proportional changes in the consumption path of each individual agent, and thus cannot be aggregated. Second, recall that only a fraction $\omega = 0.15$ of aggregate output accrues to investors. This scales down the final figure, since we compute compensations in terms of *aggregate* output.

unborn cohorts, type A agents only represent 2.5% of the wealth, so compensating them is much cheaper.

It is also important to note that, as we consider a longer horizon, the cumulative efficiency gains start to decline. This is driven by the negative compensation of type B agents: as we discussed above, eliminating inflation risk makes low probability market crashes marginally more likely, but type-B agents put extreme weight on these events. Our results indicate that this force is so strong that it outweighs both the gains from lower inflation risk, and the gains from type A agents, and thus the efficiency gains stemming from consumption changes far out in the future turn negative.

This occurs because, even if there is no inflation, markets are incomplete: unborn agents cannot trade with each other until they are born. If they could, unborn type B agents would like to ex-ante insure from these low probability events. Given market incompleteness, the equilibrium is inefficient. We study this in more detail in Appendix E, where we show that introducing a borrowing tax generates efficiency gains. Intuitively, when type A agents lever up, this creates larger volatility in the price of trees, and thus higher risk from the point of view of type B unborn agents. This is not internalized by type A agents, and thus inefficiencies arise. We show there that a borrowing tax can help internalize this, reduce borrowing and increase welfare. This shows that our estimate of the welfare gain is reduced by an inefficiency present due to the OLG structure of the model.

Comparison with Back of the Envelope Calculation. It is interesting to note that the estimated figure, of around 0.2% of perpetual aggregate consumption, is close to, but somewhat lower than, the back of the envelope estimate in Section 2. This is because of two main reasons. First, our model does not perfectly fit credit and the inflation risk premium. In particular, our model generates an inflation risk premium that is about half of that used in the back of the envelope calculation. Nominal credit is also slightly smaller. If we applied the same back of the envelope logic for the moments in our model economy, we would have approximated that welfare gains are around $1.17 \times 0.11 \approx 0.13\%$ of GDP, which is smaller than what our measured gain is.

The discrepancy is due to a second reason: in the dynamic model, the general equilibrium terms will not cancel out in general. In this case, eliminating inflation risk has additional welfare effects, because it improves risk-sharing among living agents: now risk-averse agents are less exposed to aggregate output risk, which increases their welfare. However, as explained above, it also worsens ex-ante risk-sharing among unborn agents. For the horizon we consider, the overall effect of these general equilibrium margins is positive.

As explained in Section 3, the size of the gains that come from general equilibrium

	Baseline	Core PCE	$\delta_b = 0.1$	$\delta_b = 0.4$	$\sigma_\pi = 0.03$	$\lambda_\pi/2$	$\psi_B = 0.1$
Moments							
Real Rate	1.06%	1.06%	1.05%	1.06%	1.04%	1.04%	1.36%
Equity R.P.	4.81%	4.83%	4.80%	4.82%	4.78%	4.78%	3.45%
IRP 5-yr	0.11%	0.08%	0.11%	0.10%	0.14%	0.14%	0.11%
Credit / GDP	117.1%	119.1%	122.6%	116.1%	112.8%	111.8%	135.7%
Welfare gains							
Aggregate	0.21%	0.18%	0.24%	0.18%	0.30%	0.30%	0.25%
A	9.70%	7.90%	10.72%	8.18%	14.75%	13.23%	10.22%
B	-0.03%	-0.02%	-0.03%	-0.03%	-0.07%	-0.03%	-0.01%

Table 4: Selected moments and welfare gains of eliminating inflation risk under different alternative calibrations. The welfare numbers correspond to the gains accrued for the first $T = 200$ years. For the $\lambda_\pi/2$ case, we also change σ_π in order to match the same stationary variance.

forces is 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 (absent other distortions that interact with borrowing) that would lead to higher welfare above and beyond the direct effect of reducing risk on existing positions. However, as our results show, different externalities that are related to borrowing also interact with this channel. A deeper dive on the implications of alternative reasons for trade is left for future research.

5.3 Sensitivity Analysis

In this section, we show how the headline welfare number varies under alternative calibrations. The results are reported in Table 4, where the first column repeats the baseline estimates in order to facilitate the comparison. The second column shows what happens if we use an alternative calibration for the process of inflation, based on Core PCE instead of headline. The estimated parameters in this case are $\lambda_\pi = 0.39$, $\bar{\pi} = 0.022$ and $\sigma_\pi = 0.009$. This makes inflation less volatile but more persistent. In this case, welfare gains of eliminating inflation risk are relatively stable, at around 0.18% of perpetual aggregate consumption.

The next two columns show what happens when we change the parameter governing duration of traded nominal bonds, to $\delta_b = 0.1$ (which corresponds to 10 years) and $\delta_b = 0.4$ (2.5 years). As we explained, this governs the effective degree of market incompleteness. As expected, aggregate welfare gains get larger (smaller), but remain relatively close to the baseline. The fifth column shows what happens if we increase the volatility of inflation to $\sigma_\pi = 0.03$. In this case, welfare gains of reducing inflation volatility go up to 0.3%. The inflation risk premium increases significantly, and credit is reduced. The next column shows what happens if we make inflation twice as persistent, but keep the unconditional volatility unchanged. This meaningfully increases the inflation risk premium and the welfare gains from eliminating inflation risk. Finally, the last column shows what happens if we use a less extreme value for ψ_B, of 0.1. In that case, the equity risk premium is significantly reduced. The welfare gains from reducing inflation risk remain similar to the baseline.

5.4 Comparison with Standard Sources of Inflation Costs

We compare the estimated cost to the main sources of inflation costs in standard macroeconomic models: transaction and pricing frictions. While these costs are mainly estimated in representative agent models without aggregate uncertainty, these are usually reported in units of perpetual aggregate consumption, so they can be compared to the costs estimated above. 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 each of these costs of trend inflation is estimated to be of the same order of magnitude as our cost of inflation risk.

Transaction Frictions. The classical cost of trend inflation relates to its distortionary effects on 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

[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 distorts risk-sharing of aggregate shocks.

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 of the same order of magnitude as our baseline estimate 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 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 the welfare cost of 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, which is of a similar order of magnitude as our estimated cost of inflation risk.

6 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

[23]In our flexible price 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.

of the data: the size of nominal positions and the inflation risk premium. We quantify it in a general equilibrium model, and find it 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 the 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 its size and sources more valuable.

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 can quantitatively match key asset pricing moments, such as the one we presented here.

Finally, we modeled the absence of inflation indexation as an exogenous friction, and showed that its welfare consequences are non-trivial. A natural question is whether the magnitude of the welfare costs depends 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 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.

[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.

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.

Auclert, A. (2019): "Monetary Policy and the Redistribution Channel," *American Economic Review*, 109, 2333–2367.

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," *The American Economic Review*, 98, 438–443.

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.

Gomez Cram, R., H. Kung, H. N. Lustig, and D. Zeke (2025): "Fiscal Redistribution Risk in Treasury Markets," NBER Working Paper 33769, National Bureau of Economic Research.

Hatchondo, J. C. and L. Martinez (2009): "Long-Duration Bonds and Sovereign Defaults," *Journal of International Economics*, 79, 117–125.

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.

Lorenzoni, G. and I. Werning (2019): "Slow moving debt crises," *American Economic Review*, 109, 3229–3263.

Lucas, Robert E., J. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445.

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

——— (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.

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," NBER Working Paper 21448, National Bureau of Economic Research.

A Empirical Evidence

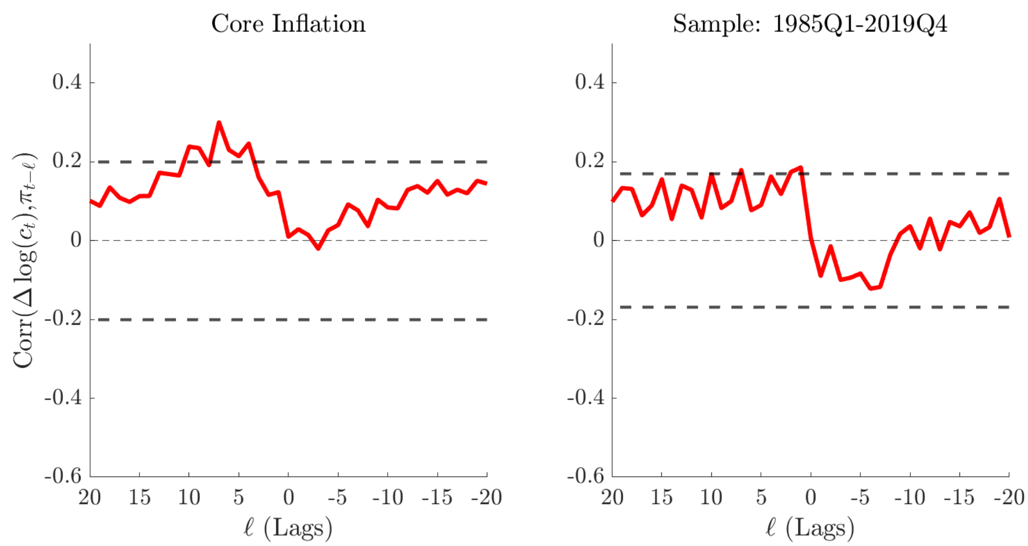

Figure A.1: Cross-correlogram between PCE real consumption growth and inflation. Left: use core CPI instead of headline CPI. Right: end the sample on 2019Q4. Dashed lines are 95% confidence intervals.

