

NBER WORKING PAPER SERIES

PRESENT BIAS AND DISCOUNT RATE RISK

Lars A. Lochstoer  
Stig R. H. Lundeby  
Zhaneta K. Tancheva

Working Paper 34453  
<http://www.nber.org/papers/w34453>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
November 2025

We appreciate helpful comments from Harjoat Bhamra, Julien Cujean, Bernard Dumas, Paul Ehling, Anastassia Fedyk, Federico Gavazzoni, Mariana Khapko, Ali Lazrak, Ella Patelli, and participants at the SFS Cavalcade Asia-Pacific, FIRS, EFA, BFWG, BI-SHoF, SAFE (Goethe University), Early Career Women in Finance Pre-WFA, SFA, Oslo Macro, and BCEA conferences, Adam Smith Workshop, BI Uncertainty Workshop, Finance and Accounting Symposium, and seminars at Boston University, BI Norwegian Business School, and Corvinus University. We acknowledge the support of the Centre for Advanced Study (CAS) in Oslo, Norway, which funded and hosted our research project (Asset Pricing with Heterogeneous Investors in Overlapping Generations) during the 2024/25 academic year. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Lars A. Lochstoer, Stig R. H. Lundeby, and Zhaneta K. Tancheva. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Present Bias and Discount Rate Risk

Lars A. Lochstoer, Stig R. H. Lundeby, and Zhaneta K. Tancheva

NBER Working Paper No. 34453

November 2025

JEL No. G12, G41

## **ABSTRACT**

Recent evidence in the psychology literature suggests that individuals' degree of present bias varies with the state of nature and increases under stress. Consistent with this notion, we document, using survey data on individuals' expected and realized consumption growth, that agents indeed tend to overconsume relative to their expectations and that the degree of such overconsumption is higher in bad times. We analyze the asset pricing implications of this bias and show that even if they control a small fraction of wealth, investors with time-varying degree of present bias cause substantial, priced discount rate risks that have first-order impact on the level and time-variation of asset risk premiums. The mechanism is distinct from that of models with time-varying preference parameters and that of models with biased expectations about aggregate outcomes.

Lars A. Lochstoer

University of California, Los Angeles

Finance Department

and NBER

[lars.lochstoer@anderson.ucla.edu](mailto:lars.lochstoer@anderson.ucla.edu)

Zhaneta K. Tancheva

BI Norwegian Business School

[zhaneta.k.tancheva@bi.no](mailto:zhaneta.k.tancheva@bi.no)

Stig R. H. Lundeby

BI Norwegian Business School

[stig.r.lundeby@bi.no](mailto:stig.r.lundeby@bi.no)

An online appendix is available at <http://www.nber.org/data-appendix/w34453>

# 1 Introduction

Present bias, or hyperbolic discounting, is one of the most widely documented and studied biases in economics (e.g., Phelps and Pollack 1968; Laibson 1997). Agents that suffer from this bias value the present too much relative to the future in a time-inconsistent manner. Although present bias alters consumption–savings choices and hence equilibrium discounting, the prevailing view is that the bias is not important for understanding standard asset pricing puzzles (e.g., Luttmer and Mariotti 2003).

We incorporate recent experimental and neuroscientific findings on present bias into a standard asset-pricing framework and show that present bias has first-order effects on the level and variation of asset risk premia. First, we model the degree of present bias and agents’ ability to forecast their future tastes as time-varying. Evidence from neuroscience and psychology shows that stress impairs individuals’ executive functions (planning and self-control), amplifying present-biased behavior (Arnsten 2009; Sapolsky 2017). Consequently, shifts in economic or technological conditions, recessions, and traumatic events can exacerbate cognitive distortions, raising present bias. Second, we assume that agents recognize others’ time inconsistency but not their own, consistent with Fedyk (2025). This misperception generates forecast errors about their own future consumption, which allows us to calibrate the time variation in bias using survey data.

Our first contribution is empirical. We show that individuals’ expectations about their own consumption growth display systematic, state-dependent pessimism consistent with time-varying present bias and time-varying awareness of that bias. To see the link between present bias awareness (often referred to as partial naiveté) and expectational errors, consider the following. A present-biased agent overconsumes today. If naïve about this bias, the agent forecasts that future selves will not overconsume; when the future arrives, overconsumption recurs. Thus, consumption growth forecasts appear pessimistic as expected growth on average is below realized growth. If stress amplifies this bias, agents will exhibit more pessimistic forecasts of their own consumption growth in bad times. A fully sophisticated agent anticipates future overconsumption and therefore does not exhibit systematically pessimistic forecasts.

We test for systematic forecast errors using the New York Fed’s Survey of Consumer Expectations (SCE). The SCE provides roughly ten years of panel data with individual-level expectations and realizations of real consumption growth. These data allow a direct test of predictability in individual forecast errors. Under the rational expectations benchmark forecast errors are orthogonal to information available at the time of the forecast and thus are not predictable.

Individuals’ realized consumption growth exceeds their own forecasts, yielding an average annual forecast error of about 2 percentage points. The bias is concentrated among more bias-susceptible groups, notably those with lower education and income. Using a cross-sectional difference-in-differences design, we show that forecast pessimism intensifies in periods of economic stress for these groups, consistent with evidence on stress-induced present bias. For the most susceptible agents, the average annual error is about 6 percentage points and ranges from 2 to 10 percentage points across stress states; for the least susceptible agents, mean errors are statistically indistinguishable from zero. The pattern is inconsistent with simple extrapolation, which would generate both predictably positive and negative forecast errors, whereas we observe predominantly predictably negative forecast errors.

At first glance, these errors need not imply asset-pricing consequences. The evidence concerns expectations about own consumption, not aggregates. If agents understand others’ present bias, they will not misperceive aggregate outcomes or payoffs.<sup>1</sup> Nevertheless, we show, as a second contribution of the paper, that this bias gives rise to a novel source of discount rate risks that can have first-order asset pricing implications.

We first analyze the implications of this bias in the simplest possible setting: agents have log utility, aggregate consumption is constant, and there is a continuum of agents with time-varying degree of present bias. At time  $t$  an agent assigns probability  $\theta_t$  that her  $t + 1$  self remains present-biased and probability  $1 - \theta_t$  that from  $t + 1$  onwards she becomes time-consistent with exponential discounting.<sup>2</sup> We allow  $\theta_t$  to vary over time and to be common across agents,

---

<sup>1</sup>This is not to say the literature has not found evidence of bias in expectations about aggregate outcomes (see, e.g. Hirshleifer, Li, and Yu 2015; Nagel and Xu 2022), but the effect of naive present-biased agents is on their individual consumption expectations which is the focus of this paper.

<sup>2</sup>This modeling choice follows the literature on partial awareness of the present bias, except that we allow the subjective probability to be time-varying. In particular, the agent displays partial naiveté as in O’Donoghue and Rabin (2001), Eliaz and Spiegler (2006) and Heidehues and Kőszegi (2010). In the setting of Eliaz and Spiegler

capturing systematic movements in awareness. Thus, the effect of an increase in  $\theta_t$  is that all agents believe they are more likely to remain present-biased in the future which raises economy-wide impatience. In reality, all agents remain present-biased forever.

Even though there is no aggregate risk in this economy, a risk premium on the aggregate wealth claim arises. The driver is agents' beliefs about their own future consumption. In reality, all agents consume a constant amount as no agent ever actually switch type. Subjectively, however, an agent does consider the possibility that they will change type. If an agent were to become time-consistent next period, she would choose to consume a fixed share of wealth given the standard log utility preferences; hence, when an unexpected rise in  $\theta_{t+1}$  raises impatience economy-wide, discount rates rise, asset values fall, and wealth declines. Conditional on switching, the agent would then consume less because consumption is proportional to wealth. Thus, the agent subjectively believes that her own consumption growth comoves positively with the return on aggregate wealth, implying a positive premium on the aggregate claim for markets to clear. Moreover, the premium is time-varying as the subjective probability of switching is  $1 - \theta_t$ . For example, if  $\theta_t = 1$ , the agent is sure they will not switch type, and therefore the previous mechanism is not at play so the risk premium is zero. By contrast, conventional hyperbolic discounting with a constant bias (e.g., Luttmer and Mariotti 2003) and standard time-varying discounting with rational agents and constant aggregates (e.g., Albuquerque et al. 2016) deliver risk-neutral pricing in this environment.

The economic mechanism put forth in this paper delivers a novel discount-rate risk. Beyond this baseline model, we analyze (i) a heterogeneous-agent log-utility economy with a rational subset that discounts exponentially, and (ii) an overlapping-generations model with Epstein–Zin preferences and aggregate consumption risk. We also in the online appendix consider a limited-participation setting in which present-biased agents do not trade equities, and different specifications with time-varying present-bias, including a case where agents are sophisticated about their bias. The core intuition survives in all environments. Calibrating the level and dynamics of present bias to the survey evidence, even a small wealth share of present-biased agents generates

---

(2006) agents with  $\theta = 0$  are defined as fully naive, while agents with  $\theta = 1$  are sophisticated.

large, priced discount-rate risk and improves the fit for the equity premium, excess volatility, the positive slope of the real term structure, and return predictability. Similar asset-pricing results obtain when present-biased agents are fully sophisticated, although a model with agents that are sophisticated about their bias cannot account for the survey evidence we document. We conclude that time-varying present bias is a plausible first-order driver of asset risk premiums.

*Related literature.* Research on present-biased preferences and hyperbolic (or quasi-hyperbolic) discounting dates to Strotz (1956) and Laibson (1997). A core prediction is that sophisticated agents demand illiquidity as a commitment device to curb overconsumption. Recent work refines this link between present bias and intertemporal choice: Maxted (2025), building on Harris and Laibson (2013), analyzes consumption and illiquid-asset demand and the welfare implications in a general consumption–savings environment; Maxted, Laibson, and Moll (2025) show that present bias amplifies the effect of monetary policy while slowing its transmission. These studies keep the degree of present bias fixed. Complementary microfoundations can generate state variation in the strength of present bias. Gabaix and Laibson (2022) model imperfect foresight in which delayed consequences are harder to anticipate than immediate ones, producing as-if hyperbolic discounting whose intensity varies across states. Hertzberg (2024) shows that even time-consistent household members can exhibit overconsumption through a dynamic commons problem, with stronger incentives to pool savings when intra-household relative-wealth risk is high. If stress increases imperfect foresight or intra-household risk, these mechanisms rationalize time-variation in present bias. Our survey evidence indicates that forecasting one’s future decisions is harder in bad times.

Our paper also relates to asset-pricing models with present-biased agents. Luttmer and Mariotti (2003) show that constant present bias does not affect risk premiums. Khapko (2023) theoretically studies state-dependence in both present bias and risk aversion in a representative agent setting, where the agent is sophisticated about their bias and does not make forecast errors. Andries, Eisenbach, and Schmalz (2024) analyze time-inconsistency in risk aversion and address the term structure of risk premia. Contract-theoretic models with time inconsistency and partial naiveté examine principal–agent interactions and menu design rather than competitive

asset trading and aggregate price dynamics (e.g., DellaVigna and Malmendier 2004; Eliaz and Spiegler 2006; Heidhues and Kőszegi 2010; Gottlieb and Zhang 2021; Citanna and Siconolfi 2022).

The model is also connected to work on time-varying discount-rate risk. In our framework the driver is not fundamentals (e.g., Bansal and Yaron 2004; Wachter 2013), nor time-varying or heterogeneous effective preferences at the aggregate level (e.g., Campbell and Cochrane 1999; Bhamra and Uppal 2009; Albuquerque et al. 2016), nor biased beliefs or learning about aggregates (e.g., Dumas, Kurshev, and Uppal 2009; Collin-Dufresne, Johannes, and Lochstoer 2016; Nagel and Xu 2022). Instead assets load on agents' time-varying *subjective* consumption risk that arises from shocks to awareness of present bias. Using a general empirical approach, Kozak and Santosh (2020) document a large negative price of discount-rate risk that helps explain risk premia across stocks and bonds. Our mechanism is consistent with this factor and reproduces its negative price of risk.

The rest of the paper is organized as follows. Section 2 presents a log-utility benchmark that isolates the channel. Section 3 describes the survey evidence on time-varying present bias. Section 4 develops the overlapping-generations model with Epstein–Zin preferences and aggregate risk and states equilibrium conditions. Section 5 reports the calibration disciplined by the survey evidence and the asset-pricing implications. Section 6 discusses robustness to alternative model assumptions. Section 7 concludes.

## 2 Model setup and economic channel

Before the full heterogeneous-agent overlapping-generations model with Epstein–Zin utility and aggregate risk, we study two log-utility benchmarks with no aggregate risk: (i) a representative present-biased agent and (ii) a heterogeneous-agent economy. These environments isolate the mechanism. Time variation in the degree of present bias generates discount-rate risk and a time-varying premium on the aggregate consumption (wealth) claim, even though aggregate consumption is constant. This channel is absent under conventional hyperbolic discounting with a constant bias (e.g., Luttmer and Mariotti 2003) and under time-varying discounting with rational beliefs (e.g., Albuquerque, et al. 2016), which both yield risk-neutral pricing in these

benchmark cases.