Figure A.1 shows the robustness of the results to alternative choices. The left panel shows what happens if we use 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 cut the sample to exclude the Covid period. The finding of a zero or weakly positive correlation is robust to both choices.

B Simple model

B.1 Cashless Limit

In order to explicitly add money, consider a government that issues nominal money and taxes the agents in lump sum by T_0, T_1. 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:

$$
\begin{aligned}
\max_{c_{0i},a_i,b_i,b_i^{\$},m_{0i},m_{1i}} \quad & u_i(c_{0i}) + a\log(m_{0i}) + \mathbb{E}[u_i(c_{1i}) + a\log(m_{1i})] \\
& \text{s.t. } c_{0i} + m_{0i} + q_s(a_i - \omega_i) + q_b b_i + q_b^{\$} b_i^{\$} = y_{0i} - T_0 \\
& \quad c_{1i} + m_{1i} = a_i(\bar{y} + \sigma_y \varepsilon_y) + (b_i^{\$} + m_{0i})(\Pi^{-1} - \sigma_\pi \varepsilon_\pi) + b_i - T_1 \\
& \quad b_i \geq 0.
\end{aligned}
$$

The FOCs now read:

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

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

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

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:

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

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 concave. 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:

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

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, i.e., $|b_i^{\$}| \neq 0$. 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. We again proceed 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 substituting 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 cannot 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 output. 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, we 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).

[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 will be clear below, 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 continuous time set-up, everything will be locally linear and thus the intuition derived under normality actually applies.

C Solving the model

C.1 The HJB of the investors

Given the wealth of an agent n 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

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

The first constraint captures the no-short-selling constraint on real, short-term bonds that we impose. A levered portfolio in stocks $\theta^i_q > 1$, where none of the borrowing is done in nominal terms $\theta^i_b = 0$ and instead is done in real terms, is not allowed.

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

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

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

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

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.

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

The first order conditions with respect to portfolio shares read

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

where ξ_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 excess 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 portfolio 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

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

Here, $\theta_t^A = [\theta_{q,t}^A,\ \theta_{b,t}^A]^T$ is the portfolio 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, valued at q_t. Thus, $q_t\frac{\delta_q}{\varphi}\varphi\bar{\alpha}$. Now, divide the expression by n_t^A and use that $n_t^A = \alpha_t^A n_t$, and total wealth is the price of capital q_t.

Next, the evolution of total wealth satisfies

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

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 existing trees is already taken into account inside the term μ_t.

Applying the Itô rule for divisions and with some rearrangement, 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 equilibrium object by applying Itô's lemma.

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

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

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

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

C.4 The Stochastic Discount Factor

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

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

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:

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

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 $\frac{dV}{V} = \mu_V(x)dt + \Sigma^T_V(x)dW$. 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 $(\theta^i(x))^T\mathbf{1} = 1$

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

the objects μ, Σ, θ^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 difference 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, $\alpha V^A + (1-\alpha)V^B = \omega/\tilde{q}$.

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 returns 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:

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

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

$$\tilde{U}^{i,t}(x_0, y_0) = \mathbb{E}_0 \left[\frac{\left(\frac{\delta_q}{\varphi} y_t \tilde{q}(x_t)\right)^{1-\gamma_i}}{1-\gamma_i} (V^i(x_t))^{\frac{1-\gamma_i}{1-\psi_i}} | x_0, y_0 \right]. \tag{D.2}$$

The agent receives δ_q/φ trees, and trees are valued at $y_t \tilde{q}(x_t)$. For initially alive agents, their utility is

$$\tilde{U}^{i,\text{initially alive}}(x_0, y_0) = \mathbb{E}_0 \left[\frac{\left(\frac{\alpha^i}{\bar{\alpha}^i} y_t \tilde{q}(x_t)\right)^{1-\gamma_i}}{1-\gamma_i} (V^i(x_t))^{\frac{1-\gamma_i}{1-\psi_i}} | x_0, y_0 \right]. \tag{D.3}$$

this is the same expression as (D.2) with different endowment of trees. Agents of type i, as a group, hold α^i of the trees and there are $\bar{\alpha}^i$ of such agents.

To simplify notation, we now define the utility of an unborn agent when $y_0 = 1$ and receives one tree.

$$U^{i,t}(x_0) = \mathbb{E}_0 \left[\frac{(y_t \tilde{q}(x_t))^{1-\gamma_i}}{1-\gamma_i} (V^i(x_t))^{\frac{1-\gamma_i}{1-\psi_i}} | x_0, y_0 = 1 \right]. \tag{D.4}$$

Next, define

$$U^i(y_t, x_t, t+\tau) = \mathbb{E}_t \left[\frac{(y_{t+\tau} \tilde{q}(x_{t+\tau}))^{1-\gamma^i}}{1-\gamma_i} (V^i(x_{t+\tau}))^{\frac{1-\gamma^i}{1-\psi^i}} | x_t, y_t \right] \tag{D.5}$$

this is the utility of an agent born at $t+\tau$ when the state at time t is (x_t, y_t). We can decompose it as

$$U^i(y_t, x_t, t+\tau) = \frac{y_t^{1-\gamma_i}}{1-\gamma_i} u^i(x_t, t+\tau) \tag{D.6}$$

Where $u^i(x_t, t+\tau) \equiv \mathbb{E}\left[(\frac{y_{t+\tau}}{y_t} \tilde{q}(x_{t+\tau}))^{1-\gamma^i} (V^i(x_{t+\tau}))^{\frac{1-\gamma^i}{1-\psi^i}} | x_t \right]$. We do so because $u^i(x_t, t+\tau)$ 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:

$$U(y_{t+dt}, x_{t+dt}, T) = U(y_t, x_t, T) + dU_t \tag{D.7}$$

and taking expectations conditional on the information at time t

$$\mathbb{E}_t[U(y_{t+dt}, x_{t+dt}, T)] = U(y_t, x_t, T) + \mathbb{E}_t[dU_t] \tag{D.8}$$

and by the law of iterated expectations,

$$\mathbb{E}_t[U(y_{t+dt}, x_{t+dt}, T)] = \mathbb{E}_t\mathbb{E}_{t+dt}[U(y_T, x_T, T)|y_{t+dt}, x_{t+dt}] = U(y_t, x_t, T) \tag{D.9}$$

which implies $\mathbb{E}_t\left[dU\right] = 0$. 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 initial conditions $u(x, 0) = (\tilde{q}(x))^{1-\gamma^i}(V^i(x))^{\frac{1-\gamma^i}{1-\psi^i}}$. The equation is defined analogously for the alternative economy with no inflation, we just need to replace the appropriate equilibrium laws of motion 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 The Term Structure of Welfare Gains

This subsection characterizes when the marginal welfare effects of a change in inflation volatility accrue. Throughout, we treat all equilibrium objects as functions of σ_π, and every derivative ∂_{σ_π} is evaluated at the baseline calibration. We use t for a cohort's birth date, h for its age, $\tau = t+h$ for the calendar date at which welfare accrues, and T for a terminal calendar date. Brownian motion is denoted by W, and all time-0 expectations are conditional on the initial state (x_0, y_0).

D.3.1 Definitions

Consider an agent of type i born at time t, with the convention that the initially alive agents are the cohort $t = 0$. Denote by $c^{i,t}_{t+h}(s^{t+h})$ her consumption at age $h \geq 0$, as a function of the history of shocks s^{t+h}, and denote by

$$M^{i,t}(s^{t+h}) = \exp\left(\int_0^h f_{\mathcal{U}}(c^{i,t}_{t+s}, \mathcal{U}^{i,t}_{t+s})ds\right) f_c(c^{i,t}_{t+h}, \mathcal{U}^{i,t}_{t+h}) \tag{D.10}$$

her marginal utility of consumption at age h in that history (Schröder and Skiadas, 1999). The building block of the term structure is the value, in units of the welfare numeraire, of the marginal change in consumption at age h:

$$\chi^{i,t}(x_0, h) = \frac{\mathbb{E}_0\left[M^{i,t}(s^{t+h})\,\partial_{\sigma_\pi} c^{i,t}_{t+h}(s^{t+h})\right]}{\lambda^{i,t}(x_0, y_0)}, \tag{D.11}$$

where $\partial_{\sigma_\pi} c^{i,t}_{t+h}$ is the derivative of consumption with respect to σ_π, holding the history of shocks fixed, and $\lambda^{i,t}(x_0, y_0)$, defined in (D.18) below, is the utility value of the welfare numeraire: dividing by it expresses the change in units that can be compared and added across agents. The welfare gains of cohort (i,t) accruing up to date T are

$$W^{i,t}(x_0,T) = \int_0^{T-t} \chi^{i,t}(x_0,h)\,dh, \tag{D.12}$$

and by the utility-gradient representation of recursive preferences, $\lim_{T\to\infty} W^{i,t}(x_0,T) = \frac{\partial_{\sigma\pi} U^{i,t}}{\lambda^{i,t}(x_0,y_0)}$: as the horizon grows, the term structure integrates to the total welfare gain of the cohort. Averaging across types with birth weights, and aggregating across cohorts (the initially alive cohort has mass one, and new cohorts are born at rate φ), we obtain

$$W^t(x_0,T) = \bar{\alpha} W^{A,t}(x_0,T) + (1-\bar{\alpha}) W^{B,t}(x_0,T), \tag{D.13}$$

$$W(x_0,T) = W^0(x_0,T) + \varphi \int_0^T W^t(x_0,T)dt. \tag{D.14}$$

$W(x_0,T)$ measures the efficiency gains generated by consumption changes up to date T, counting all cohorts born before T; as $T\to\infty$, it converges to the total efficiency gains.