## 2.1 Agents

### 2.1.1 Time-varying present bias

We consider a continuum of infinitesimal present-biased agents who value current consumption too much relative to future consumption in a time-inconsistent manner (see Phelps and Pollack 1968; Laibson 1997). Recent literature documents prominent features inherent to an individual who displays this bias. First, Fedyk (2025) shows empirically that such an individual is aware of other agents' bias but is naive about the extent of their own. Second, recent research in psychology and neuroscience (see Arnsten 2009; Sapolsky 2017) shows that the level of present bias varies with the state of nature and increases under stress. To model these features we follow Eliaz and Spiegler (2006) and let the agent have the following preferences:

$$U_t = \log(C_t) + \beta \delta \mathbb{E}_t^S[\tilde{U}_{t+1}], \quad \text{where} \quad (1)$$

$$\tilde{U}_{t+1} = \begin{cases} U_{t+1}, & \text{if the agent remains present-biased} \\ \mathbb{E}_{t+1} \sum_{j=0}^{\infty} \beta^j \log(C_{t+1+j}), & \text{if the agent becomes time-consistent.} \end{cases}$$

Every agent believes that they will remain present-biased and impatient next period with probability  $\theta_t$  (using discount factor  $\beta\delta$ , where  $\delta \in (0, 1)$  captures the present bias) or become time consistent and more patient with probability  $1 - \theta_t$  (using discount factor  $\beta$ ). Following the literature on present-bias, the agent believes that if they become time consistent they will remain so forever. Importantly, we model the agents as naive about their bias in the sense that in reality the agent *always* remains present-biased and never becomes time consistent.<sup>3</sup> This expectational error is captured by the subjective expectation  $\mathbb{E}_t^S[\cdot]$  and the fact that  $\theta_t < 1$  for at least some  $t$ . The degree of present bias  $\theta_t$  varies over time, but every period it is identical for all individuals,

---

<sup>3</sup>In the online appendix we consider the case where the agents are sophisticated about their bias.

capturing the systematic changes in the bias.<sup>4</sup> The agents have rational expectations about all other shocks in the economy. With  $\mathbb{E}[\cdot]$  denoting expectations taken under the true probabilities, the preferences can be written:<sup>5</sup>

$$U_t = \log(C_t) + \beta \delta \mathbb{E}_t[\theta_t U_{t+1} + (1 - \theta_t) \sum_{j=0}^{\infty} \beta^j \log(C_{t+1+j})]. \quad (2)$$

To capture the recent experimental evidence documented by Fedyk (2025), we assume that agents are not aware of their own bias, but anticipate the bias of others. Hence, every agent believes that while they individually may become time consistent at  $t + 1$  with probability  $1 - \theta_t$ , all other agents remain present-biased with probability 1. Thus, each agent believes aggregate wealth is unaffected if they turn time consistent, since no one else will. For this reason each agent  $i$  understands the actual wealth they will have,  $W_{i,t+1}$ , as a function of the aggregate state, even though they may mispredict their own consumption  $C_{i,t+1}$ . In other words, the fact that agents are naive about their future type results in expectation errors about their individual consumption growths, but does not lead to expectation errors about aggregate quantities.

Next, we derive implications for asset prices and the stochastic discount factor in this simple economy where there is no aggregate risk. The time variation in the agent's subjective belief about their own future type,  $\theta_t \in [0, 1]$ , is then the only source of uncertainty in the economy.

For analytical and numerical convenience we assume that  $\theta_t$  can take values  $\hat{\theta}$ , where

$$\hat{\theta} = \begin{pmatrix} \hat{\theta}_1 \\ \vdots \\ \hat{\theta}_K \end{pmatrix} \quad (3)$$

---

<sup>4</sup>Eliaz and Spiegler (2006) term this feature partial naivete and assume a constant  $\theta$  that is different for each individual agent. O'Donoghue and Rabin (2001) model partial naivete as an underestimation of the individual's degree of present bias in the future  $\hat{\delta}$ , where  $0 < \delta \leq \hat{\delta} \leq 1$ . In the online appendix we consider a case with time-varying  $\delta_t$ .

<sup>5</sup>Note that, unlike the case where agents are sophisticated about their own bias (e.g., Luttmer and Mariotti, 2001), solving the value function does not involve solving a game between current and future selves.

with transition probability matrix  $\Pi$  and individual transition probabilities given by  $\pi_{k,l} \equiv \mathbb{P}(\theta_{t+1} = \hat{\theta}_l | \theta_t = \hat{\theta}_k)$ . Note that we only consider state-contingent claims for each possible realization of  $\theta_{t+1}$ , which in this case means that state-contingent claims with payoff  $\mathbf{1}_{\hat{\theta}_k}$  (an indicator that takes the value 1 if  $\theta_{t+1} = \hat{\theta}_k$  and 0 otherwise) are available for all possible states  $\hat{\theta}_k$  next period. We do not allow claims with payoffs contingent on whether the agent becomes time consistent<sup>6</sup>.

Given this process for  $\theta_t$ , we can explore the consumption-savings choice of the agent. In particular, the agent's wealth-consumption ratio,  $\phi$ , as a function of the possible values for  $\theta_t$ ,  $\hat{\theta}$  (see appendix A for derivations), is:

$$\phi(\hat{\theta}) = \left( I - \beta \delta D(\hat{\theta}) \Pi \right)^{-1} \left( \mathbf{1} + \beta \delta (\mathbf{1} - \hat{\theta}) \phi^{TC} \right), \quad (4)$$

where  $D(x) = \text{diag}(x)$  and  $\phi^{TC} = (1 - \beta)^{-1}$  is the constant wealth-consumption ratio the agent would have if she became time consistent. Thus, unlike the standard log utility case, the wealth-consumption ratio of a TI agent in our model varies over time with the current value of  $\theta_t$ . A natural assumption is that the  $\theta_t$  process is persistent.<sup>7</sup> Thus, a high value of  $\theta_t$  means the agent expects to be impatient also in the near future, whereas a low  $\theta_t$  means the agent believes it is highly likely they will become time consistent and therefore more patient. Higher patience pushes the wealth-consumption ratio up as agents are more willing to forego consumption today relative to the future (i.e.  $\phi(\theta_t)$  and  $\theta_t$  are negatively correlated). Through market clearing, this pushes the price of the consumption claim up and discount rates down.

It may seem that individuals' time-varying degree of naiveté about their own present-bias is isomorphic to a model with time-varying (exponential) discounting, as in Albuquerque et al. (2016). We show in the following that this is not the case. In particular, the stochastic discount factor reflects the beliefs of the TI agent and therefore is based on a weighted average of the

---

<sup>6</sup>Since each agent correctly perceives that the probability of others becoming time consistent is zero, allowing these claims would lead to arbitrage. Furthermore, if such claims were traded, it is not clear how anyone could verify that the TI agent in question indeed did remain TI. This is because the agent would have strong incentives to pretend to be a TC agent if she had bought claims that would pay off in the event she became TC. Thus, the claims would need to be contracts tailored for each individual to ensure that the individual would always choose to reveal her type truthfully. Although interesting, this is beyond the scope of this paper.

<sup>7</sup>We give technical conditions on  $\Pi$  in appendix A.

consumption growth in case the agent remains TI or becomes TC next period (see appendix A for derivations):

$$\begin{aligned} M_{t+1} &= \beta \delta \left( \theta_t \frac{C_t}{C_{t+1}^{TI}} + (1 - \theta_t) \frac{C_t}{C_{t+1}^{TC}} \right) \\ &= \beta \delta \left( \theta_t + (1 - \theta_t) \frac{\phi^{TC}}{\phi(\theta_{t+1})} \right). \end{aligned} \quad (5)$$

For intuition, consider an infinitesimal agent  $i$ . Today this agent faces the same problem as all other agents and will in equilibrium choose to hold the consumption claim. The agent's wealth next period will therefore be proportional to aggregate wealth. The agent correctly understands that all other agents will remain present-biased next period and that they in equilibrium will choose consumption proportional to aggregate consumption. Thus, if she remains present-biased, which she believes will happen with probability  $\theta_t$ , her consumption growth is the same as that of the other agents and equal to  $\frac{C_{t+1}^{TI}}{C_t} = 1$ .<sup>8</sup> However, if she becomes time consistent, which she (wrongly) believes will happen with probability  $1 - \theta_t$ , her  $t + 1$  consumption will be  $C_{t+1}^{TC} = W_{i,t+1}/\phi^{TC}$ . Her consumption growth can then be written  $\frac{C_{t+1}^{TC}}{C_t} = \phi(\theta_{t+1})/\phi^{TC} < 1$ . Since  $\phi(\theta_{t+1})$  and  $\theta_{t+1}$  are negatively correlated, a positive shock to  $\theta_{t+1}$  is associated with high marginal utility for this agent. On the other hand, the return on aggregate wealth decreases with  $\theta_{t+1}$ :

$$R_{C,t+1} = \frac{\phi(\theta_{t+1})}{\phi(\theta_t) - 1}. \quad (6)$$

In sum, the agent believes that with probability  $1 - \theta_t$  her next period consumption growth will be low and her marginal utility will be high, exactly when the return to aggregate wealth  $R_{C,t+1}$  is low. This risk is reflected in the above pricing kernel via the term  $(1 - \theta_t) \frac{\phi^{TC}}{\phi(\theta_{t+1})}$  as every agent has the same subjective beliefs  $\theta_t$ . Since the return on the consumption claim is negatively correlated with the pricing kernel, a positive aggregate risk premium arises even though there is no aggregate consumption risk.

---

<sup>8</sup>Recall that aggregate consumption is assumed constant in this section, and thus aggregate consumption growth equals 1.

Furthermore, observe that the risk premium will be time-varying. To see this, it is enough to consider two cases:  $\theta_t = 1$  and  $\theta_t = 0$ . In the first case, each agent correctly understands that they will remain present-biased next period and therefore faces no consumption risk. As a consequence, when  $\theta_t = 1$ , the dependence on  $\phi(\theta_{t+1})$  drops out and there is no risk premium. On the other hand, when  $\theta_t = 0$ , each agent is sure they will become time consistent next period and that their consumption growth will be perfectly correlated with the return on the consumption claim, which gives rise to a positive risk premium. More generally, the risk premium will be positive if and only if  $\theta_t < 1$ .

This novel discount rate risk is generated by the time-varying present bias of the agent. When  $\theta$  is constant, as in the conventional hyperbolic discounting case, the wealth-consumption ratio is constant, resulting in risk-neutral pricing and no additional risk due to present bias. Next, we show that the economic channel causing discount rate risk is distinct from the case with time-consistent agents with time-varying degree of exponential discounting.

### 2.1.2 Time-varying discounting

As a benchmark, we consider a model with time-varying, but time-consistent, exponential discounting similar to Albuquerque, et al. (2016). The continuum of agents have preferences:

$$U_t^{TV} = \log(C_t) + \beta_t \mathbb{E}_t[U_{t+1}^{TV}], \quad (7)$$

where we assume that  $\beta_t$  follows a similar process to  $\theta_t$  and can take values  $\hat{\beta} = (\hat{\beta}_1, \dots, \hat{\beta}_K)'$ , with transition probability matrix  $\Pi_\beta$ .

In contrast to the present-biased agents, this agent has rational expectations about future discounting and does not make predictable forecast errors about her own future consumption growth. The wealth-consumption ratio only depends on the current  $\beta_t$  (see online appendix for derivations):

$$\phi^{TV}(\hat{\beta}) = (I - D(\hat{\beta})\Pi_\beta)^{-1} \mathbf{1}. \quad (8)$$

The stochastic discount factor with these preferences is  $M_{t+1} = \beta_t C_t / C_{t+1}$ . With constant

aggregate consumption, the stochastic discount factor is simply  $M_{t+1} = \beta_t$ , which implies risk-neutral pricing. The returns to the consumption claim are not constant, however, as the risk-free rate is time-varying:

$$R_{C,t+1} = \frac{\phi^{TV}(\beta_{t+1})}{\phi^{TV}(\beta_t) - 1}. \quad (9)$$

Nevertheless, this variation in discount rates is not priced in this case. Intuitively, in an economy with a continuum of agents that apply standard time-varying discounting, each agent is rational and understands that their consumption will be the same as that of all other agents. In the absence of aggregate risk, this consumption is risk-free, and therefore the agent does not require a risk premium.

### 2.1.3 Individual consumption growth expectations

As discussed above, an individual present-biased agent  $i$  believes she will become time consistent with probability  $1 - \theta_t$ , but also believes no other agent will change type. In reality, the agent always remains present-biased. This leads to predictable forecast errors in individual consumption growth. In particular, following the earlier discussion, under agent  $i$ 's beliefs their consumption growth forecast can be written as:

$$F_{i,t} \left( \frac{C_{i,t+1}}{C_{i,t}} \right) = \theta_t \times 1 + (1 - \theta_t) \times \mathbb{E}_t [\phi(\theta_{t+1})] / \phi^{TC}. \quad (10)$$

Since her actual consumption growth simply equals 1, the expected consumption growth forecast error is:

$$\mathbb{E}_t \left[ \frac{C_{i,t+1}}{C_{i,t}} - F_{i,t} \left( \frac{C_{i,t+1}}{C_{i,t}} \right) \right] = (1 - \theta_t) (1 - \mathbb{E}_t [\phi(\theta_{t+1})] / \phi^{TC}). \quad (11)$$

In words, since  $\phi(\theta_{t+1}) < \phi^{TC}$  (see appendix A) forecast errors are predictably positive, which implies that agents are too pessimistic about their own future consumption. Based on the expression, we can see that the forecast error is driven by two channels. First, higher  $\theta_t$  lowers forecast errors, keeping  $\mathbb{E}_t [\phi(\theta_{t+1})]$  constant, as the agent becomes more aware of her true future

type. Second, forecast errors are decreasing in  $\mathbb{E}_t[\phi(\theta_{t+1})]$ , keeping  $\theta_t$  constant: since  $\phi(\theta_t)$  is decreasing in  $\theta_t$ ,  $\mathbb{E}_t[\phi(\theta_{t+1})]$  is decreasing in  $\theta_t$  if  $\phi(\theta_t)$  is persistent. The expected forecast error is 0 only if  $\theta_t = 1$ . Thus, the forecast error will initially grow as we lower  $\theta_t$  from 1, but the pattern might be hump-shaped, i.e. depending on the persistence of  $\theta_t$ , the forecast errors might reach their maximum value at  $\theta^* \in (0, 1)$  rather than at  $\theta_t = 0$ .

We emphasize that the forecast errors induced by present bias refer only to the agents' own consumption growth, not for aggregate consumption, as each agent correctly anticipates that other agents will never actually become time consistent. Further, it is immediate that the time-varying discount rate model does not imply such predictable forecast errors as in this model agents have rational expectations about their own future preferences.

## 2.2 Heterogeneous agents economy

We next introduce time-consistent agents with a constant rate of time preference into the economy. This facilitates risk-sharing and endogenous consumption choices that in turn affect the risk-return trade-off. We for now maintain the assumptions of log felicity functions and no aggregate risk.

In particular, let there be a continuum of two types of infinitesimal infinitely-lived agents – time inconsistent (TI) and time consistent (TC). The preferences of the TI and TC agents are:

$$U_t^{TI} = \log(C_t^{TI}) + \beta \delta \mathbb{E}_t[\theta_t U_{t+1}^{TI} + (1 - \theta_t) U_{t+1}^{TC}] \quad (12)$$

$$U_t^{TC} = \log(C_t^{TC}) + \beta \mathbb{E}_t[U_{t+1}^{TC}], \quad (13)$$

where market clearing implies that the two agents' consumption sums up to aggregate consumption ( $C_t^{TI} + C_t^{TC} = C_t$ ).

The aggregate wealth-consumption ratio in this economy is:

$$\phi_t = \left( \frac{s_t}{\phi^{TI}(\theta_t)} + \frac{1 - s_t}{\phi^{TC}} \right)^{-1}, \quad (14)$$

where  $s_t$  is the wealth share of the TI agents,  $\phi^{TI}$  is given in equation (4) and  $\phi^{TC} = (1 - \beta)^{-1}$ .

See appendix A for derivations.

Since there is no aggregate risk, one might have expected that TI and TC agents would trade such that their marginal utilities are identical across all states  $\theta_{t+1}$  next period, implying no risk premium in the economy. This does not happen as each TI agent wrongly believes that with probability  $1 - \theta_t$  they themselves will become TC and prefer the TC portfolio. In particular, the TI agent believes that with probability  $\theta_t$  their marginal utility next period is that of a TI agent and therefore decreasing in  $\theta_{t+1}$ , whereas with a probability  $1 - \theta_t$  it is that of a TC agent and therefore increasing in  $\theta_{t+1}$ . As a result, their actual portfolio positions are between the portfolio position the agent would have taken if she knew she would remain TI next period and the portfolio position of a TC agent.

These dynamics are reflected in the equilibrium stochastic discount factor (see appendix A for derivations):

$$M_{t+1} = \beta \frac{(1 - s_t)\phi^{TI}(\theta_t) + s_t \delta \left( \theta_t + (1 - \theta_t) \frac{\phi^{TC}}{\phi^{TI}(\theta_{t+1})} \right) \phi^{TC}}{s_t \phi^{TC} + (1 - s_t) \phi^{TI}(\theta_t)}, \quad (15)$$

where  $s_t$  is the wealth share of the TI agents. Thus, the stochastic discount factor correlates negatively with the TI agent's future wealth-consumption ratio  $\phi_{t+1}^{TI}$ . From equation (14) we see that the aggregate wealth-consumption ratio is positively related to  $\phi^{TI}(\theta_t)$  which in turn is negatively related to  $\theta_t$ . Hence, the stochastic discount factor and the return on the consumption claim are negatively correlated, which implies a positive market risk premium.

In contrast, if we replace the TI agents in this economy with agents with standard, but time-varying, exponential discounting, the shocks to their utility discount factor is traded away with the TC agents. Under the maintained assumption of no aggregate consumption risk, this results in a locally deterministic stochastic discount factor (risk-neutral pricing) and no risk premiums (see the online appendix for derivations and detailed discussion of this case). Thus, time-varying present bias gives rise to a novel source of discount rate risk relative to existing preference shock channels.

### 3 Survey evidence on own consumption expectations

The present-biased agents in our model make systematic forecast errors when forecasting their own consumption growth, as discussed in the previous section. In particular, they tend to overconsume next period relative to their expectations. In this section, we show that expectations data from a survey of U.S. consumers are consistent with this implication of the model. We estimate both the level and the time-variation of the bias and use these results in the subsequent calibration of our quantitative model.

#### 3.1 Data

We use data from the New York Fed’s Survey of Consumer Expectations, which is a survey of a population-representative rotating panel of 1,300 U.S. household heads, who own, buy or rent their home. The Core module includes monthly records of a total of 21,222 unique individuals over the sample period June 2013 until July 2023. Each consumer is observed for up to 12 consecutive months and there are about 1,000 respondents on average every month. The Spending module of the survey spans the period December 2014 until December 2022, but only takes place three times per year and includes 12,579 unique individuals.

We construct individuals’ realized real consumption growth and their expected real consumption growth to assess if the agents exhibit predictable forecast errors. The expected consumption growth is available in the Core module under the question: “By about what percentage do you expect your total household spending to [increase/decrease] over the next 12 months?”. In order to obtain expected real consumption growth we subtract expected inflation, which is also reported in the survey under the question: “What do you expect the rate of [inflation/deflation] to be over the next 12 months?”. The realized nominal consumption growth is recorded in the Spending module under the combined answers to the questions “How does your current monthly household spending compare with your household’s monthly spending 12 months ago?” and ”In percentage terms, by how much has your current monthly household spending [increased/decreased] compared to 12 months ago?”. We obtain the realized real consumption growth after subtracting the actual inflation rate from FRED.

Table 1 reports the summary statistics of the main variables of interest from the raw, unfiltered New York Fed’s Survey data. In the first two rows we can already see that individuals’ average realized real consumption growth is higher than their expected real consumption growth (1.05% vs.  $-1.12\%$ ), which suggests that agents on average underestimate the amount they will consume next period. This is exactly the kind of expectation error we would expect from an agent that suffers from present bias. However, before conducting our formal statistical tests, we note that both variables have extreme outliers (e.g., maximum values above 9,000%, and minimum values below  $-100\%$ ). We therefore trim the data cross-sectionally at the 1% and 99% levels each time period, where the respective values reported in Table 1 are reasonable. Finally, we limit the sample to individuals who respond to both the Core and the Spending modules to make sure we observe both their expected and realized consumption growth. After filtering the final sample includes 11,928 individuals.

The Core module also contains information about the individuals’ demographic characteristics. Education is recorded as the highest obtained degree ranging from Less than High School (1) to Professional Degree (8). We classify them into three categories, High School, Some College and College education, consistent with the classification of the New York Fed’s Survey, normalized to 1, 0.5 and 0, respectively. To ensure that the individuals have already completed their education at the time of the survey and to clean errors such as the maximum reported age of 891, we exclude individuals below the age of 25 and above or equal the age of 80. Afterwards, we classify individuals in age cohorts below 40, between 40 and 60, and above 60, consistent with the categories provided by the survey. Income ranges from Less than \$10,000 (1) to \$200,000 or more (11), that we classify in three cohorts consistent with the survey categories: Below \$50,000, Between 50,000 and \$100,000, and Above \$100,000, encoded as 1, 0.5 and 0, respectively. The risk tolerance is available in the Core survey under the question “On a scale from 1 to 7, how would you rate your willingness to take risks regarding financial matters?”. We group the individuals in three categories, low, medium and high risk tolerance: 1, 0.5 and 0, respectively. The last variable reported in Table 1 is an indicator of saving in general available in the Spending module under the question “People budget in different ways. Do you (and your family) generally

try to focus more on trying to save regular amounts of money?”. We define it such that it takes the value of 1, if individuals do not save in general, and 0 otherwise.

We also analyze how the individuals’ consumption growth forecast errors vary with stress factors. In particular, we focus on state-level annual changes in unemployment rates available in FRED, as well as state-level economic conditions indicators. The unemployment rate is available on a monthly frequency at the beginning of the month, so we use the lagged value in order to make sure the information is available to the agents at the time of completing the survey. The economic indicators are developed by Baumeister, Leiva-León and Sims (2024) and available online at a weekly frequency. They are broad indices capturing the mobility, labor market, real activity, and financial conditions of households in the 50 U.S. states. The database does not include the District of Columbia. A value of zero indicates growth equal to the national long-run growth, negative values of the indicators correspond to lower than average growth and positive values denote higher than average growth. We multiply the variable by -1 such that positive values indicate lower than average growth and therefore higher degree of stress. We use the two-month lagged observations of the last week of the month in order to make sure the data is available at the time agents complete the survey.

To show that our results are not driven by financially constrained agents, we run our tests in a sub-sample of only financially unconstrained individuals. Questions that are relevant for defining financial constraints are available in the Credit access module of the survey. In addition to our standard data filters, we limit our sample to the ones who participate in that module, which consists of 7,546 unique consumers. We define individuals as financially constrained and drop them from the sample, if they satisfy one of the following criteria. First, they answer “Yes” to the question: “Over the past 12 months, did you max out (borrow up to the limit) on any of your credit cards?”. Second, they respond “I did not think I would get approved” to the question “You just indicated that you did not apply for any new loans or credit cards over the past twelve months, nor did you make any request for an increase in limits, or refinancing. What is the reason for that?”. Third, they respond “I did not think I would get approved” to the question “You just indicated that it is very unlikely that you will apply for any new loans

or credit cards over the past twelve months, nor did you make any request for an increase in limits, or refinancing. What is the reason for that?”. Fourth, they answer “No, my request was rejected” to the question “Was your request for [new loans or credit cards] granted?”. Finally, to show that biased consumers are not more pessimistic about aggregate quantities, we use the answers to the question “What do you think is the percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now?”, available in the Core module of the survey. Higher values correspond to higher degree of unemployment pessimism.

### 3.2 Individual consumption expectations and realizations

In this section we show that the individuals’ expectations about their real consumption growth are on average negatively biased. Less sophisticated individuals exhibit stronger average bias and increasing bias in times of economic stress.

The final sample gives us expected and realized annual real consumption growth for a broad set of individuals over the period from December 2014 to December 2022. To establish departures from rationality, we would ideally run regressions with the forecast error,  $\Delta c_{t+1}^i - E_t^i(\Delta c_{t+1}^i)$ , on the left hand side of the regression and predictors known at time  $t$ ,  $x_t^i$ , on the right hand side. However, survey participants report their realized consumption growth over the last year along with the expectation of their consumption for the next year. This limits the specifications we are able to run to document the bias as we now have  $\Delta c_t^i - E_t^i(\Delta c_{t+1}^i)$  on the left hand side of the regression, which can be predictable based on variables known at time  $t$  also under rational expectations.

To overcome this challenge, we first establish that agents on average are pessimistic about future consumption and that this bias is stronger for individuals that are likely to be less sophisticated. In particular, columns (1) and (2) of Table 2 give the estimate of the average forecast error  $\mu$  from the regression

$$\Delta c_t^i - E_t^i(\Delta c_{t+1}^i) = \mu + \varepsilon_t^i. \quad (16)$$

The estimates are 1.6% and 2%, respectively, where the latter restricts the sample to only include the pre-Covid period. Thus, agents on average, across time and agents, expect their consumption

growth to be substantially lower than the realized growth is. Since we only have 8 years of time series data available, one may worry that there are time trends in the data that the standard errors do not accurately capture. Columns (3) and (4) of the same table show regressions

$$\Delta c_t^i - E_t^i(\Delta c_{t+1}^i) = \mu_t + \mu_s + \beta^\top x_t^i + \varepsilon_t^i, \quad (17)$$

where  $\mu_t$  and  $\mu_s$  refer to time and state fixed effects, respectively, so the identification is cross-sectional, within-state. The  $x_t^i$  are agent  $i$ 's level of education, income, whether they save generally, age, and risk tolerance. The three former demographic variables are statistically significant at the 5% level or lower, while the two latter are not. All variables, except age, are coded into brackets between 0 and 1, as discussed in the previous section, where a higher value arguably indicates less sophisticated individuals. For instance, no high school education is coded as a 1, while a graduate degree is coded as a 0. The regression coefficients on the three significant coefficients are all positive, which indicates that less sophisticated agents are more pessimistic, where the difference in pessimism from the most to the least sophisticated is about 3% per year for education, 2% per year for income and 0.5% per year for save generally.

The timing issue that we highlight above is not an issue for this regression as long as  $Cov(\Delta c_t^i, x_t^i) = Cov(\Delta c_{t+1}^i, x_t^i)$ , where the covariance is taken across time and agents, controlling for time and state fixed effects. If the level of education, for instance, is set years before an agent enters the sample, this condition is likely to be satisfied. Similarly, the “Do you save regularly?” question refers to a longer-run condition of the agent that is unlikely to affect this covariance. The income brackets could be problematic, however, as an agent that had unexpectedly high income last year also might have consumed more than expected, but this would not necessarily affect the forward-looking expectation. In this case, the high income agent would appear pessimistic relative to the low income agent, which is the opposite of what we find. That said, the income brackets are very coarse, making it unlikely that an agent changes brackets so these shifts are in any case a small part of the variation in the data. Nevertheless, some agents likely do change income brackets (we cannot assess this directly with the data at hand as the survey participants only answer the income once), so this coefficient should be interpreted with

caution. In the online appendix, we show that all our results are robust to removing the income variable from the regressions.

Next, we turn to whether this bias is varying with economic conditions. This makes the timing issue in the data more salient as it entails conditioning on variables that affect the conditional distribution at time  $t$  of  $t + 1$  outcomes. To address this, we create cohorts based on the characteristics we have shown are related to the bias – education, income, and saving generally. Importantly, these characteristics are not changing with the conditioning variables. Income and education take values 0, 0.5, and 1, while saving generally takes values 0 and 1. We create cohorts by state for each of these variables individually, as well as for a combination of all three variables where we use the predicted bias from a regression predicting individual pessimism, similar to that in Table 2, but only using the three variables (i.e., not age or risk tolerance). An example of this cohort, which we term the bias cohort, is individuals with no high school education (the education variable equal to 1), income bracket 0.5, and that save generally. Thus, the agents who are most susceptible to the bias are indicated by 1 and the least biased by 0. At each time  $t$ , we compute the average realized and expected consumption growth for each cohort for each state,  $\Delta c_{s,t}^j$  and  $E_t^j(\Delta c_{s,t+1})$ , where  $j$  refers to the cohort.