The decomposition above fixes a cohort t and integrates over its ages h. Alternatively, we can fix a calendar date τ and aggregate the gains accruing at τ across all cohorts alive at that date. Since a cohort born at $t \le \tau$ has age $\tau - t$ at date τ, we define

$$\mathcal{W}^i(x_0,\tau) = \chi^{i,0}(x_0,\tau) + \varphi \int_0^\tau \chi^{i,t}(x_0,\tau - t)\,dt, \tag{D.15}$$

$$\mathcal{W}(x_0,\tau) = \bar{\alpha}\mathcal{W}^A(x_0,\tau) + (1-\bar{\alpha})\mathcal{W}^B(x_0,\tau). \tag{D.16}$$

That is, we take the changes in consumption at date τ in all states, value them with the appropriate marginal utility $M^{i,t}$, divide by the appropriate marginal valuation $\lambda^{i,t}$ so that everything is in the same units, and integrate across all cohorts born before τ (cohorts born after τ contribute zero). The object $\mathcal{W}(x_0,\tau)$ is the flow of welfare gains accruing at date τ, in units of the welfare numeraire: this is what we refer to as the term structure of welfare gains. The two decompositions are consistent: changing the order of integration over the triangle $\{(t,\tau) : 0 \le t \le \tau \le T\}$, with $\tau = t + h$,

$$W(x_0,T) = \int_0^T \mathcal{W}(x_0,\tau)d\tau. \tag{D.17}$$

Computing the term structure therefore requires two inputs: the marginal valuations $\lambda^{i,t}(x_0,y_0)$ in the denominator of (D.11), and the consumption-change strips $\mathbb{E}_0[M^{i,t}\,\partial_{\sigma\pi} c^{i,t}_{t+h}]$ in its numerator. We compute each in turn.

D.3.2 Computing $\lambda^{i,t}$

The object $\lambda^{i,t}(x_0, y_0)$ is the value, in utils of an agent of type i born at time t, of a claim to permanent consumption, i.e. a claim that delivers aggregate output at every date and state, when the initial state is (x_0, y_0):

$$\lambda^{i,t}(x_0, y_0) = \mathbb{E}_0 \left[\int_0^\infty M^{i,t}(s^{t+h}) y_{t+h} dh \right]. \tag{D.18}$$

This is the welfare numeraire of Barcons et al. (2026). We compute it in two steps: we first value the output claim at the cohort's birth date, and then bring that value back to time 0.

The value of the output claim at birth. By the law of iterated expectations, conditioning on the state at birth,

$$\lambda^{i,t}(x_0, y_0) = \mathbb{E}_0 \left[M_t^{i,t}\, y_t\, q^{a,i}(x_t) \right], \qquad q^{a,i}(x) = \mathbb{E} \left[\int_0^\infty \frac{M_{t+h}^{i,t}}{M_t^{i,t}} \frac{y_{t+h}}{y_t} dh \,\middle|\, x_t = x \right], \tag{D.19}$$

where $y_t q^{a,i}(x_t)$ is the value in goods, at birth, of the claim to aggregate output from birth onwards, and $M_t^{i,t}$ converts goods at birth into utils at time 0. The ratio $M_{t+h}^{i,t}/M_t^{i,t}$ is the same for all agents of type i and depends only on the path of the aggregate state, with dynamics

$$\frac{dM_u^i}{M_u^i} = -(r^i(x_u) + \varphi) du - (\kappa^i(x_u))^\top dW_u, \tag{D.20}$$

where r^i and κ^i are computed in Appendix C.4. Note that this is the marginal utility process of the agent, whose drift includes the mortality discount φ; it differs from the stochastic discount factor that prices traded assets, whose drift is $-r^i$, because annuities compensate surviving agents for mortality risk (Gârleanu and Panageas, 2015).

Computationally, we work with the finite-horizon claim $q^{a,i}(x, \tau') = \mathbb{E}[\int_0^{\tau'} \frac{M_{t+h}^{i,t}}{M_t^{i,t}} \frac{y_{t+h}}{y_t} dh | x_t = x]$, the value of aggregate output over the next τ' years. It satisfies the pricing equation

$$0 = y_u du + \mathbb{E}_u \left[d\left(M_u^i\, y_u\, q^{a,i}(x_u, \tau') \right) \right] / M_u^i, \tag{D.21}$$

where the remaining horizon τ' decreases one for one with time. Applying Itô's lemma to the product, and writing $\Sigma_y = [\sigma, 0]^\top$ for the loading of output growth on the shocks,

$$\mathbb{E}_u \left[\frac{d(M^i y q^{a,i})}{M^i y q^{a,i}} \right] = \left(-\frac{q_{\tau'}^{a,i}}{q^{a,i}} - (r^i + \varphi) + \mu - \Sigma_y^\top \kappa^i + \frac{(\mu_x - \Sigma_x \kappa^i + \Sigma_x \Sigma_y)^\top \nabla_x q^{a,i}}{q^{a,i}} + \frac{1}{2} \frac{Tr(\Sigma_x^\top (\nabla_x^2 q^{a,i}) \Sigma_x)}{q^{a,i}} \right)$$

so replacing in (D.21) we obtain

$$q^{a,i}_{\tau'} = 1 + \left(\mu - r^i - \varphi - \Sigma_y^\top \kappa^i\right) q^{a,i} + \left(\mu_x - \Sigma_x \kappa^i + \Sigma_x \Sigma_y\right)^\top \nabla_x q^{a,i} + \frac{1}{2} Tr(\Sigma_x^\top (\nabla_x^2 q^{a,i}) \Sigma_x), \tag{D.22}$$

which we solve forward in τ' from the initial condition $q^{a,i}(x,0) = 0$, iterating until $q^{a,i}(x,\tau')$ converges; the limit is $q^{a,i}(x)$.

From birth to time 0. Marginal utility at birth is $M^{i,t}_t = f_c(c^{i,t}_t, \mathcal{U}^{i,t}_t) = (n^{i,t}_t)^{-\gamma_i} (V^i(x_t))^{\frac{1-\gamma_i}{1-\psi_i}}$, using the expression for f_c derived in Appendix C.4. Wealth at birth is the value of the trees the agent is born with, $n^{i,t}_t = s^{i,t} y_t \tilde{q}(x_t)$, where $s^{i,t}$ denotes the number of trees: $s^{i,t} = \delta_q/\varphi$ for cohorts $t > 0$, and $s^{A,0} = \alpha_0/\bar{\alpha}$, $s^{B,0} = (1-\alpha_0)/(1-\bar{\alpha})$ for the initially alive agents. Substituting into (D.19),

$$\lambda^{i,t}(x_0, y_0) = (s^{i,t})^{-\gamma_i} \mathbb{E}_0 \left[y_t^{1-\gamma_i} k^i(x_t) \right], \qquad k^i(x) \equiv \tilde{q}(x)^{-\gamma_i} (V^i(x))^{\frac{1-\gamma_i}{1-\psi_i}} q^{a,i}(x). \tag{D.23}$$

Next, define

$$\ell^i(x, \tau') = \mathbb{E} \left[\left(\frac{y_{t+\tau'}}{y_t} \right)^{1-\gamma_i} k^i(x_{t+\tau'}) \,\middle|\, x_t = x \right], \tag{D.24}$$

which does not depend on y_t because $y_{t+\tau'}/y_t$ only depends on Brownian increments after t, which are independent of (x_t, y_t). Then

$$\lambda^{i,t}(x_0, y_0) = (s^{i,t})^{-\gamma_i}\, y_0^{1-\gamma_i}\, \ell^i(x_0, t). \tag{D.25}$$

To characterize ℓ^i, fix the birth date t and let $L^i(x_u, y_u, \tau') = y_u^{1-\gamma_i} \ell^i(x_u, \tau')$ for $u \le t$, where $\tau' = t - u$ is the time left to the birth date. By the law of iterated expectations, L^i is a martingale along the path, so $\mathbb{E}_u[dL^i] = 0$. Suppressing the i index, Itô's lemma gives

$$\begin{aligned} d(y^{1-\gamma_i}) &= y^{1-\gamma_i} \left((1-\gamma_i) \left(\mu - \gamma_i \frac{\sigma^2}{2} \right) du + (1-\gamma_i) \Sigma_y^\top dW \right), \\ d\ell &= \left(-\ell_{\tau'} + (\nabla_x \ell)^\top \mu_x + \frac{1}{2} Tr(\Sigma_x^\top (\nabla_x^2 \ell) \Sigma_x) \right) du + (\nabla_x \ell)^\top \Sigma_x dW, \end{aligned}$$

and the product rule $dL = d(y^{1-\gamma_i})\ell + y^{1-\gamma_i} d\ell + d(y^{1-\gamma_i}) d\ell$ together with $\mathbb{E}_u[dL] = 0$ yields the PDE

$$\ell_{\tau'} = (1-\gamma_i) \left(\mu - \gamma_i \frac{\sigma^2}{2} \right) \ell + (\nabla_x \ell)^\top \left(\mu_x + (1-\gamma_i) \Sigma_x \Sigma_y \right) + \frac{1}{2} Tr(\Sigma_x^\top (\nabla_x^2 \ell) \Sigma_x), \tag{D.26}$$

which we solve forward in τ' from the initial condition $\ell^i(x,0) = k^i(x)$. Solving up to $\tau' = t$ delivers $\lambda^{i,t}(x_0, y_0)$ for every cohort t in a single pass. The same PDE, with the equilibrium objects of the economy without inflation risk, delivers the corresponding valuations in the counterfactual economy.