We then run regressions of the form:

$$\Delta c_{s,t+1}^j - E_t^j(\Delta c_{s,t+1}) = \mu_t + \mu_s + \beta_1 x_t^j + \beta_2 z_{st} + \beta_3 x_t^j z_{st} + \varepsilon_{s,t+1}^j, \quad (18)$$

where  $x_t^j$  is the variable that defines a cohort, e.g. education level, and  $z_{st}$  is a state-level indicator of economic stress. The economic stress variable is either a state-level economic condition indicator (see, Baumeister, Leiva-Leon, and Sims 2021) or state-level change in unemployment. In the regressions, we sign both variables such that high values indicate bad times. As before, the cohort variable is high for the less sophisticated. Note that in this regression the timing of the left hand side variable is the usual definition of a forecast error due to the use of representative cohorts to construct the realized and expected consumption growth.

The main coefficient of interest is  $\beta_3$ , which asks whether the sensitivity of the bias to the cohort variable is changing with economic conditions. A positive  $\beta_3$  implies that less sophisticated

agents get more pessimistic about their future consumption growth in bad times. This is indeed the case for each variable – cohorts associated with proxies for less sophistication are more pessimistic on average ( $\beta_1$ ) and even more so in bad times ( $\beta_3$ ). Thus, in Column (4) of Table 3, we can see that the most biased agents have an average annual forecast error of about 6%, relative to the least biased agents who do not make any mistakes on average, and that error increases to about 10%, if the stress indicators rise by 2 standard deviations. In Table 5, we show summary statistics of the implied forecast errors. The annualized average error is the average prediction based on the significant explanatory variables from the regression model in Table 3, Panel B, Column 4 (based on bias cohorts and the economic condition indicator using all the available data on a weekly frequency). The annualized standard deviation of the forecast errors is estimated as the projected standard deviation of the economic conditions indicator in the regression model. Confidence intervals are shown in brackets below. The persistence of the forecast errors is based on the annualized autocorrelation of the economic conditions indicator. We use these values as an upper bound for the forecast errors we calibrate our model to in the next section (see Column (2), Table 5). Notably, the coefficient on the economic indicators are insignificant, indicating that the bias is not related to aggregate (in this case, state level) economic outcomes in general. In the online appendix, we show that the results are not driven by agents that are financially constrained, the use of trimming in consumption outcomes, the inclusion of the income variable, or the pessimism in aggregate quantities.

## 4 Model with Epstein-Zin preferences

In this section, we consider quantitative implications of the model in a stationary equilibrium where the agents have recursive preferences. We still keep the model ingredients simple, however, to emphasize the economics and facilitate easy benchmarking to existing models.

## 4.1 The Economy

In order to get a stationary equilibrium, we consider a very simple overlapping-generations (OLG) model where all wealth is financial and the wealth of agents who die is equally redistributed among newborn agents<sup>9</sup>. Each time  $t$  a mass of  $\lambda$  agents are born and a randomly chosen proportion  $\lambda$  die, such that the survival probability at each point in time is  $1 - \lambda$ . Thus, agents born at time  $b \leq t$  represent a  $\lambda(1 - \lambda)^{t-b}$  fraction of the population and at each date  $t$  the total population of all agents born between  $-\infty$  and  $t$  sums to 1.

The economy is populated by two types of agents: time-consistent (TC) and time-inconsistent (TI). The TI agents make up a fraction  $\zeta$  of the overall population. Since the wealth of agents who die are equally redistributed among newborn agents, the wealth share of newborn TI and TC agents are  $\lambda\zeta$  and  $\lambda(1 - \zeta)$ , respectively.

Aggregate consumption growth is assumed to be i.i.d. log-normal with growth rate  $\mu_c$  and volatility  $\sigma_c$ :

$$\log \left( \frac{C_{t+1}}{C_t} \right) = \mu_c - \frac{\sigma_c^2}{2} + \sigma_c \varepsilon_{t+1}. \quad (19)$$

Let  $\phi_t \equiv \frac{W_t}{C_t}$  denote the aggregate wealth-consumption ratio at time  $t$ , where wealth is measured cum-consumption. The return on the aggregate wealth portfolio (i.e. the claim that pays aggregate consumption as dividends) is then:

$$R_{C,t+1} = \frac{\phi_{t+1}}{\phi_t - 1} \frac{C_{t+1}}{C_t}. \quad (20)$$

We assume that the financial markets are complete with respect to shocks to  $\theta_t$  and aggregate consumption. In equilibrium, the optimal portfolio of TI agents at time  $t$  is one that pays off a fraction  $g_{TI}(\theta_{t+1}; \theta_t, s_t) \in (0, 1)$  of aggregate wealth  $W_{t+1}$  at time  $t + 1$ , where  $s_t$  denotes the current wealth share of TI agents. The equilibrium dynamics of TI agents wealth share will therefore be:

$$s_{t+1} = (1 - \lambda)g_{TI}(\theta_{t+1}; \theta_t, s_t) + \lambda\zeta. \quad (21)$$

---

<sup>9</sup>The results are very similar in a discrete time version of the Blanchard (1985) and Gârleanu and Panageas (2015) overlapping-generations model.

Finally, we also consider the pricing of a levered claim to aggregate output, which proxies for the equity market in our model. The log dividend growth of this portfolio is:

$$\log\left(\frac{D_{t+1}}{D_t}\right) = \mu_d + \varrho_d \log\left(\frac{C_{t+1}}{C_t}\right) - \frac{\sigma_d^2}{2} + \sigma_d \varepsilon_{d,t+1}, \quad (22)$$

where  $\varrho_d > 1$  represents leverage and  $\varepsilon_{d,t+1} \sim N(0, 1)$  i.i.d. The expected growth rate of dividends is  $\mu_d + \varrho_d \mu_c$  and  $\sigma_d$  denotes the “idiosyncratic” volatility.

## 4.2 Investor preferences

Both types of agents have Epstein-Zin preferences (Epstein and Zin 1989; Weil 1989) with identical elasticity of intertemporal substitution (EIS) parameter  $\psi$ , where  $\rho = 1 - \frac{1}{\psi}$ , risk aversion level  $\gamma$ , where  $\alpha = 1 - \gamma$ , and time discount factor  $\beta$ . An individual  $i$  who is a time-consistent (TC) agent therefore has the following value function:

$$U_{TC,t}(W_{i,t}) = \max \left[ C_{i,t}^\rho + (1 - \lambda) \beta \mathbb{E}_t \left[ U_{TC,t+1}^\alpha(W_{i,t+1}) \right]^\frac{\rho}{\alpha} \right]^\frac{1}{\rho}, \quad (23)$$

where  $\mathbb{E}_t$  denotes the rational expectation conditional on all available information at time  $t$ , and  $C_{i,t}$  denotes the TC agent  $i$ ’s consumption. It is also useful to denote the wealth of this agent as  $W_{i,t}$ .

Next, the value function of an individual TI agent  $i$  is:

$$U_{TI,t}(W_{i,t}) = \max \left[ C_{i,t}^\rho + (1 - \lambda) \beta \delta \mathbb{E}_t \left[ \theta_t U_{TI,t+1}^\alpha(W_{i,t+1}) + (1 - \theta_t) U_{TC,t+1}^\alpha(W_{i,t+1}) \right]^\frac{\rho}{\alpha} \right]^\frac{1}{\rho}, \quad (24)$$

where  $C_{i,t}$  denotes the TI agent  $i$ ’s consumption. The agent’s wealth is denoted by  $W_{i,t}$ .

As in section 2, TI agents anticipate the bias of others, but not of themselves. Hence, each (representative) TI agent believes that while she will be time consistent at  $t + 1$  with probability  $1 - \theta_t$ , the rest of the agents of her type will remain present-biased with probability 1. Since we rule out claims whose payoffs are contingent on whether an individual becomes TC or not, this again means that each TI agent is rational about the dynamics of asset returns and thus their

wealth,  $W_{i,t+1}$ .

### 4.3 Equilibrium

We find the equilibrium consumption and portfolio choice according to the TI and TC agents' beliefs as of time  $t$ .<sup>10</sup> It is convenient to denote their wealth-consumption ratios by  $\phi_{TC,t}$  and  $\phi_{TI,t}$ , respectively.

Since all agents of the same type are identical except for the level of initial wealth, the equilibrium wealth-consumption ratio of an individual agent is equal to the wealth-consumption ratio of her type. As usual with Epstein-Zin preferences, we can express the agents' value functions at any period  $t$  using a general recursive formulation of their wealth-consumption ratios. We provide a proof of the proposition in appendix B.

**Proposition 1.** *The maximized value functions of an agent  $i$  who is either time-consistent or time-inconsistent with time-varying degree of present bias at any period  $t$  are given by:*

$$U_{TI,t}(W_{i,t}) = \Psi_{TI,t}W_{i,t} \quad (25)$$

$$U_{TC,t}(W_{i,t}) = \Psi_{TC,t}W_{i,t}, \quad (26)$$

where  $\Psi_{TI,t} = \phi_{TI,t}^{\frac{1-\rho}{\rho}}$  and  $\Psi_{TC,t} = \phi_{TC,t}^{\frac{1-\rho}{\rho}}$  and the wealth-consumption ratios take the following form:

$$\phi_{TI,t} \equiv 1 + [(1 - \lambda)\beta\delta]^{\frac{1}{1-\rho}} \mathbb{E}_t \left[ \left( \theta_t \phi_{TI,t+1}^{\frac{\alpha(1-\rho)}{\rho}} + (1 - \theta_t) \phi_{TC,t+1}^{\frac{\alpha(1-\rho)}{\rho}} \right) R_{TI,t+1}^\alpha \right]^{\frac{\rho}{\alpha(1-\rho)}} \quad (27)$$

$$\phi_{TC,t} \equiv 1 + [(1 - \lambda)\beta]^{\frac{1}{1-\rho}} \mathbb{E}_t \left[ \phi_{TC,t+1}^{\frac{\alpha(1-\rho)}{\rho}} R_{TC,t+1}^\alpha \right]^{\frac{\rho}{\alpha(1-\rho)}}, \quad (28)$$

where  $R_{TI,t+1}$  and  $R_{TC,t+1}$  denote the return on the TI and TC agents' portfolios, respectively.

As in the log utility case in Section 2, we only consider time  $t$  state-contingent claims for each possible realization of  $(\theta_{t+1}, \varepsilon_{t+1})$ , which we refer to as a semi-complete market. Since the market is complete with respect to these shocks, the agents' intertemporal marginal rates of substitution

---

<sup>10</sup>O'Donoghue and Rabin (1999, 2001) term this a perception-perfect equilibrium.

are equalized state-by-state with respect to these shocks. In the below proposition, we give the equilibrium stochastic discount factor of each type of agent projected onto  $(\theta_{t+1}, \varepsilon_{t+1})$ . For the TI agent, this is effectively a projection of their marginal intertemporal rates of substitution onto these shocks conditional on the current state vector  $(\theta_t, s_t)$ , where the projection integrates out the (perceived) uncertainty related to switching type from TI to TC.

**Proposition 2.** *The stochastic discount factors (projected onto the state-space generated by  $\theta_{t+1}$  and  $\varepsilon_{t+1}$ ) of the time-consistent agent and time-inconsistent agent with time-varying degree of present bias in semi-complete markets are equal state by state and given by:*

$$M_{TI,t+1}(\theta_{t+1}, \varepsilon_{t+1}; \theta_t, s_t) = M_{TC,t+1}(\theta_{t+1}, \varepsilon_{t+1}; \theta_t, s_t), \quad (29)$$

where

$$M_{TI,t+1} = [(1 - \lambda)\delta\beta]^{\frac{\alpha}{\rho}} \left[ \theta_t \left( \frac{\phi_{TI,t+1}}{\phi_{TI,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} + (1 - \theta_t) \left( \frac{\phi_{TC,t+1}}{\phi_{TC,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} \right] R_{TI,t+1}^{\alpha-1} \quad (30)$$

$$M_{TC,t+1} = [(1 - \lambda)\beta]^{\frac{\alpha}{\rho}} \left( \frac{\phi_{TC,t+1}}{\phi_{TC,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} R_{TC,t+1}^{\alpha-1}. \quad (31)$$

The term  $(1 - \theta_t) \left( \frac{\phi_{TC,t+1}}{\phi_{TI,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}}$  reflects risk perceived by the TI agent associated with a change of type to TC.

## 5 Numerical results: Epstein-Zin preferences

In this section we report numerical results for our heterogeneous agents general equilibrium model with Epstein-Zin utility, disciplined to match the consumption growth forecast errors documented in the New York Fed's survey of consumer expectations. Based on this calibration we show that the novel discount rate risk channel induced by time variation in the present bias has first-order asset pricing implications.

## 5.1 Calibration and consumption growth forecast errors

We choose standard preference, aggregate consumption growth, and OLG parameter values, given in Table 4. Aggregate consumption and dividend growth are i.i.d.. The volatility of consumption growth is set to 2.7%, while the dividend leverage parameter is set to 3. The risk aversion  $\gamma$  of both types of agents is 10, their EIS parameter  $\psi$  is 2, and their discount factor  $\beta$  is 0.984 per year. We set the annual mortality probability to 0.02%, which gives a life-expectancy of 50 years. The fraction of newborn TI agents is  $\zeta = 0.5$  (as in Halevy, 2015). We solve the model on a grid of  $S = 201$  TI wealth share values  $s_t \in (0, 1)$ .