With $\lambda^{i,t}$ in hand, the denominator of (D.11) is known. It remains to compute the numerator, the consumption-change strips.

D.3.3 Present Value of a Consumption Perturbation

By iterated expectations, the numerator of (D.11) equals $\mathbb{E}_0[\mathbb{E}_t[M^{i,t}_{t+h}\,\partial_{\sigma_\pi} c^{i,t}_{t+h}]]$: we first characterize the expectation conditional on the state at birth, and then take its date-zero expectation.

Consumption and its pathwise derivative. Consumption at age h satisfies

$$c^{i,t}_{t+h} = V^i(x_{t+h}) n^{i,t}_{t+h}. \tag{D.27}$$

The drift and volatility of consumption are the same for all surviving agents of type i, irrespective of their birth date. We write the state, level-consumption, and output processes as

$$dx_u = \mu_x(x_u;\sigma_\pi)\,du + \Sigma_x(x_u;\sigma_\pi)\,dW_u, \tag{D.28}$$

$$\frac{dc^{i,t}_u}{c^{i,t}_u} = \mu^i_c(x_u;\sigma_\pi)\,du + \Big(\Sigma^i_c(x_u;\sigma_\pi)\Big)^\top dW_u, \tag{D.29}$$

$$\frac{dy_u}{y_u} = \mu\,du + \Sigma_y^\top dW_u. \tag{D.30}$$

Analogously to $\tilde{q} = q/y$, define consumption normalized by aggregate output as

$$\tilde{c}^{i,t}_u \equiv \frac{c^{i,t}_u}{y_u}. \tag{D.31}$$

Itô's lemma yields

$$\frac{d\tilde{c}^{i,t}_u}{\tilde{c}^{i,t}_u} = \Big(\mu^i_c - \mu + \sigma^2 - (\Sigma^i_c)^\top \Sigma_y\Big)\,du + \Big(\Sigma^i_c - \Sigma_y\Big)^\top dW_u. \tag{D.32}$$

Aggregate output does not depend on σ_π. Hence

$$\partial_{\sigma_\pi} \log c^{i,t}_u = \partial_{\sigma_\pi} \log \tilde{c}^{i,t}_u.$$

Thus the level and output-normalized consumption processes have the same proportional sensitivity to σ_π.

Hold the Brownian path fixed and define the state and proportional-consumption sensitivities directly:

$$z_u \equiv \partial_{\sigma_\pi} x_u, \qquad \eta_u^i \equiv \partial_{\sigma_\pi} \log c_u^{i,t} = \partial_{\sigma_\pi} \log \tilde{c}_u^{i,t}. \tag{D.33}$$

For any differentiable vector- or matrix-valued function f, let

$$D_x f[z] \equiv \frac{d}{d\varepsilon} f(x + \varepsilon z)\bigg|_{\varepsilon=0}$$

denote its directional derivative along z. Since

$$d \log c_u^{i,t} = \left[\mu_c^i - \frac{1}{2}(\Sigma_c^i)^\top \Sigma_c^i\right] du + (\Sigma_c^i)^\top dW_u, \tag{D.34}$$

differentiating the state and log-consumption equations path by path gives

$$dz_u = [\partial_{\sigma_\pi}\mu_x + D_x\mu_x[z_u]]\, du + [\partial_{\sigma_\pi}\Sigma_x + D_x\Sigma_x[z_u]]\, dW_u, \tag{D.35}$$

$$\begin{aligned} d\eta_u^i = &\left[\partial_{\sigma_\pi}\mu_c^i + D_x\mu_c^i[z_u] - (\Sigma_c^i)^\top \left(\partial_{\sigma_\pi}\Sigma_c^i + D_x\Sigma_c^i[z_u]\right)\right] du \\ &+ \left[\partial_{\sigma_\pi}\Sigma_c^i + D_x\Sigma_c^i[z_u]\right]^\top dW_u. \end{aligned} \tag{D.36}$$

All coefficients from this point onward are evaluated at the baseline equilibrium, although we suppress that argument.

Because $x = (x_1, x_2) = (\alpha, \pi)$ and there are two Brownian shocks, z_u and a_z are 2×1 vectors, while b_z and c_z are 2×2 matrices. Define

$$a_z \equiv \partial_{\sigma_\pi}\mu_x, \qquad b_z \equiv \nabla_x \mu_x, \qquad c_z \equiv \partial_{\sigma_\pi}\Sigma_x, \qquad D_{z,j} \equiv \partial_{x_j}\Sigma_x, \quad j = 1, 2. \tag{D.37}$$

The two matrices $D_{z,1}$ and $D_{z,2}$ are the state-coordinate slices of a $2 \times 2 \times 2$ tensor.[26] More

[26]Numerically, the state-gradient coefficients (b_z, $D_{z,j}$, the analogous consumption-drift and volatility gradients below, and $\nabla_x \log f^i$ in the pre-birth correction) are damped to zero over the outermost grid cells: at a reflecting boundary the state cannot move through the edge, and one-sided finite-difference gradients there run through solver boundary artifacts.

explicitly,

$$a_z = \begin{pmatrix} \partial_{\sigma_\pi}\mu_\alpha \\ \partial_{\sigma_\pi}\mu_\pi \end{pmatrix}, \qquad b_z = \begin{pmatrix} \partial_\alpha\mu_\alpha & \partial_\pi\mu_\alpha \\ \partial_\alpha\mu_\pi & \partial_\pi\mu_\pi \end{pmatrix}, \tag{D.38}$$

$$c_z = \begin{pmatrix} \partial_{\sigma_\pi}\Sigma_x^{1,1} & \partial_{\sigma_\pi}\Sigma_x^{1,2} \\ \partial_{\sigma_\pi}\Sigma_x^{2,1} & \partial_{\sigma_\pi}\Sigma_x^{2,2} \end{pmatrix}, \qquad D_{z,j} = \begin{pmatrix} \partial_{x_j}\Sigma_x^{1,1} & \partial_{x_j}\Sigma_x^{1,2} \\ \partial_{x_j}\Sigma_x^{2,1} & \partial_{x_j}\Sigma_x^{2,2} \end{pmatrix}. \tag{D.39}$$

Thus equation (D.35) becomes

$$dz_u = (a_z + b_z z_u)\, du + (c_z + D_{z,1}z_{u,1} + D_{z,2}z_{u,2})\, dW_u. \tag{D.40}$$

Conditional on the cohort's birth state, $z_t = 0$: the strip below prices the perturbation for a given state at birth. For cohorts born at $t > 0$ the birth state itself responds to σ_π; we return to this when we bring the strip to date zero.

Define

$$\begin{aligned} a_\eta^i &\equiv \partial_{\sigma_\pi}\mu_c^i - (\Sigma_c^i)^\top \partial_{\sigma_\pi}\Sigma_c^i, & (b_\eta^i)_j &\equiv \partial_{x_j}\mu_c^i - (\Sigma_c^i)^\top \partial_{x_j}\Sigma_c^i, \\ c_\eta^i &\equiv \partial_{\sigma_\pi}\Sigma_c^i, & (D_\eta^i)_{j\cdot} &\equiv (\partial_{x_j}\Sigma_c^i)^\top. \end{aligned} \tag{D.41}$$

Here a_η^i is a scalar, b_η^i and c_η^i are 2×1 vectors, and D_η^i is a 2×2 matrix. Equation (D.36) can then be written as

$$d\eta_u^i = \left[a_\eta^i + z_u^\top b_\eta^i\right] du + \left[(c_\eta^i)^\top + z_u^\top D_\eta^i\right] dW_u. \tag{D.42}$$

Conditional on the birth state x_t, a newborn begins with

$$c_t^{i,t} = s^{i,t} y_t \tilde{q}(x_t) V^i(x_t), \qquad z_t = 0, \qquad \eta_t^i = \eta_0^i(x_t), \tag{D.43}$$

where

$$\eta_0^i(x) \equiv \partial_{\sigma_\pi}\log\tilde{q}(x) + \partial_{\sigma_\pi}\log V^i(x). \tag{D.44}$$

The endowment $s^{i,t}$ and current output y_t are held fixed in this local derivative.

Pricing the perturbation. The marginal utility process of a surviving type-i agent satisfies

$$\frac{dM_u^{i,t}}{M_u^{i,t}} = -\widehat{r}^i(x_u)\, du - (\kappa^i(x_u))^\top dW_u, \qquad \widehat{r}^i \equiv r^i + \varphi. \tag{D.45}$$

The survival adjustment is required here: the drift μ_c^i includes the annuity income φ, while the marginal utility of a particular cohort is discounted by its mortality rate.

Using $\partial_{\sigma_\pi} c^{i,t}_{t+h} = c^{i,t}_{t+h}\eta^i_{t+h}$ and factoring out date-t marginal utility and consumption gives

$$\begin{aligned}
&\mathbb{E}_t\left[M^{i,t}_{t+h}\,\partial_{\sigma_\pi} c^{i,t}_{t+h}\right] \\
&\quad = \left(s^{i,t} y_t \tilde{q}(x_t)\right)^{1-\gamma_i}\left(V^i(x_t)\right)^{1+\frac{1-\gamma_i}{1-\psi_i}} G^i(x_t, \eta^i_t, z_t, h),
\end{aligned} \tag{D.46}$$

where, letting v measure age from the cohort's birth date,

$$\begin{aligned}
G^i(x,\eta,z,h) \equiv \mathbb{E}\Bigg[\exp\Bigg\{&\int_0^h \left[-\widehat{r}^i - \frac{1}{2}(\kappa^i)^\top\kappa^i + \mu^i_c - \frac{1}{2}(\Sigma^i_c)^\top\Sigma^i_c\right](x_{t+v})\,dv \\
&+ \int_0^h \left(-\kappa^i + \Sigma^i_c\right)^\top (x_{t+v})\,dW_{t+v}\Bigg\}\eta^i_{t+h} \;\Bigg|\; x_t = x, \eta^i_t = \eta, z_t = z\Bigg].
\end{aligned} \tag{D.47}$$

The drift and volatility of η^i do not depend on its current level, and the SDE for z is affine in z. This motivates the affine conjecture

$$G^i(x,\eta,z,h) = A^i(x,h)\eta + (B^i(x,h))^\top z + C^i(x,h). \tag{D.48}$$

We now derive, rather than simply state, the equations for these coefficients. For the local expansion, reset the clock at the conditioning date and denote the infinitesimal increment by du. The exponential factor in equation (D.47) satisfies

$$\begin{aligned}
&\exp\left\{\left[-\widehat{r}^i - \frac{1}{2}(\kappa^i)^\top\kappa^i + \mu^i_c - \frac{1}{2}(\Sigma^i_c)^\top\Sigma^i_c\right]du + (\Sigma^i_c - \kappa^i)^\top dW_u\right\} \\
&\quad = 1 + \left[\mu^i_c - \widehat{r}^i - (\kappa^i)^\top\Sigma^i_c\right]du + (\Sigma^i_c - \kappa^i)^\top dW_u + o(du).
\end{aligned} \tag{D.49}$$

The dynamic-programming step is therefore

$$\begin{aligned}
G^i(x,\eta,z,h) = \mathbb{E}\Bigg[\Bigg(1 &+ \left[\mu^i_c - \widehat{r}^i - (\kappa^i)^\top\Sigma^i_c\right]du \\
&+ (\Sigma^i_c - \kappa^i)^\top dW_u\Bigg) \\
&\times G^i(x+dx, \eta+d\eta, z+dz, h-du)\Bigg] + o(du).
\end{aligned} \tag{D.50}$$

Under the affine conjecture, $G^i_{\eta\eta}$, G^i_{zz}, and $G^i_{\eta z}$ are zero. The Itô expansion inside equation (D.50) is

$$
\begin{aligned}
G^i(x+dx,\eta+d\eta,z+dz,h-du) = G^i &+ (G^i_x)^\top dx + \frac{1}{2}\operatorname{Tr}\left(\Sigma_x\Sigma_x^\top G^i_{xx}\right) du \\
&+ (G^i_z)^\top dz + G^i_\eta\, d\eta - G^i_h\, du \\
&+ (dx)^\top G^i_{xz}\, dz + (dx)^\top G^i_{x\eta}\, d\eta + o(du).
\end{aligned} \tag{D.51}
$$

The jth column of G^i_{xz} is $\nabla_x G^i_{z_j}$. Multiplying equations (D.49) and (D.51), taking conditional expectations, cancelling G^i from both sides, and dividing by du yields the general equation

$$
\begin{aligned}
G^i_h = (G^i_x)^\top &\left[\mu_x + \Sigma_x(\Sigma^i_c - \kappa^i)\right] + \frac{1}{2}\operatorname{Tr}\left(\Sigma_x\Sigma_x^\top G^i_{xx}\right) \\
&+ (G^i_z)^\top(a_z + b_z z) + G^i_\eta\left[a^i_\eta + (b^i_\eta)^\top z\right] \\
&+ \operatorname{Tr}\left[(c_z + D_{z,1}z_1 + D_{z,2}z_2)\,\Sigma_x^\top G^i_{xz}\right] \\
&+ (G^i_{x\eta})^\top\Sigma_x\left[c^i_\eta + (D^i_\eta)^\top z\right] + \left[\mu^i_c - \hat{r}^i - (\kappa^i)^\top\Sigma^i_c\right] G^i \\
&+ (G^i_z)^\top\,(c_z + D_{z,1}z_1 + D_{z,2}z_2)\,(\Sigma^i_c - \kappa^i) + G^i_\eta\left[(c^i_\eta)^\top + z^\top D^i_\eta\right](\Sigma^i_c - \kappa^i).
\end{aligned} \tag{D.52}
$$

We can now match the coefficients on η, z_1, z_2, and the constant. The terms proportional to η give

$$
\begin{aligned}
\partial_h A^i = (\nabla_x A^i)^\top &\left[\mu_x + \Sigma_x(\Sigma^i_c - \kappa^i)\right] + \frac{1}{2}\operatorname{Tr}\left(\Sigma_x\Sigma_x^\top\nabla^2_x A^i\right) \\
&+ \left[\mu^i_c - \hat{r}^i - (\kappa^i)^\top\Sigma^i_c\right] A^i, \qquad A^i(x,0) = 1.
\end{aligned} \tag{D.53}
$$

This is the price of a unit perturbation already present at the beginning of the strip.

For the terms proportional to z, write

$$
(B^i)^\top z = B^i_1 z_1 + B^i_2 z_2, \qquad J^i_B \equiv \left[\nabla_x B^i_1 \quad \nabla_x B^i_2\right].
$$

It is useful to identify each contribution before collecting them. First,

$$
(B^i)^\top b_z z = \left(B^i_1 b^{11}_z + B^i_2 b^{21}_z\right) z_1 + \left(B^i_1 b^{12}_z + B^i_2 b^{22}_z\right) z_2, \tag{D.54}
$$

while $A^i(b^i_\eta)^\top z$ contributes $A^i(b^i_\eta)_j$ to the coefficient on z_j. Second, the x–z covariance term expands as

$$
\begin{aligned}
&\operatorname{Tr}\left[(c_z + D_{z,1}z_1 + D_{z,2}z_2)\,\Sigma_x^\top J^i_B\right] \\
&\quad= \operatorname{Tr}\left(c_z\Sigma_x^\top J^i_B\right) + \operatorname{Tr}\left(D_{z,1}\Sigma_x^\top J^i_B\right) z_1 + \operatorname{Tr}\left(D_{z,2}\Sigma_x^\top J^i_B\right) z_2.
\end{aligned} \tag{D.55}
$$

Third, the x–η covariance contributes

$$(\nabla_x A^i)^\top \Sigma_x c^i_\eta + (\nabla_x A^i)^\top \Sigma_x (D^i_\eta)^\top_{j\cdot} z_j,$$

where repeated j in this displayed expression denotes the sum over the two state coordinates. Finally, the covariance with the exponential factor contributes

$$\begin{aligned}&(B^i)^\top \left(c_z + D_{z,1} z_1 + D_{z,2} z_2\right)\left(\Sigma^i_c - \kappa^i\right)\\ &\quad + A^i \left[(c^i_\eta)^\top + z^\top D^i_\eta\right]\left(\Sigma^i_c - \kappa^i\right).\end{aligned} \tag{D.56}$$

Collecting the coefficient on z_1 gives

$$\begin{aligned}\partial_h B^i_1 = {} & (\nabla_x B^i_1)^\top \left[\mu_x + \Sigma_x(\Sigma^i_c - \kappa^i)\right] + \frac{1}{2}\operatorname{Tr}\left(\Sigma_x \Sigma_x^\top \nabla^2_x B^i_1\right) + \left[\mu^i_c - \widehat{r}^i - (\kappa^i)^\top \Sigma^i_c\right] B^i_1\\ & + B^i_1 b^{11}_z + B^i_2 b^{21}_z + A^i (b^i_\eta)_1 + \operatorname{Tr}\left(D_{z,1}\Sigma_x^\top J^i_B\right)\\ & + (\nabla_x A^i)^\top \Sigma_x (D^i_\eta)^\top_{1\cdot} + \left[(B^i)^\top D_{z,1} + A^i (D^i_\eta)_{1\cdot}\right]\left(\Sigma^i_c - \kappa^i\right), \qquad B^i_1(x,0) = 0,\end{aligned} \tag{D.57}$$

and the coefficient on z_2 gives

$$\begin{aligned}\partial_h B^i_2 = {} & (\nabla_x B^i_2)^\top \left[\mu_x + \Sigma_x(\Sigma^i_c - \kappa^i)\right] + \frac{1}{2}\operatorname{Tr}\left(\Sigma_x \Sigma_x^\top \nabla^2_x B^i_2\right) + \left[\mu^i_c - \widehat{r}^i - (\kappa^i)^\top \Sigma^i_c\right] B^i_2\\ & + B^i_1 b^{12}_z + B^i_2 b^{22}_z + A^i (b^i_\eta)_2 + \operatorname{Tr}\left(D_{z,2}\Sigma_x^\top J^i_B\right)\\ & + (\nabla_x A^i)^\top \Sigma_x (D^i_\eta)^\top_{2\cdot} + \left[(B^i)^\top D_{z,2} + A^i (D^i_\eta)_{2\cdot}\right]\left(\Sigma^i_c - \kappa^i\right), \qquad B^i_2(x,0) = 0.\end{aligned} \tag{D.58}$$

Thus B^i_1 and B^i_2 must be solved jointly.