The novel part of our paper relates to the dynamics of  $\theta_t$  and its calibration. The log utility case in Section 2 describes how TI agents make predictable forecast errors about their own future consumption growth, measured as the difference between their objective and subjective consumption growth expectations. This feature carries over to the Epstein-Zin case. We calibrate the parameters governing the dynamics of  $\theta_t$ , conditional on the other standard parameters given above, such that the model is consistent with the magnitude of agents' consumption growth forecast errors documented in the survey. We set  $\delta = 0.9$  on an annual basis, corresponding to an extra discounting of about 10.5% per year, which is on the conservative side based on the time inconsistency literature. We model the subjective belief of the TI agent about their own future type,  $\theta_t$ , as a Markov process bounded between 0 and 1 and consistent with the persistence of the forecast errors in the survey:

$$\theta_t = \left( \frac{e^{x_t}}{1 + e^{x_t}} \right)^\varphi. \quad (32)$$

The shape of the  $\theta_t$  function is determined by  $\varphi$ , while  $x_t$  is a first-order Gaussian autoregressive process with parameters  $\kappa_x = 0.85$  and  $\sigma_x = 5$  governing its persistence and unconditional volatility. At the average value of  $x_t$ ,  $\theta(\mu_x)$  is 0.10, which means that the TI agents are quite naive about their own present bias, consistent with the evidence in Fedyk (2025). To solve the model numerically we assume that  $x_t$  can take  $K = 101$  values  $\hat{x}$  with transition probability matrix  $\Pi_x$  calibrated to match its AR(1) dynamics (e.g., Rouwenhorst 1995) with individual transition probabilities given by  $\pi_{k,l} \equiv \mathbb{P}(x_{t+1} = \hat{x}_l | x_t = \hat{x}_k)$ . This implicitly also defines the transition probabilities for  $\theta$ ,  $\Pi$ .

Table 5 gives the annual mean, standard deviation and autocorrelation of forecast errors both from the New York Fed survey and those implied by our calibrated model. The numbers from the survey correspond to the “Bias” column of Table 3 when the bias variable is set to 1 and the stress variable is the “economic conditions”-variable. The model-implied forecast error is calculated as the TI agents’ expectation of their own consumption growth subtracted from the rational expectation of their consumption growth. In terms of magnitude, the average forecast error of the TI agents in the model is about 6%, similar to that in the survey, which means that TI agents’ forecasts are too low relative to actual consumption. The standard deviation of model-implied forecast errors is about 8%, which is higher than the standard deviation of 2% implied by the survey regressions. However, this is the total variation in forecast errors, while the survey regressions only capture the part projected onto the “economic conditions”-variable. When we do a similar projection within the model, projecting the forecast errors onto measures of the state of the economy – the conditional variance of market returns – the projection has a volatility slightly higher than 1%, which is well within the standard errors in the data. We conclude that the mean and volatility of the forecast errors in the model are consistent with those in the data.

To understand better the interaction between the forecast errors and the degree of present bias, we plot the forecast error as a function of  $\theta$  after integrating out the TI agent wealth share,  $s_t$ . Figure 1 reveals a hump-shaped pattern. On the left hand side of the graph, when  $\theta$  is low, the forecast error is small. This seems at first counter-intuitive, but is happening as a low  $\theta_t$  today means  $\theta_{t+1}$  also is likely to be low given it is a persistent process. In these states, the next period TI and TC wealth-consumption ratios are quite similar, since the agent, if she remains TI, will believe she is likely to become TC very soon given a low  $\theta_{t+1}$ . Thus, next period’s consumption growth is perceived to be similar whether the agent is TI or TC, and thus the forecast error is low. As  $\theta$  increases, there are two opposing effects on the forecast error that drive the hump shape: (1) the TI and TC agents’ wealth-consumption ratios diverge more as the agent perceives that they will likely not change to TC for a while, which means their consumption growth diverge, and (2) the TI agent becomes less wrong about her type next period. Once  $\theta$  reaches 1, the

TI agent correctly understands she will in fact remain TI next month and therefore correctly forecasts one month ahead consumption growth. However, since there is a positive probability that  $\theta$  two months from now will be lower than 1, she thinks there is a positive probability to become TC in the future. As a consequence, the TI agent with current  $\theta = 1$  will still have forecast errors for longer forecast horizons, e.g. annual as in Figure 1.

The figure also gives the unconditional distribution of  $\theta$  in the gray area – most of the time  $\theta$  is low, but there is a long right tail of higher values for  $\theta$  that gives rise to larger forecast errors.

## 5.2 Asset pricing implications

To obtain model predictions, we simulate 20,000 years of monthly data from our model. After discussing the properties of the main calibration, we consider model sensitivity to the share of TI agents in the economy, a constant  $\theta$  case, and the benchmark case where all agents are TC. In the latter case, risk premiums are low and discount rates are constant since consumption growth is i.i.d. This is a natural benchmark against which to contrast the impact of time-varying present-bias.

Table 6 gives asset pricing moments of our various calibrations. Panel A gives the unconditional mean and volatility of the TI wealth share, aggregate consumption-claim return, levered equity return, risk-free rate, and slope of the default-free yield curve. Panel B gives the volatility of conditional risk premiums, as well as the correlation between risk premiums and the price-dividend ratio and the slope of the term structure.

The mean wealth share of the TI agents in our baseline calibration (column (1)) is about 15%, despite half of all agents being TI. This is because their present-bias leads to over-consumption and low savings, which means their wealth is relatively low. Despite the low wealth share their presence still has strong effects on asset prices. For instance, the risk premium on the consumption and dividend claim are about 2.2% and 3.7% per year, respectively. This is compared to 0.7% and 2.2% for the benchmark case when all agents are TC (given in Column (5) of the table). Part of this increase is due to higher stock return volatility. In particular, the benchmark TC case has 2.7% and 9.6% return volatilities for the consumption and dividend claims, respec-

tively, while the baseline calibration yields 7.0% and 11.8%. This excess volatility is due to the time-variation in the present bias, which induces time-varying discount rates. This can be seen in Panel B which shows that the volatilities of the conditional risk premiums are 2.5% and 2.7% for the consumption and dividend claim, respectively. Further, this discount rate risk is priced, as we can see from the increase in the Sharpe ratio of the dividend claim, which is 0.23 when all agents are TC and 0.31 in the baseline calibration. The price of discount rate risk generated in our model is negative, consistent with the empirical evidence in Kozak and Santosh (2020) and as explained in Section 2.

Notably, when the degree of present bias is constant, as shown in columns (3) and (4) of the table, there is no increase in the risk premiums or return volatilities relative to the benchmark case when all agents are TC. This is because, as can be seen from Panel B of Table 6, the constant  $\theta$  cases yield no time-variation in discount rates. This is consistent with Luttmer and Mariotti (2003), who make this point for constant hyperbolic discounting in the case where agents are sophisticated about their present-bias. Columns (3) and (4) show the asset pricing implications of constant  $\theta$ , equal to 0.1 and 0.9, respectively, where the former is the median  $\theta$  in our baseline calibration. The moments that depend strongly on the level of  $\theta$  are the wealth shares of TI agents and the level of the risk-free rate. With  $\theta = 0.1$  the agent believes they with 90% probability become TC next period and forever thereafter. Therefore their consumption and portfolio choice are both close to the TC agent. For this reason, the average wealth share is high and the risk-free rate low relative to the high  $\theta$  case. In the  $\theta = 0.9$  case, the agent over-consumes and under-saves much more strongly, which leads to a lower average wealth share and a higher risk-free rate.

Consistent with stylized facts, the baseline calibration also has an on average upward-sloping term structure, where the average slope of the yield curve is 0.6% with volatility 1.7%. This result is consistent with the findings of Albuquerque, et al. (2016), Duffee (2018), Gomez-Cram and Yaron (2020), and Chernov, Lochstoer, and Song (2025), who show that variation in real rates rather than inflation expectations is the main driver of the positive slope in nominal yields. Again the positive bond risk premium is due to negatively priced discount rate risk. Column

(2) reports asset pricing moments for the time-varying  $\theta_t$  case when there are more TI agents in the economy (80%). In this case, as is intuitive, there is even more variation in discount rates, which leads to a further increase in risk premiums, return volatilities, and Sharpe ratios. Panel B further documents that, consistent with the data, the price-dividend ratio is negatively correlated with conditional expected excess market returns, and the slope of the term structure is positively correlated with conditional expected excess bond returns. Thus, the model is broadly consistent with stylized facts on excess return predictability.

### 5.3 Inspecting the mechanism

To further inspect the mechanism driving discount rate risks, Figure 2 shows the conditional asset pricing moments in our baseline calibration as a function of the time  $t$  degree of present bias  $\theta_t$  after integrating out the TI agent wealth share  $s_t$ . We can see that the equity risk premium, just like the forecast error, is time-varying and increasing in the degree of present bias,  $\theta_t$  (except for when  $\theta_t$  is very close to one). For the most frequently observed values ( $\theta_t < 0.5$ ) the conditional risk premium ranges between 2% and 7% per year. However, in more rare states with higher  $\theta_t$  it can exceed 20%. Note that an important driver of fluctuations in the conditional risk premium is variation in the price of risk (maximum Sharpe ratio), which ranges between 0.3 and 0.8, rather than just variation in the conditional volatility of returns. The bond term premium is hump-shaped in  $\theta_t$ , ranging between 0% and 2% per year.

The increasing discount rates and Sharpe ratios in  $\theta_t$  is due to increasing, priced discount rate risk. To understand this, it is useful to take the perspective of the TI agent. This agent believes they will become TC next period with probability  $1 - \theta_t$ . If we for argument's sake pretend  $\theta_t$  is a constant  $\theta$ , the expected time to becoming TC is  $1/(1 - \theta)$ . The derivative of this "duration" metric is  $1/(1 - \theta)^2$ , which is small when  $\theta$  is close to zero and large as  $\theta$  approaches one. Thus, the expected time to becoming TC is very sensitive to shocks to  $\theta_t$  when  $\theta_t$  is high. As is well-understood, the wealth-consumption ratio of an agent is strongly impacted by the time-discounting, which in turn depends on the expected time to becoming TC. Thus, shocks to  $\theta_t$  have a stronger impact on the TI agent's wealth-consumption ratio, and thus marginal utility,

when  $\theta_t$  is high than when  $\theta_t$  is low.

Next, we calculate impulse-responses from a positive shock to  $\theta_t$  to a selection of endogenous variables.<sup>11</sup> We calculate these impulse response functions for all values of  $\theta_t$  and for three different levels of the TI wealth-share at time  $t$  (30th, 50th, and 70th percentiles) conditioning out the dependence on  $\theta_t$  using its unconditional distribution. Figure 3 plots the resulting “average” impulse response functions. There are two general takeaways worth noting. First, the patterns for all variables are qualitatively similar across initial TI wealth shares. Secondly, higher TI wealth-shares magnify the impulse-responses.

From the top-left panel in Figure 3 we see the impact of the  $\theta_t$ -shock on future TI wealth-shares. At impact, the TI wealth-share goes up significantly, e.g. in the case of initial TI wealth-share at the 70th percentile it increases by about 1 percentage point. This is due to TI agents choosing portfolios that pay off more when there are positive shocks to  $\theta_t$ . However, after the initial impact, the TI wealth share falls each of the following months until it actually drops below its initial starting point. This is due to the TI wealth-consumption ratio on “average” being a decreasing function of  $\theta_t$ . Thus, the positive shock to  $\theta_t$  initially increases TI wealth through a high portfolio return, but then draws down that wealth thereafter due to higher TI consumption.

From the remaining panels in Figure 3 we see that the risk-premium, volatility, Sharpe ratio and the variance risk premium on the dividend claim all jump up at the impact of the shock before gradually reverting. For instance, the immediate impact in the risk-premium and volatility when the initial TI wealth-share is at its median is about 2 and 4 percentage points annualized, respectively. We also see that the log dividend-price ratio and the risk-free rate declines at the impact of the shock before gradually reverting back. Finally, the yield curve slope (the difference in yield-to-maturity on a 10 year bond and the 1 month bond) also jumps up at the impact of the shock.

In sum, time-variation in  $\theta_t$  gives rise to economically significant priced discount rate risks across all asset prices. Next, we discuss how these risks relate to the existing literature on

---

<sup>11</sup>The magnitude of the shock corresponds to a one standard deviation positive shock to  $x_t$ , which is driving  $\theta_t$  per Equation (32).

discount rate risk and the stochastic discount factor.

## 5.4 Present-bias and discount rate risk: decomposing the SDF

To further understand the mechanisms that give rise to our results and to relate it to the existing literature, it is useful to decompose shocks to the log SDF as follows, letting  $\widetilde{z}_{t+1}$  denote  $z_{t+1} - E_t(z_{t+1})$ :

$$\begin{aligned}\widetilde{m}_{t+1} &= -\frac{\gamma-1}{\psi-1} \ln \widetilde{\phi}_{TC,t+1} - \gamma \widetilde{r}_{TC,t+1} \\ &= -\frac{\gamma-1}{\psi-1} \ln \widetilde{\phi}_{TC,t+1} - \gamma \ln(1 - \widetilde{g}_{TI,t+1}) - \gamma \ln \widetilde{\phi}_{t+1} - \gamma \ln \widetilde{C}_{t+1},\end{aligned}\quad (33)$$

where  $\frac{\gamma-1}{\psi-1} > 0$  since  $\gamma > 1$  and  $\psi > 1$  in our calibration. From this decomposition, we see that the marginal utility of the agents at time  $t+1$  is high when either the growth in the TC wealth-consumption ratio is low, or the portfolio return of the TC agent is low. The latter can be low because: 1) the “active portfolio bet” made by TC agents results in a lower TC wealth-share (of TC agents alive at time  $t$ )  $1 - g_{TI,t+1}$  at  $t+1$ ; 2) the aggregate wealth-consumption ratio  $\phi_{t+1}$  is low; or 3) aggregate consumption growth is low. Only the last component, shocks to aggregate consumption growth, remains as a risk factor in the case where all agents are TC. This is also the cash flow risk of the consumption claim. The other components are thus, from the aggregate perspective, drivers of discount rate risk.

Campbell, Giglio, Polk and Turley (2018; CGPT) argue that shocks to stock market cash flows, expected return and return volatility are all priced and thus correlate with the stochastic discount factor. In Table 7 we project the log SDF onto shocks to expected returns, return volatility, and cash flows of either the consumption or dividend claim (columns (1) or (2), respectively). Consistent with CGPT, the sign in the projections on cash flows are negative, while the signs on discount rates are positive. Thus, higher discount rates are associated with bad states, while high cash flows are associated with good states. The fraction of variation explained by each component is given in parenthesis under the projection coefficients. We see that while cash flows are the dominant component, discount rate and volatility risk are important drivers of risk

in the economy. These projections, however, do not capture all variation in the SDF. The  $R^2$ 's of the projections are 75% and 58%, respectively. This is because the model is highly nonlinear with strong conditional dynamics. Thus, the one-period ahead shocks we project linearly onto are not sufficient to capture these elements. Column (3) shows a projection of the log SDF onto the aggregate consumption shock and discount rate components of the SDF – that is the shock to aggregate consumption versus the other components of the SDF. This projection gives a 100%  $R^2$  by construction and reveals that about 57% of the variation in the SDF are due to cash flow shocks while 43% are due to the components that give rise to discount rate shocks in the model.