The remaining constant terms give

$$\begin{aligned}\partial_h C^i = {} & (\nabla_x C^i)^\top \left[\mu_x + \Sigma_x(\Sigma^i_c - \kappa^i)\right] + \frac{1}{2}\operatorname{Tr}\left(\Sigma_x \Sigma_x^\top \nabla^2_x C^i\right) + \left[\mu^i_c - \widehat{r}^i - (\kappa^i)^\top \Sigma^i_c\right] C^i\\ & + (B^i)^\top a_z + A^i a^i_\eta + \operatorname{Tr}\left(c_z \Sigma_x^\top J^i_B\right) + (\nabla_x A^i)^\top \Sigma_x c^i_\eta\\ & + \left[(B^i)^\top c_z + A^i (c^i_\eta)^\top\right]\left(\Sigma^i_c - \kappa^i\right), \qquad C^i(x,0) = 0.\end{aligned} \tag{D.59}$$

The recursive order is therefore to solve A^i first, solve B^i_1 and B^i_2 jointly, and then solve C^i. Equivalently, all four equations can be advanced in a single horizon loop.

Conditional on the birth state, $z_t = 0$ and $\eta_t^i = \eta_0^i(x_t)$. Substitution into equation (D.46) yields

$$\mathbb{E}_t\left[M_{t+h}^{i,t}\,\partial_{\sigma_\pi} c_{t+h}^{i,t}\right] = (s^{i,t}y_t)^{1-\gamma_i} F^i(x_t, h), \tag{D.60}$$

where

$$F^i(x, h) \equiv \tilde{q}(x)^{1-\gamma_i} V^i(x)^{1+\frac{1-\gamma_i}{1-\psi_i}} \left[A^i(x, h)\eta_0^i(x) + C^i(x, h)\right]. \tag{D.61}$$

Thus, $F^i(x, h)$ is the birth-date value of type i's marginal consumption response at age h, conditional on the cohort being born in state x and net of the scale factor $(s^{i,t}y_t)^{1-\gamma_i}$. It is a welfare strip: it values the response at the single date $t+h$, rather than the cumulative response up to that date. Although B^i does not appear directly in equation (D.61) because $z_t = 0$, it enters through the equation for C^i, and it reappears directly in the pre-birth correction for unborn cohorts below. The term containing η_0^i captures the initial change in consumption, while C^i captures the subsequent changes in consumption growth and in the state path.

The pre-birth state sensitivity. It remains to take the outer expectation in equation (D.11). The calculation so far prices the strip conditional on the cohort's birth state: setting $z_t = 0$ at birth treats x_t as given. That is the correct initial condition for the initially alive cohort, whose state x_0 is fixed. It is not the correct initial condition for a cohort born at $t > 0$: holding the history of shocks fixed, a change in σ_π alters the path of the state, so the perturbed economy reaches the cohort's birth date in a different state. The pre-birth sensitivity $z_t = \partial_{\sigma_\pi} x_t$ is generated by equation (D.40) from $z_0 = 0$, and it is not zero at $t > 0$. Because consumption at birth is $c_t^{i,t} = s^{i,t}y_t f^i(x_t)$ with $f^i \equiv \tilde{q}V^i$, the correct birth conditions for a cohort born at $t > 0$ are

$$z_t = \partial_{\sigma_\pi} x_t, \qquad \eta_t^i = \eta_0^i(x_t) + \left(\nabla_x \log f^i(x_t)\right)^\top z_t. \tag{D.62}$$

Substituting these into equation (D.46) with the affine representation (D.48) gives

$$\mathbb{E}_t\left[M_{t+h}^{i,t}\,\partial_{\sigma_\pi} c_{t+h}^{i,t}\right] = (s^{i,t}y_t)^{1-\gamma_i}\left[F^i(x_t, h) + \left(\tilde{B}^i(x_t, h)\right)^\top z_t\right], \tag{D.63}$$

where

$$\tilde{B}^i(x, h) \equiv \tilde{q}(x)^{1-\gamma_i} V^i(x)^{1+\frac{1-\gamma_i}{1-\psi_i}} \left[A^i(x, h)\,\nabla_x \log f^i(x) + B^i(x, h)\right]. \tag{D.64}$$

Setting $z_t = 0$ keeps only the F^i term. This is where the coefficients B^i finally enter directly: they price the effect of a birth-state displacement on the entire subsequent consumption path, while the $A^i \nabla_x \log f^i$ term prices its effect on consumption at birth.

Bringing the perturbation to date zero. Taking the date-zero expectation of equation (D.63) and dividing by $\lambda^{i,t}(x_0, y_0) = (s^{i,t})^{-\gamma_i} y_0^{1-\gamma_i} \ell^i(x_0, t)$ from equation (D.25),

$$\chi^{i,t}(x_0, h) = s^{i,t} \, \frac{K^i(x_0, t, h) + R^i(x_0, t, h)}{\ell^i(x_0, t)}, \tag{D.65}$$

with

$$K^i(x_0, t, h) \equiv \mathbb{E}_0 \left[\left(\frac{y_t}{y_0} \right)^{1-\gamma_i} F^i(x_t, h) \;\middle|\; x_0 \right], \qquad R^i(x_0, t, h) \equiv \mathbb{E}_0 \left[\left(\frac{y_t}{y_0} \right)^{1-\gamma_i} \left(\tilde{B}^i(x_t, h) \right)^\top z_t \;\middle|\; x_0 \right]. \tag{D.66}$$

For each fixed age h, K^i has exactly the same Feynman–Kac structure as ℓ^i. In particular,

$$\begin{aligned} \partial_t K^i = (1 - \gamma_i) \left(\mu - \frac{\gamma_i \sigma^2}{2} \right) K^i + (\nabla_x K^i)^\top \left[\mu_x + (1 - \gamma_i) \Sigma_x \Sigma_y \right] \\ + \frac{1}{2} \operatorname{Tr} \left(\Sigma_x \Sigma_x^\top \nabla_x^2 K^i \right), \qquad K^i(x, 0, h) = F^i(x, h). \end{aligned} \tag{D.67}$$

This is the backward pricing operator, advanced from zero to the cohort birth date t; it is not a recursion in the age h.

The correction R^i is not an expectation of a function of x_t alone: z_t depends on the whole path of the state. The affine structure saves us once more. The SDE (D.40) for z is affine in z, so the function

$$N^i(x, z, t, h) \equiv \mathbb{E} \left[\left(\frac{y_t}{y_0} \right)^{1-\gamma_i} \left(\tilde{B}^i(x_t, h) \right)^\top z_t \;\middle|\; x_0 = x, \, z_0 = z \right] \tag{D.68}$$

is affine in the initial sensitivity,

$$N^i(x, z, t, h) = \left(P^i(x, t, h) \right)^\top z + R^i(x, t, h), \tag{D.69}$$

and the correction we need is its intercept, $R^i(x_0, t, h) = N^i(x_0, 0, t, h)$. Repeating the dynamic-programming argument used for G^i — now under the output-tilted expectation, whose effective drift is $\mu_x + (1 - \gamma_i) \Sigma_x \Sigma_y$, whose growth rate is $(1 - \gamma_i)(\mu - \gamma_i \sigma^2 / 2)$, and whose Brownian tilt contracts volatilities against $(1 - \gamma_i) \Sigma_y$ — and matching the coefficients on z_1, z_2, and the constant gives, writing $(P^i)^\top z = P_1^i z_1 + P_2^i z_2$ and $J_P^i \equiv [\nabla_x P_1^i \;\; \nabla_x P_2^i]$,

$$\begin{aligned}\partial_t P_j^i &= (\nabla_x P_j^i)^\top [\mu_x + (1-\gamma_i)\Sigma_x \Sigma_y] + \frac{1}{2} \operatorname{Tr}\left(\Sigma_x \Sigma_x^\top \nabla_x^2 P_j^i\right) + (1-\gamma_i)\left(\mu - \frac{\gamma_i \sigma^2}{2}\right) P_j^i \\ &\quad + P_1^i b_z^{1j} + P_2^i b_z^{2j} + \operatorname{Tr}\left(D_{z,j} \Sigma_x^\top J_P^i\right) + (1-\gamma_i)\left[(P^i)^\top D_{z,j}\right]\Sigma_y, \qquad P_j^i(x,0,h) = \tilde{B}_j^i(x,h),\end{aligned} \tag{D.70}$$

for $j = 1, 2$, and

$$\begin{aligned}\partial_t R^i &= (\nabla_x R^i)^\top [\mu_x + (1-\gamma_i)\Sigma_x \Sigma_y] + \frac{1}{2} \operatorname{Tr}\left(\Sigma_x \Sigma_x^\top \nabla_x^2 R^i\right) + (1-\gamma_i)\left(\mu - \frac{\gamma_i \sigma^2}{2}\right) R^i \\ &\quad + (P^i)^\top a_z + \operatorname{Tr}\left(c_z \Sigma_x^\top J_P^i\right) + (1-\gamma_i)(P^i)^\top c_z \Sigma_y, \qquad R^i(x,0,h) = 0.\end{aligned} \tag{D.71}$$

These are the B^i and C^i equations of the strip system with the pricing drift $\mu_x + \Sigma_x(\Sigma_c^i - \kappa^i)$ replaced by the tilted drift, the discount rate $\mu_c^i - \widehat{r}^i - (\kappa^i)^\top \Sigma_c^i$ replaced by the tilt growth rate, the trailing $(\Sigma_c^i - \kappa^i)$ contractions replaced by $(1-\gamma_i)\Sigma_y$, and the η-sensitivity sources removed: the η channel is already inside $\tilde{B}^i$. The system is advanced forward in the birth date t, and the age h enters only through the initial condition, so all age slices propagate independently. Since $z_0 = 0$ and $R^i(x, 0, h) = 0$, the initially alive cohort is unaffected: the pre-birth correction is a property of unborn cohorts only. Quantitatively it matters: omitting R^i flips the sign of the long-run accrual of type-B agents relative to the exact cohort welfare change.

D.3.4 Term Structure of Welfare Gains

We now fix a calendar date τ and sum the welfare effects that accrue at that date to all cohorts already born. This date contribution will be denoted by $\mathcal{W}$, reserving W for welfare accumulated over dates.

Initially alive cohort. For type i, since $K^i(x_0, 0, \tau) = F^i(x_0, \tau)$ and $R^i(x_0, 0, \tau) = 0$, the date-τ contribution of an agent alive at date zero is

$$
\begin{aligned}
\mathcal{W}^i_{\text{alive}}(x_0, \tau) &\equiv \chi^{i,0}(x_0, \tau) \\
&= s^{i,0} \frac{F^i(x_0, \tau)}{\ell^i(x_0, 0)},
\end{aligned}
\tag{D.72}
$$

with

$$
s^{A,0} = \frac{\alpha_0}{\bar{\alpha}}, \qquad s^{B,0} = \frac{1-\alpha_0}{1-\bar{\alpha}}. \tag{D.73}
$$

The strip calculation has a direct lifetime-utility check:

$$
\int_0^\infty F^i(x, h)\, dh = \frac{1}{1-\gamma_i} \partial_{\sigma_\pi} \left[\tilde{q}(x)^{1-\gamma_i} V^i(x)^{\frac{1-\gamma_i}{1-\psi_i}} \right]. \tag{D.74}
$$

Multiplying both sides by $(s^{i,t} y_t)^{1-\gamma_i}$ gives the corresponding derivative of the cohort's utility at birth.

Unborn cohorts. Now fix a type i and consider agents born after date zero but before τ. Every such representative newborn receives δ_q/φ trees, so $s^{i,t} = \delta_q/\varphi$. Before multiplying by the flow of newborn cohorts, their date-τ contributions satisfy

$$
\int_0^\tau \chi^{i,t}(x_0, \tau - t)\, dt = \frac{\delta_q}{\varphi} \int_0^\tau \frac{K^i(x_0, t, \tau - t) + R^i(x_0, t, \tau - t)}{\ell^i(x_0, t)}\, dt. \tag{D.75}
$$

Future cohorts arrive at rate φ. Hence their contribution to social welfare is

$$
\mathcal{W}^i_{\text{future}}(x_0, \tau) = \delta_q \int_0^\tau \frac{K^i(x_0, t, \tau - t) + R^i(x_0, t, \tau - t)}{\ell^i(x_0, t)}\, dt. \tag{D.76}
$$

The cohort weight δ_q is distinct from the across-type weights $\bar{\alpha}$ and $1 - \bar{\alpha}$. The corrected numerator also restores the adding-up property of the utility-gradient representation cohort by cohort: integrating over all ages,

$$
\int_0^\infty \left[K^i + R^i \right] (x_0, t, h)\, dh = \frac{\partial_{\sigma_\pi} u^i(x_0, t)}{1-\gamma_i}, \tag{D.77}
$$

where u^i is the utility-before-birth function of Appendix D.2. This generalizes equation (D.74) to $t > 0$ and is the consistency check we use to validate the computation: without R^i, the left-hand side does not integrate to the cohort's utility derivative.

Computing $K^i(x_0, t, h)$ and $R^i(x_0, t, h)$ separately for every pair (t, h) would require a fine double loop. When evaluating a small set of initial states, both terms can be assembled without one: K^i through a forward weighted density, and R^i through the adjoint of the (P^i, R^i) system, as follows.

FORWARD REPRESENTATION. For a fixed x_0, define the type-specific output-tilted density $\nu_t^i(x \mid x_0)$ through

$$\mathbb{E}_0\left[\left(\frac{y_t}{y_0}\right)^{1-\gamma_i} \phi(x_t, t) \;\middle|\; x_0\right] = \int_{\mathcal{X}} \phi(x,t)\nu_t^i(x \mid x_0)\, dx \tag{D.78}$$

for every smooth test function ϕ. We now derive its forward equation. Let

$$Y_t^i \equiv \left(\frac{y_t}{y_0}\right)^{1-\gamma_i}, \qquad \beta_i \equiv (1-\gamma_i)\left(\mu - \frac{\gamma_i \sigma^2}{2}\right), \qquad a \equiv \Sigma_x \Sigma_x^\top.$$

Itô's lemma gives

$$dY_t^i = Y_t^i \left[\beta_i\, dt + (1-\gamma_i)\Sigma_y^\top dW_t\right], \tag{D.79}$$

$$d\phi(x_t,t) = \left[\phi_t + (\nabla_x \phi)^\top \mu_x + \frac{1}{2}\operatorname{Tr}\left(a \nabla_x^2 \phi\right)\right] dt + (\nabla_x \phi)^\top \Sigma_x dW_t, \tag{D.80}$$

$$dY_t^i\, d\phi(x_t,t) = Y_t^i (1-\gamma_i)(\nabla_x \phi)^\top \Sigma_x \Sigma_y\, dt. \tag{D.81}$$

Therefore

$$d\left[Y_t^i \phi(x_t,t)\right] = Y_t^i \left[\phi_t + (\nabla_x \phi)^\top b_i^Y + \frac{1}{2}\operatorname{Tr}\left(a\nabla_x^2 \phi\right) + \beta_i \phi\right] dt + d\mathcal{M}_t, \tag{D.82}$$

where $d\mathcal{M}_t$ is a martingale increment and

$$b_i^Y \equiv \mu_x + (1-\gamma_i)\Sigma_x \Sigma_y. \tag{D.83}$$

Taking expectations, using equation (D.78), and cancelling the common ϕ_t terms gives the weak equation

$$\int_0^T \int_{\mathcal{X}} \phi(x,t)\partial_t \nu_t^i(x \mid x_0)\, dx\, dt = \int_0^T \int_{\mathcal{X}} \nu_t^i(x \mid x_0) \left[(\nabla_x \phi)^\top b_i^Y + \frac{1}{2}\operatorname{Tr}\left(a\nabla_x^2 \phi\right) + \beta_i \phi\right] dx\, dt. \tag{D.84}$$

With reflecting, no-flux boundaries, integration by parts gives

$$\int_{\mathcal{X}} \nu_t^i (\nabla_x \phi)^\top b_i^Y \, dx = -\int_{\mathcal{X}} \phi \sum_j \partial_{x_j} \left(b_{i,j}^Y \nu_t^i \right) dx, \tag{D.85}$$

$$\frac{1}{2} \int_{\mathcal{X}} \nu_t^i \operatorname{Tr} \left(a \nabla_x^2 \phi \right) dx = \frac{1}{2} \int_{\mathcal{X}} \phi \sum_{j,k} \partial_{x_j x_k} \left(a_{jk} \nu_t^i \right) dx. \tag{D.86}$$

Since the test function is arbitrary, the density solves

$$\partial_t \nu_t^i = -\sum_j \partial_{x_j} \left(b_{i,j}^Y \nu_t^i \right) + \frac{1}{2} \sum_{j,k} \partial_{x_j x_k} \left(a_{jk} \nu_t^i \right) + \beta_i \nu_t^i, \qquad \nu_0^i(x \mid x_0) = \delta_{x_0}(x). \tag{D.87}$$

The factor $(1-\gamma_i)$ in b_i^Y and the factor $1/2$ in the diffusion adjoint are essential. Moreover, ν_t^i depends on the preference type through γ_i; it is not a common density for the two types. It is a weighted measure rather than a probability density, with total mass $\exp(\beta_i t)$.

The forward representation now gives

$$K^i(x_0, t, h) = \int_{\mathcal{X}} F^i(x, h) \nu_t^i(x \mid x_0) \, dx, \tag{D.88}$$

$$\ell^i(x_0, t) = \int_{\mathcal{X}} k^i(x) \nu_t^i(x \mid x_0) \, dx, \tag{D.89}$$

where k^i is the initial condition for ℓ^i defined above.

The correction term admits the same economy of computation. The (P^i, R^i) system (D.70)–(D.71) is linear and time-invariant, and the age h enters only through the initial condition $P^i(x, 0, h) = \tilde{B}^i(x, h)$. Propagating the adjoint of its one-step operator from a unit mass placed on the R-component at x_0 therefore yields, per type, a path of adjoint weights ψ_t^i carrying one pair of entries per state-grid node — one for each component of z, the two P-blocks of the adjoint state — such that, contracting over both the grid and the two sensitivity components,

$$R^i(x_0, t, h) = \langle \psi_t^i, \tilde{B}^i(\cdot, h) \rangle, \tag{D.90}$$

so one adjoint path per type replaces the full (t, h) tableau, exactly as one forward density ν_t^i per type replaces the backward K^i scans. Consequently,

$$\begin{aligned} \mathcal{W}^i(x_0, \tau) = {} & s^{i,0} \frac{F^i(x_0, \tau)}{\ell^i(x_0, 0)} \\ & + \delta_q \int_0^\tau \frac{\int_{\mathcal{X}} F^i(x, \tau - t) \nu_t^i(x \mid x_0) \, dx + \langle \psi_t^i, \tilde{B}^i(\cdot, \tau - t) \rangle}{\ell^i(x_0, t)} \, dt. \end{aligned} \tag{D.91}$$

This is the convolution in cohort birth date and cohort age that delivers the term structure at a fixed initial state.