## 5.5 Discount-rate risk and portfolio choice

Given that discount rates are time-varying and priced in our model, a relevant question becomes who bears this risk and why. Figure 4 plots the unconditional correlations between log returns on the TI and TC portfolios, as well as on the aggregate consumption claim, with shocks to consumption growth and one period expected return and return variance of the consumption claim return. In the final panel, we also plot the correlation with shocks to the “discount rate component” of the SDF, as defined above.

The TI agent's portfolio return is positively correlated with shocks to discount rates (top left and bottom right panels of Figure 4) and shocks to the return variance of the consumption claim, (bottom left panel), while both the TC portfolio return and the consumption claim return are negatively correlated with these shocks. The reason is that discount rates and variances tend to spike when the TI agents become more impatient. However, when TI agents are impatient, their value function is also low for a given level of consumption, which is a high marginal utility state of the world for the TI agents. The TI agents seek to hedge this risk, and therefore buy a portfolio that pays off more when discount rates increase. TC agents on the other hand, face better investment opportunities when discount rates are higher, resulting in a larger value function. The TC agents are therefore willing to accommodate TI agents by taking on more discount rate risk than just passively holding the aggregate consumption claim.

The sharing of risk varies with  $\theta_t$ . When  $\theta_t$  is close to 0, the TI agents are almost sure

they will be TC agents next period, and therefore wish to have similar risk exposures in their portfolio. As a consequence, TI and TC agents hold very similar portfolios, which is very close to the aggregate consumption claim. However, since  $\theta_t$  is the subjective probability of remaining TI, and therefore having the marginal utility of a TI rather than TC agent next period, a higher  $\theta_t$  results in diverging portfolio choices. As a consequence, when TI agents are more impatient, they are also less willing to take on discount rate risk, and thus their portfolio return will be more positively correlated with innovations to discount rates.

## 6 Robustness to alternative model assumptions

In the online appendix we discuss the effects of alternative modeling choices. We consider (1) time-varying present bias through a time-varying  $\delta_t$  instead of using our current time-varying  $\theta_t$  specification, (2) the case where TI investors are precluded from investing in stocks (limited market participation), and (3) the case where TI agents are sophisticated about their bias. The asset pricing implications are qualitatively the same — time-varying present-bias remains a novel source of substantial, priced discount rate risks.

## 7 Conclusion

Present bias is well documented. The prevailing view in benchmark models is that it does not affect risk premiums (Luttmer and Mariotti 2003). Recent empirical and experimental findings emphasize that present-bias is time-varying and stronger during times of stress. Incorporating such time-varying present-bias generates priced discount-rate risk that shifts both conditional and unconditional asset-pricing moments in otherwise standard environments. Assets load on this risk, altering Sharpe ratios and the dynamics of expected returns. This mechanism aligns with evidence that discount-rate risk is central to the risk–return trade-off and asset-price dynamics.

We discipline the model by calibrating the bias to survey evidence on individual consumption-growth forecast errors. These data show that a subset of agents exhibits state-dependent pessimism about own future consumption, consistent with our mechanism. Calibrated to this ev-

idence, quantitative effects are large even when such agents control a small wealth share. We show theoretically that time-varying present bias is distinct from time-varying but exponential discounting and time-varying risk aversion. Unlike these alternatives, time-varying present bias generates priced discount-rate risk even in the absence of aggregate risks.

The paper offers an alternative mechanism for the large swings in discount rates observed in financial markets (e.g., Cochrane 2011), grounded in evidence from neuroscience and psychology and in household consumption–expectation data. An interesting question for future research is to what extent investors indeed show substantial heterogeneity in holdings of discount rate sensitive assets, as predicted by the model.

## References

- [1] Albuquerque, R., M. Eichenbaum, V. Luo, and S. Rebelo. 2016. Valuation Risk and Asset Pricing. *Journal of Finance* 71, 2861-2904.
- [2] Andries, M., T. M. Eisenbach, and M. C. Schmalz. 2024. Horizon-Dependent Risk Aversion and the Timing and Pricing of Uncertainty. *Review of Financial Studies* 37, 3272-3334.
- [3] Arnsten, A. F. T. 2009. Stress Signaling Pathways that Impair Prefrontal Cortex Structure and Function. *Nature Reviews Neuroscience* 10, 410-422.
- [4] Baumeister, C., D. Leiva-León, and E. Sims. 2024. Tracking Weekly State-Level Economic Conditions. *Review of Economics and Statistics* 106, 483-504.
- [5] 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.
- [6] Bhamra, H. S., and R. Uppal. 2009. The Effect of Introducing a Non-redundant Derivative on the Volatility of Stock-market Returns when Agents Differ in Risk Aversion. *Review of Financial Studies* 22, 2303-2330.
- [7] Blanchard, O. J. 1985. Debt, Deficits, and Finite Horizons. *Journal of Political Economy* 93, 223-247.
- [8] 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.
- [9] Campbell, J. Y., S. Giglio, C. Polk, and R. Turley. 2018. An Intertemporal CAPM with Stochastic Volatility. *Journal of Financial Economics* 128, 207-233.
- [10] Citanna, A., and P. Siconolfi. 2022. Naive Stochastic Present Bias and Credit. *Working Paper*.
- [11] Chernov, M., L. A. Lochstoer, and D. Song. 2025. The Real Channel for Nominal Bond-Stock Puzzle. *Journal of Finance*, forthcoming.

[12] Cochrane, J. H. 2011. Presidential Address: Discount Rates. *Journal of Finance* 66, 1047-1108.

[13] Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer. 2016. Parameter Learning in General Equilibrium: The Asset Pricing Implications. *American Economic Review* 106, 664–98.

[14] DellaVigna, S., and U. Malmendier. 2004. Contract Design and Self-Control: Theory and Evidence. *Quarterly Journal of Economics* 119, 353–402.

[15] Duffee, G. R. 2018. Expected Inflation and Other Determinants of Treasury Yields. *Journal of Finance* 73, 2139-2180.

[16] Dumas, B., A. Kurshev, and R. Uppal. 2009. Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility. *Journal of Finance* 64, 579-629.

[17] Eliaz, K., and R. Spiegler. 2006. Contracting with Diversely Naive Agents. *Review of Economic Studies* 73, 689-714.

[18] Epstein, L., and S. Zin. 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica* 57, 937-969.

[19] Fedyk, A. 2025. Asymmetric Naiveté: Beliefs about Self-Control. *Management Science* 71, 6047–6068.

[20] Gabaix, X., and D. Laibson. 2022. Myopia and Discounting. *NBER working paper*.

[21] Gârleanu, N., and S. Panageas. 2015. Young, Old, Conservative, and Bold: The Implication of Finite Lives and Heterogeneity for Asset Pricing. *Journal of Political Economy* 123, 670-695.

[22] Gomez-Cram, R., and A. Yaron. 2020. How Important Are Inflation Expectations for the Nominal Yield Curve? *Review of Financial Studies* 34, 985-1045.

[23] Gottlieb, D., and X. Zhang. 2021. Long-Term Contracting with Time-Inconsistent Agents. *Econometrica* 123, 793-824.

[24] Groneck, M., A. Ludwig, and A. Zimper. 2024. Who Saves More, the Naive or the Sophisticated Agent? *Journal of Economic Theory* 219, 1-33.

[25] Halevy, Y. 2015. Time Inconsistency: Stationarity and Time Invariance. *Econometrica* 83, 335-352.

[26] Harris, C., and D. Laibson. 2013. Instantaneous Gratification. *Quarterly Journal of Economics* 128, 205-248.

[27] Heidehues, P., and B. Kőszegi. 2010. Exploiting Naivete About Self-Control in the Credit Market. *American Economic Review* 100, 2279-2303.

[28] Hertzberg, A. 2024. Time-Consistent Individuals, Time-Inconsistent Households. *Journal of Finance* 79, 3821-3857.

[29] Hirshleifer, D., J. Li, and J. Yu, 2015. Asset Pricing in Production Economies with Extrapolative Expectations. *Journal of Monetary Economics* 76, 87-106.

[30] Khapko, M. 2023. Asset Pricing with Dynamically Inconsistent Agents. *Finance and Stochastics* 27, 1017-1046.

[31] Kozak, S., and S. Santosh. 2020. Why Do Discount Rates Vary? *Journal of Financial Economics* 137, 740-751.

[32] Laibson, D. 1997. Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics* 112, 443-477.

[33] Luttmer, E.G.J., and T. Mariotti. 2003. Subjective Discounting in an Exchange Economy. *Journal of Political Economy* 111, 959-989.

[34] Maxted, P. 2025. Present Bias Unconstrained: Consumption, Welfare, and the Present-Bias Dilemma. *Quarterly Journal of Economics* 140, 2963-3013.

[35] Maxted, P., D. Laibson, and B. Moll. 2025. Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy. *Quarterly Journal of Economics* 140, 691-743.

[36] Nagel, S., and Z. Xu. 2022. Asset Pricing with Fading Memory. *Review of Financial Studies* 35, 2190-2245.

[37] O'Donoghue, T., and M. Rabin. 1999. Doing Now or Later. *American Economic Review* 89, 103-124.

[38] O'Donoghue, T., and M. Rabin. 2001. Choice and Procrastination. *Quarterly Journal of Economics* 116, 121-160.

[39] Phelps, E. S., and R. A. Pollak. 1968. On Second-best National Saving and Game-equilibrium Growth. *Review of Financial Studies* 35, 185-199.

[40] Pollak, R. A. 1968. Consistent Planning. *Review of Economic Studies* 35, 201-208.

[41] Rouwenhorst, K. G. 1995. Asset Pricing Implications of Equilibrium Business Cycle Models. *Frontiers of Business Cycle Research*. Princeton University Press, Princeton.

[42] Sapolsky, R. 2017. *Behave: The Biology of Humans at Our Best and Worst*. London: Penguin Press.

[43] Strotz, R. H. 1956. Myopia and Inconsistency in Dynamic Utility Maximization. *Review of Economic Studies* 23, 165-180.

[44] Wachter, J. 2013. Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility? *Journal of Finance* 67, 987-1035.

[45] Weil, P. 1989. The Equity Premium Puzzle and the Risk-Free Rate Puzzle. *Journal of Monetary Economics* 24, 401-421.

**Table 1: Summary statistics**

The table reports the summary statistics of the New York Fed's Survey of Consumer Expectations raw, unfiltered data. Detailed explanations of the variables are provided in Section 3.1.

|                    | <b>Mean</b><br>(1) | <b>St. dev.</b><br>(2) | <b>N</b><br>(3) | <b>Min</b><br>(4) | <b>P1</b><br>(5) | <b>P99</b><br>(6) | <b>Max</b><br>(7) |
|--------------------|--------------------|------------------------|-----------------|-------------------|------------------|-------------------|-------------------|
| $\Delta c$         | 1.05               | 63.60                  | 27,819          | -477.77           | -41.85           | 42.86             | 9,997.86          |
| $F[\Delta c]$      | -1.12              | 90.11                  | 156,756         | -11,190           | -65              | 50                | 24,965            |
| Education level    | 0.28               | 0.34                   | 157,041         | 0                 | 0                | 1                 | 1                 |
| Income level       | 0.53               | 0.40                   | 155,920         | 0                 | 0                | 1                 | 1                 |
| Age                | 50.54              | 15.54                  | 157,399         | 0                 | 23               | 82                | 891               |
| Risk tolerance     | 0.39               | 0.30                   | 121,145         | 0                 | 0                | 1                 | 1                 |
| Not save generally | 0.66               | 0.47                   | 24,604          | 0                 | 0                | 1                 | 1                 |

**Table 2: Consumption growth forecast errors and demographics**

The table reports the difference between individuals' realized and expected consumption growth based on the New York Fed's Survey of Consumer Expectations data. Column (1) presents the average consumption growth forecast error based on the full sample period (December 2014 until December 2022), while Column (2) focuses on the Pre-Covid (pre-2020) period. Columns (3) and (4) report the estimates of a regression with time and state fixed effects of consumption growth forecast errors on demographic characteristics: education, income, age, risk tolerance (each classified in three normalized categories 0, 0.5 and 1, consistent with the New York Fed's categorization), and an indicator of not saving in general. Individuals with the lowest level of education and income are assigned the value of 1. The realized and expected consumption growth are trimmed at 1% by time. Individuals between 25 and 80 years old are included in the final sample of 11,928 unique individuals. The t-statistics, reported in brackets below, are based on Driscoll-Kraay standard errors using 5 lags. Significance at 10%, 5% and 1% is denoted by \*, \*\*, and \*\*\*, respectively.

|                    | (1)                 | (2)                 | (3)                 | (4)                 |
|--------------------|---------------------|---------------------|---------------------|---------------------|
| Education          |                     |                     | 3.461***<br>(3.725) | 3.367***<br>(3.516) |
| Income             |                     |                     | 1.727***<br>(4.591) | 1.732***<br>(5.288) |
| Not save generally |                     |                     | 0.357***<br>(4.156) | 0.341***<br>(2.911) |
| Age                |                     |                     |                     | 0.293<br>(0.972)    |
| Risk tolerance     |                     |                     |                     | -0.529*<br>(-1.904) |
| Constant           | 1.596***<br>(5.000) | 2.025***<br>(6.065) |                     |                     |
| Time and State FE  | N                   | N                   | Y                   | Y                   |
| Driscoll-Kraay SE  | Y                   | Y                   | Y                   | Y                   |
| Pre-Covid period   | N                   | Y                   | N                   | N                   |
| R-squared          | 0.000               | 0.000               | 0.020               | 0.020               |
| N                  | 25,817              | 16,494              | 23,398              | 22,529              |