Welfare decomposition. With $\chi^{i,t}$ in hand, the objects defined in Appendix D.3.1 follow by aggregation: equation (D.65) feeds the cohort decomposition (D.12)–(5.9) and the calendar-date decomposition (D.15)–(D.17), with $\mathcal{W}^i(x_0,\tau) = \mathcal{W}^i_{\text{alive}}(x_0,\tau) + \mathcal{W}^i_{\text{future}}(x_0,\tau)$.

All expressions above are derivatives with respect to an increase in σ_π. The first-order gain from eliminating inflation risk is therefore

$$\text{Gain}^{\text{removal}}(x_0,T) \simeq -\sigma_\pi W(x_0,T). \tag{D.92}$$

When the numerical denominator is implemented as the claim to the tree cash flow ωy rather than to aggregate output y, the reported output-unit gain is equivalently

$$\text{Gain}^{\text{removal}}(x_0,T) \simeq -\sigma_\pi \omega W_{\text{code}}(x_0,T). \tag{D.93}$$

E Efficiency Gains from a Tax on Borrowing

This appendix considers the model of Section 4 abstracting from inflation risk, and introduces a tax on borrowing. The purpose of the exercise is to show that the efficiency gains from such a tax are positive, indicating that the competitive equilibrium of the economy without inflation risk exhibits excessive borrowing. As we argued in the main text, this force reduces the welfare gains of eliminating inflation risk, since lower inflation risk increases borrowing. To isolate the corrective role of the tax from any fiscal or redistributive side effect, its proceeds are rebated lump-sum to the investors that pay it.

E.1 Environment

The economy is that of Section 4 with inflation risk shut down: $\sigma_\pi = 0$ and $\pi_0 = \bar{\pi}$, so inflation is constant at $\bar{\pi}$ and the aggregate state reduces to the wealth share of type-A investors, with law of motion $d\alpha_t = \mu_\alpha(\alpha_t)\,dt + \sigma_\alpha(\alpha_t)\,dW_{y,t}$ given by the same expressions (C.9) as in the full model. The market structure is unchanged, but returns now load on the single output shock: $\mu = [\mu_q, \mu_b]^T$ and $\Sigma = [\sigma_q, \sigma_b]^T \in \mathbb{R}^2$, all functions of α. With one shock and two freely traded assets, markets are complete among living investors: the only remaining friction is that, due to the overlapping-generations structure, unborn agents cannot trade until they are born.

The government levies a proportional tax at rate τ on borrowing (on negative positions in the long-term bond) and rebates the proceeds to the taxed investors, lump-sum in proportion to their wealth. The tax is levied on borrowing, not on saving: an investor with $\theta_b \geq 0$ pays nothing. We guess and verify that for τ low enough type-A investors borrow, $\theta_b^A < 0$, while type-B investors save, $\theta_b^B > 0$, so in equilibrium the tax falls on type-A investors.[27] The budget constraint of a (borrowing) type-A investor becomes

$$\frac{dn_t^A}{n_t^A} = r_t\,dt + (\theta_t^A)^T\left((\mu_t - r_t)\,dt + \Sigma_t\,dW_t\right) + \tau\theta_{b,t}^A\,dt + T_t\,dt - \frac{c_t^A}{n_t^A}\,dt + \varphi\,dt, \tag{E.1}$$

where $T_t = -\tau\bar{\theta}_{b,t}^A$ is the rebate per unit of wealth and $\bar{\theta}_{b,t}^A$ denotes the average bond share of type-A investors, which each investor takes as given. Because all type-A investors choose the same portfolio share, in equilibrium $\bar{\theta}_b^A = \theta_b^A$ and every investor receives back exactly what they pay: the government budget balances investor by investor, and no resources flow across types. What the investor does *not* internalize is that borrowing one unit less also lowers their rebate by τ. Type-B investors, being savers, pay no tax and keep the budget constraint (4.7). Since the rebate is proportional to wealth and depends only on aggregates, the homotheticity of the problem — and with it the scaling results of Appendix C.1 — is preserved.

E.2 Optimality Conditions

Under the guess $\theta_b^A < 0$ and $\theta_b^B > 0$, only type-A investors face the tax. The HJB equation of a type-A investor is that of Appendix C.1 with the two additional drift terms of (E.1), of which only the first, $\tau\theta_b^A$, involves a choice variable. After the same scaling $\mathcal{U}^A(n;\alpha) = \frac{n^{1-\gamma^A}}{1-\gamma^A} V^A(\alpha)^{\frac{1-\gamma^A}{1-\psi^A}}$, consumption optimality is unchanged, $c^A/n^A = V^A(\alpha)$, and the first-order condition for the portfolio becomes

$$\mu + \tau\mathbf{e}_b - r\mathbf{1} - \gamma^A\Sigma\Sigma^T\theta^A + \frac{1-\gamma^A}{1-\psi^A}\left(\frac{V^{A\prime}(\alpha)}{V^A(\alpha)}\,\sigma_\alpha\,\Sigma^T\right)^T = \xi\mathbf{1}, \tag{E.2}$$

where $\mathbf{e}_b = [0,1]^T$ and ξ is the multiplier on $\theta^T\mathbf{1} \leq 1$: the investor prices the bond as if its expected return were $\mu_b + \tau$, so a short bond position carries a perceived cost of τ per unit borrowed even though the rebate returns it ex post. Multiplying (E.2) by $[1,-1]$ to

[27] Numerically we implement the tax as a linear wedge on the type-A bond position, which coincides with the borrowing tax whenever $\theta_b^A < 0$ and $\theta_b^B \geq 0$ — verified on the whole state space — and avoids the kink at $\theta_b = 0$.

eliminate the multiplier, and doing the same for the type-B condition, we have:

$$\mu_q - \mu_b - \tau = [1, -1] \left(\gamma^A \Sigma\Sigma^T \theta^A - \frac{1-\gamma^A}{1-\psi^A} \left(\frac{V^{A\prime}}{V^A} \sigma_\alpha \, \Sigma^T \right)^T \right), \tag{E.3}$$

$$\mu_q - \mu_b = [1, -1] \left(\gamma^B \Sigma\Sigma^T \theta^B - \frac{1-\gamma^B}{1-\psi^B} \left(\frac{V^{B\prime}}{V^B} \sigma_\alpha \, \Sigma^T \right)^T \right). \tag{E.4}$$

The wedge τ between (E.3) and (E.4) is the only place the tax enters the equilibrium system.

Since the tax and rebate cancel in every realized budget, the wealth dynamics of both types have the same expressions as when $\tau = 0$: the only equation that changes is the type-A portfolio condition (E.3). All effects of the tax therefore operate through the re-optimized portfolio θ^A and the induced changes in prices and in the dynamics of α.

A recursive competitive equilibrium is defined as in Section 4, with all objects functions of α alone and the type-A optimality condition replaced by (E.3). At $\tau = 0$ the economy coincides with the no-inflation-risk counterfactual of Section 5. We solve the model with the global finite-difference method of Appendix C.6.[28]

E.3 The Efficiency Gains from the Tax

We compute the efficiency term Ξ^E of Section 5.2 with the perturbation parameter now being τ rather than σ_π, approximating the derivatives by finite differences between the equilibria at $\tau = \varepsilon$ and $\tau = 0$, and decomposing the gains by the calendar date at which they accrue as in Appendix D.3. Table 5 reports the linearized gain from introducing a tax of $\tau = 0.5\%$, counting welfare accruing within the first T years, evaluated at the modal α of the baseline stationary distribution.

The efficiency gain is positive at every horizon, rising from 0.14% of aggregate output over the first fifty years to 0.65% over 200 years: a marginal tax on borrowing is welfare improving, so the competitive equilibrium of the economy without inflation risk borrows too much. Intuitively, unborn type-B agents place extreme weight on low-probability stock market crashes. If borrowing is reduced, the severity of those events is reduced, and this is extremely valuable for type-B agents. On the flip side, type A agents are also made better-off: the tax on borrowing is rebated lump-sum, so the only effect of it is through equilibrium prices. Given that the demand for funds is lower due to the tax, the equilibrium rate falls, and thus type A agents can finance themselves more cheaply. Another way to think about

[28]We use a grid for α, with unknowns $V^i(\alpha)$, $\Sigma(\alpha)$, and $q_b(\alpha)$ (five unknowns per grid point, since Σ has a single column) and verify that the $\tau = 0$ solution reproduces the $\sigma_\pi = 0$ equilibrium of the two-dimensional model on its invariant slice $\pi = \bar{\pi}$.

	Accrual horizon T (years)		
	$T = 50$	$T = 100$	$T = 200$
Total	0.140	0.338	0.649
A	7.082	11.242	15.772
B	−0.038	0.059	0.261
Initially alive	0.068	0.141	0.204
A	6.742	10.168	12.738
B	−0.103	−0.117	−0.118
Future cohorts	0.072	0.198	0.445
A	0.340	1.074	3.034
B	0.066	0.175	0.378

Table 5: Efficiency gains from introducing a borrowing tax of $\tau = 0.5\%$ in the economy without inflation risk, in percent of aggregate output (perpetual-consumption units), counting welfare accruing within the first T years. Evaluated at the modal α of the baseline ($\tau = 0$) stationary distribution.

this is that, essentially, the tax operates as if type A agents collectively had market power: they shift their demand curve for funds down, which lowers their borrowing costs in a way that increases their welfare. This also explains why type B agents initially lose: they now get lower returns for their savings, and they do not value the extra insurance this generates in equilibrium the same way unborn agents do.