**Table 3: Time variation in consumption growth forecast errors**

The table reports the estimates of a regression with time and state fixed effects of average consumption growth forecast errors on demographic characteristic cohorts, a stress indicator, and the interaction between cohort and stress. In Panel A the stress variable  $z_{j,t-1}$  is a state-level economic conditions indicator (as in Baumeister, Leiva-León, and Sims, 2021) and in Panel B – a state-level change in unemployment compared to the year before (available in FRED). Columns (1), (2), and (3) report the estimates using demographic characteristics education, income, and indicator of not saving in general, respectively. Column (4) presents the results based on a bias composed from the previous three indicators and estimated the online appendix. The expected and realized consumption growth along with demographic characteristics are available in the New York Fed's Survey of Consumer Expectations data for the sample period from December 2014 until December 2022. The realized and expected consumption growth are trimmed at 1% by time and individuals between 25 and 80 years old are included in the final sample of 11,928 unique individuals. The t-statistics, reported in brackets below, are based on clustered by cohort and state standard errors. Significance at 10%, 5% and 1% is denoted by \*, \*\*, and \*\*\*, respectively.

|   | Education<br>(1)    | Income<br>(2)       | Save<br>(3)         | Bias<br>(4)          |
|---|---------------------|---------------------|---------------------|----------------------|
| <i>Panel A: State-level change in unemployment</i>        |                     |                     |                     |                      |
| Cohort  | 3.663***<br>(9.203) | 2.252***<br>(7.162) | 1.245***<br>(4.608) | 6.250***<br>(10.809) |
| Stress indicator  | -0.342<br>(-1.278)  | -0.191<br>(-0.972)  | -0.245<br>(-1.143)  | -0.348<br>(-1.504)   |
| Cohort × Stress indicator                                 | 0.695**<br>(2.394)  | 0.480**<br>(2.105)  | 0.338**<br>(2.083)  | 0.775**<br>(2.134)   |
| Time and State FE   | Y                   | Y                   | Y                   | Y                    |
| Clustered SE  | Y                   | Y                   | Y                   | Y                    |
| R-Squared   | 0.072               | 0.072               | 0.095               | 0.049                |
| N   | 2,576               | 2,769               | 1,800               | 5,557                |
| <i>Panel B: State-level economic conditions indicator</i> |                     |                     |                     |                      |
| Cohort  | 3.628***<br>(9.035) | 2.191***<br>(6.964) | 1.263***<br>(4.700) | 6.283***<br>(10.795) |
| Stress indicator  | -0.498<br>(-1.367)  | -0.523*<br>(-1.793) | -0.403<br>(-1.258)  | -0.422<br>(-1.360)   |
| Cohort × Stress indicator                                 | 0.928***<br>(2.781) | 0.546***<br>(2.902) | 0.395**<br>(2.497)  | 0.896**<br>(2.482)   |
| Time and State FE   | Y                   | Y                   | Y                   | Y                    |
| Clustered SE  | Y                   | Y                   | Y                   | Y                    |
| R-Squared   | 0.073               | 0.075               | 0.094               | 0.050                |
| N   | 2,552               | 2,750               | 1,778               | 5,549                |

**Table 4: Parameters for numerical solution**

The table reports the calibration parameters used in the simulation based on the model with time-inconsistent (TI) and time-consistent (TC) agents and Epstein-Zin preferences. Parameter values related to the economy, preferences and OLG are organized in columns and presented at an annual frequency, where applicable. The parameters are explained in detail in Section 5.

| Economy         | Value  | Preferences | Value | OLG          | Value |
|-----------------|--------|-------------|-------|--------------|-------|
| $\mu_c$         | 0.018  | $\beta$     | 0.984 | $\lambda$    | 0.02  |
| $\mu_d$         | -0.018 | $\delta$    | 0.90  | $\zeta_{TI}$ | 0.50  |
| $\sigma_c$      | 0.027  | $\gamma$    | 10    |              |       |
| $\varrho_d$     | 3      | $\psi$      | 2     |              |       |
| $\sigma_d$      | 0.05   |             |       |              |       |
| $\theta(\mu_x)$ | 0.10   |             |       |              |       |
| $\sigma_x$      | 5      |             |       |              |       |
| $\#\sigma_x$    | 4.5    |             |       |              |       |
| $\kappa_x$      | 0.85   |             |       |              |       |
| $\varphi$       | 0.25   |             |       |              |       |
| K               | 101    |             |       |              |       |
| S               | 201    |             |       |              |       |

**Table 5: Individual consumption growth forecast errors**

The table reports summary statistics for the individual consumption growth forecast errors. Column “Survey data” presents the estimated forecast errors (with 95% confidence interval in brackets below) based on the regression analysis in Table 3 using the New York Fed’s Survey of consumer expectations data and the economic conditions indicator as a stress variable (sample period from December 2014 until December 2022). Column “Model” shows the forecast errors computed as the difference between the objective and subjective TI consumption growth expectations in the model with TI and TC agents and Epstein-Zin preferences. Column “Model projection” shows the forecast errors when projected onto the conditional variance of market returns. The calibration parameters of the model are given in Table 4.

|          | <b>Survey data</b>   | <b>Model</b> | <b>Model projection</b> |
|----------|----------------------|--------------|-------------------------|
| Mean     | 6.19<br>[5.04, 7.34] | 6.46         | 6.46                    |
| St. dev. | 2.00<br>[0.42, 3.58] | 7.83         | 1.23                    |

**Table 6: Asset pricing moments**

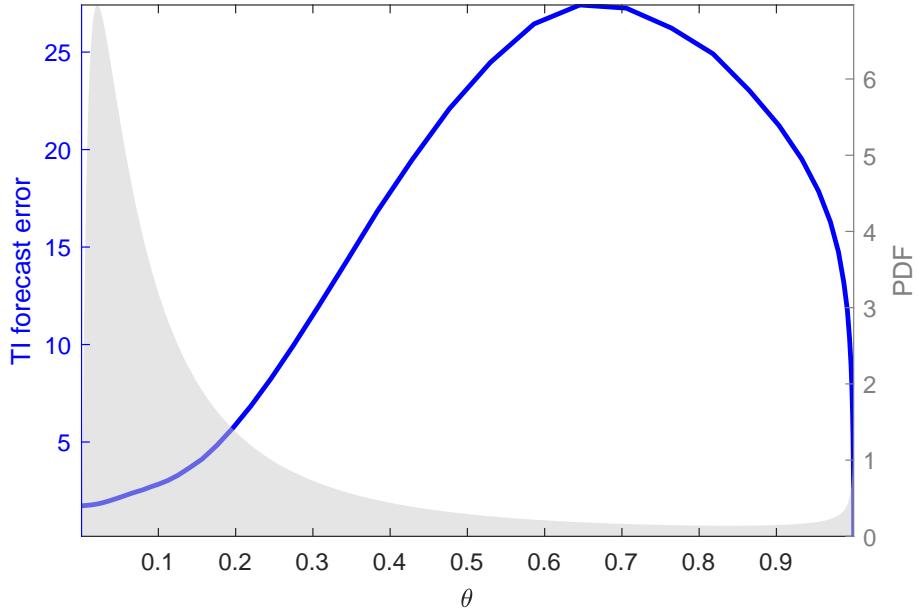
The table reports asset pricing moments for the OLG model with TI and TC agents and Epstein-Zin preferences. The calibration parameters are given in Table 4.  $E(x)$  and  $\sigma(x)$  denote the unconditional mean and variance of  $x$ , respectively, while  $\text{corr}(x, y)$  denotes the correlation between  $x$  and  $y$ .  $SR$  stands for Sharpe ratio,  $s_t$  is the TI agent wealth share,  $R_{C,t}$  is the return to the aggregate consumption claim,  $R_{m,t}$  is the return to the dividend claim,  $R_{f,t}$  is the one-period real risk-free rate,  $R_{10,t}$  is the return to a 10-year default-free, real, zero-coupon bond,  $y_t^{(n)}$  is the yield of the  $n$ -maturity, default-free, real zero-coupon bond, and  $PD_t$  is the price-dividend ratio of the dividend claim. All moments, except for correlations, are annualized.

|  | Baseline | 80% TI<br>agents | Constant<br>$\theta \equiv 0.1$ | Constant<br>$\theta \equiv 0.9$ | All TC |
|--|----------|------------------|---------------------------------|---------------------------------|--------|
|  | (1)      | (2)              | (3)                             | (4)                             | (5)    |
| <b>Panel A: Unconditional moments</b>                              |          |                  |                                 |                                 |        |
| $E(s_t)$   | 15.53    | 17.20            | 49.29                           | 12.39                           | 0.00   |
| $\sigma(s_t)$  | 6.79     | 8.82             | 0.00                            | 0.00                            | 0.00   |
| $E(R_{C,t} - R_{f,t})$   | 2.18     | 3.83             | 0.73                            | 0.73                            | 0.73   |
| $\sigma(R_{C,t} - R_{f,t})$  | 6.96     | 10.75            | 2.71                            | 2.71                            | 2.71   |
| $E(R_{m,t} - R_{f,t})$   | 3.71     | 5.54             | 2.20                            | 2.20                            | 2.20   |
| $\sigma(R_{m,t} - R_{f,t})$  | 11.79    | 14.81            | 9.58                            | 9.58                            | 9.58   |
| $SR(R_{m,t} - R_{f,t})$  | 0.31     | 0.37             | 0.23                            | 0.23                            | 0.23   |
| $E(R_{f,t})$   | 2.44     | 0.78             | 4.01                            | 4.43                            | 3.99   |
| $\sigma(R_{f,t})$  | 1.70     | 3.02             | 0.00                            | 0.00                            | 0.00   |
| $E(y_t^{(10)} - y_t^{(1)})$  | 0.55     | 0.83             | 0.00                            | 0.00                            | 0.00   |
| $\sigma(y_t^{(10)} - y_t^{(1)})$                                   | 1.71     | 3.04             | 0.00                            | 0.00                            | 0.00   |
| <b>Panel B: Conditional moments</b>                                |          |                  |                                 |                                 |        |
| $\sigma(E_t(R_{C,t+1} - R_{f,t+1}))$                               | 2.53     | 4.86             | 0.00                            | 0.00                            | 0.00   |
| $\sigma(E_t(R_{m,t+1} - R_{f,t+1}))$                               | 2.72     | 5.31             | 0.00                            | 0.00                            | 0.00   |
| $\sigma(E_t(R_{t+1}^{(10)} - R_{f,t+1}))$                          | 1.82     | 3.26             | 0.00                            | 0.00                            | 0.00   |
| $\text{Corr}(PD_t, E_t(R_{m,t+1} - R_{f,t+1}))$                    | -0.55    | -0.50            | 0.00                            | 0.00                            | 0.00   |
| $\text{Corr}(y_t^{(10)} - y_t^{(1)}, E_t(R_{10,t+1} - R_{f,t+1}))$ | 0.72     | 0.75             | 0.00                            | 0.00                            | 0.00   |

**Table 7: Price of risk estimates**

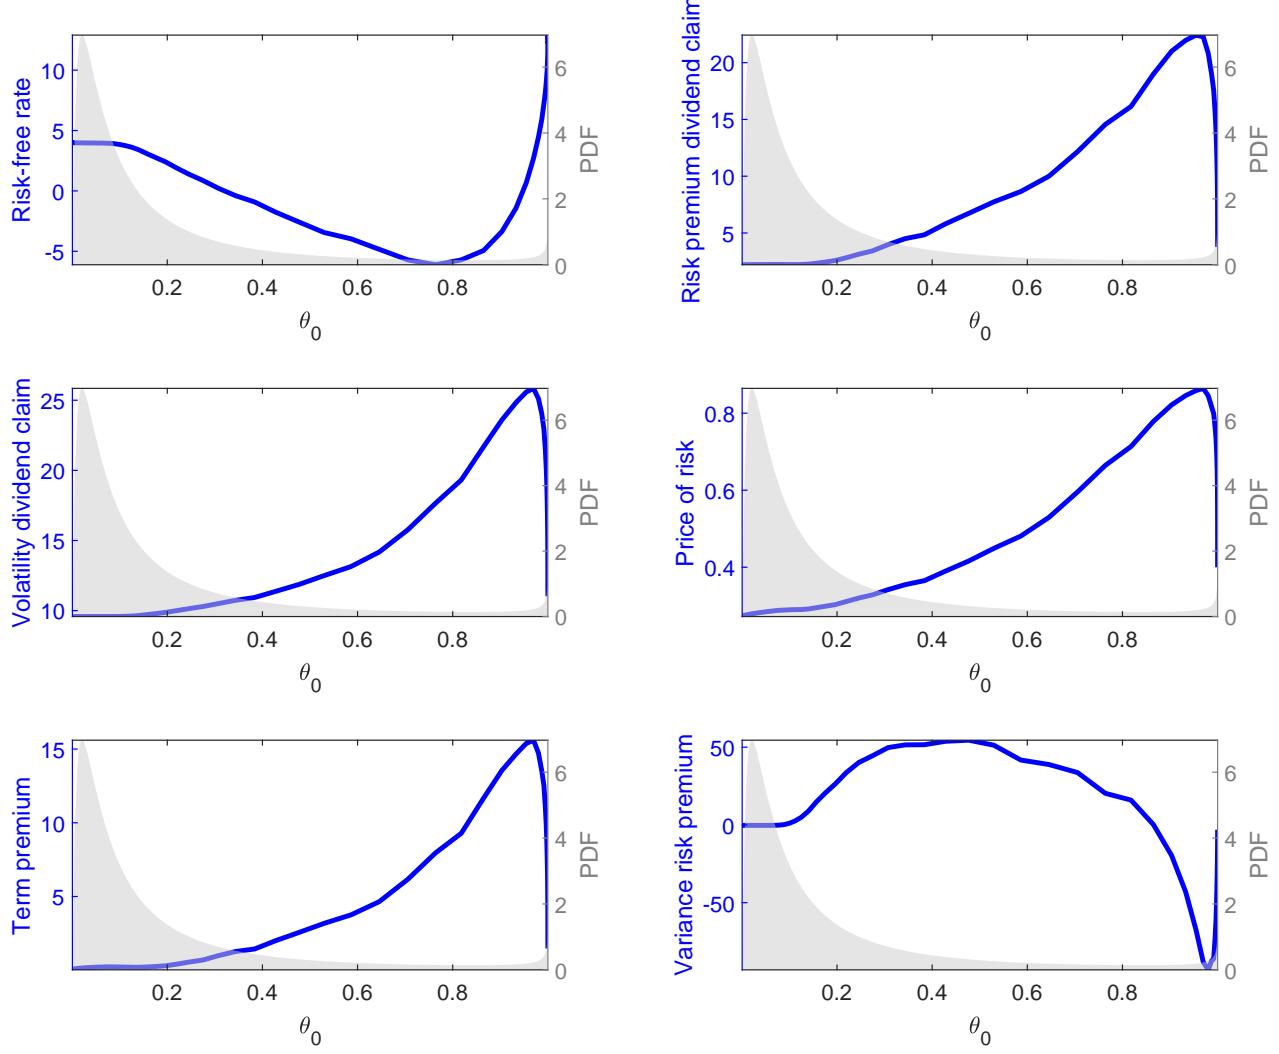
The table reports the coefficients from a projection of innovations to the log SDF onto innovations to expected log returns, cash-flows and variance of log returns. The calibration parameters of the model are given in Table 4. Column (1) and (2) present the projections of the log SDF innovations on the innovations to log cash flow, log expected return and variance based on the consumption and dividend claims, respectively. Column (3), report a decomposition of the log SDF into a discount rate component:  $\frac{(1-\rho)\alpha}{\rho} \log\left(\frac{\phi_{TC,t+1}}{\phi_{TC,t}}\right) + (\alpha - 1) \log\left(\frac{\phi_{t+1}}{\phi_t}\right) + (\alpha - 1) \log\left(\frac{1-g_{TI,t+1}}{1-s_t}\right)$  and log consumption growth. The fraction of SDF variance explained by each variable is reported in parentheses.

|                           | SDF projection    |                  |                   |
|---------------------------|-------------------|------------------|-------------------|
|                           | (1)               | (2)              | (3)               |
| Discount rate innovations | 12.02<br>(12.03)  | 12.45<br>(13.28) | 1.00<br>(42.46)   |
| Variance innovations      | 11.43<br>(3.73)   | 6.46<br>(2.39)   |                   |
| Growth innovations        | -10.00<br>(57.54) | -2.42<br>(41.70) | -10.00<br>(57.54) |
| R-Squared                 | 73.30             | 57.37            | 100.00            |



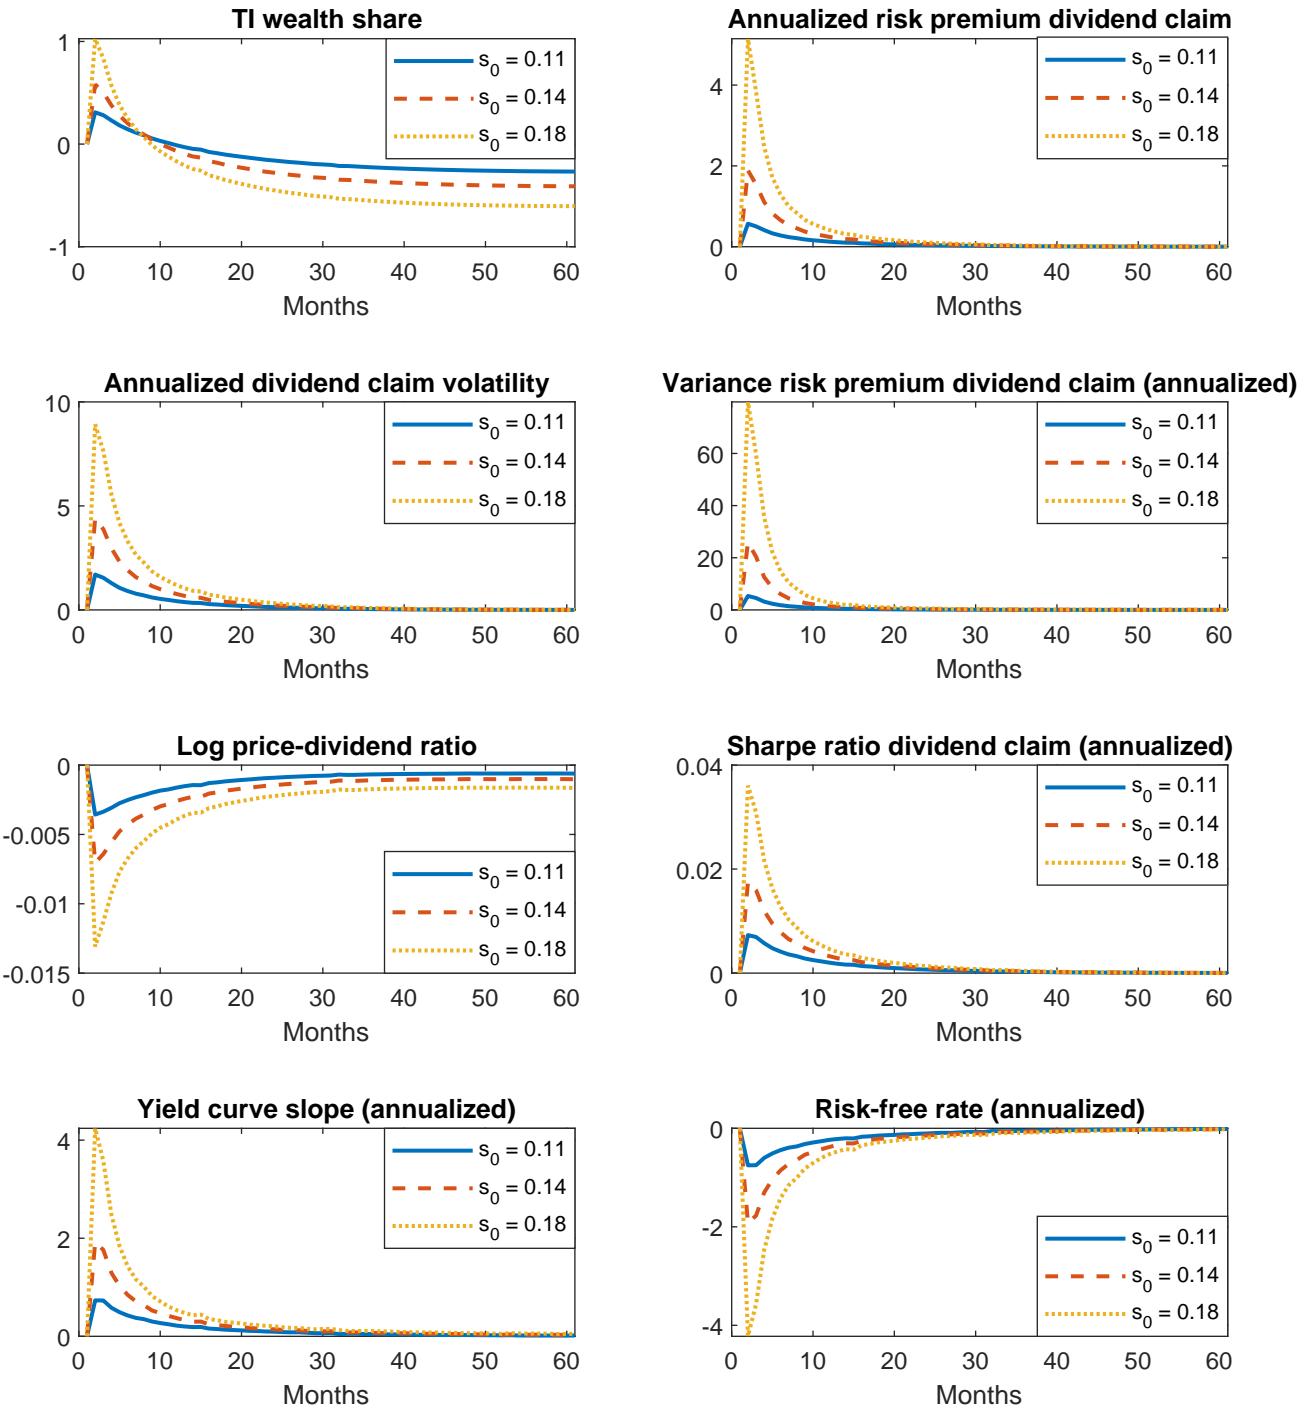
**Figure 1: Consumption growth forecast error**

This figure plots the model-implied annualized forecast error of the time-inconsistent (TI) agent expected consumption growth (left y-axis) against the degree of present bias  $\theta$  (x-axis) after integrating out the TI agent wealth share ( $s_{TI}$ ). The gray shaded area (right y-axis) shows the probability density function of  $\theta$ . The forecast error is calculated as the difference between the TI agent own realized and expected consumption growth. In this setting both the TI and TC agents have Epstein-Zin preferences and the calibration parameters are given in Table 4.



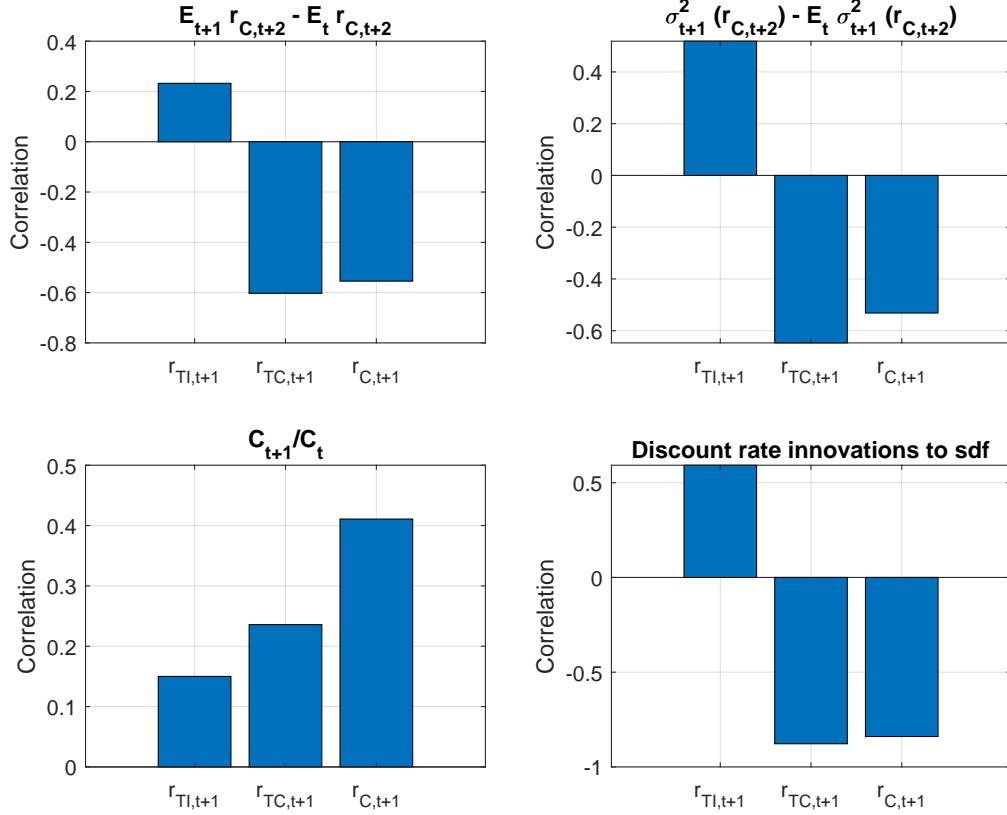
**Figure 2: Asset pricing moments**

This figure plots the annualized risk-free rate, risk premium on the dividend claim, return volatility, price of risk, term premium, and variance risk premium (left y-axis) against the present bias parameter  $\theta$  (x-axis) after integrating out the time-inconsistent agent wealth share  $s_{TI}$ . The gray shaded area (right y-axis) shows the probability density function of  $\theta$ . In this setting both the TI and TC agents have Epstein-Zin preferences and the calibration parameters are given in Table 4.



**Figure 3: Impulse response functions**

This figure plots the impulse response functions of asset pricing moments (y-axis) 0-60 month (x-axis) after a one standard deviation shock ( $\epsilon_x$ ) to  $x_t$ , where the present bias parameter is given by  $\theta_t = \left(\frac{e^{x_t}}{1+e^{x_t}}\right)^\varphi$ . The functions are presented for different initial wealth shares of the time-inconsistent agent  $s_0$ . In this setting both the TI and TC agents have Epstein-Zin preferences and the calibration parameters are given in Table 4.



**Figure 4: Return correlations**

This figure plots the correlations between log returns to the TI and TC portfolios and the consumption claim with: shocks to expected log return and variance of the consumption claim (Panels A and B), with shocks to log consumption growth (Panel C), and shocks to the discount rate component of the log SDF given by  $\frac{(1-\rho)\alpha}{\rho} \log \frac{\phi_{TC,t+1}}{\phi_{TC,t}} + (\alpha - 1) \log \frac{1 - g_{TI,t+1}}{1 - s_t} + (\alpha - 1) \log \frac{\phi_{t+1}}{\phi_t}$ . In this setting both the TI and TC agents have Epstein-Zin preferences and the calibration parameters are given in Table 4.

## A Model solution: Log utility

### A.1 General solution

#### A.1.1 TI value function

The problem of a TI agent is

$$U_t^{TI}(W_t, \theta_t) = \max_{C_t, \omega_t} \log C_t + \beta \delta \mathbb{E}_t \left[ \theta_t U_{t+1}^{TI}(W_{t+1}, \theta_{t+1}) + (1 - \theta_t) U_{t+1}^{TC}(W_{t+1}) \right], \quad (\text{A.1})$$

subject to  $W_{t+1} = (W_t - C_t)(R_{f,t} + \omega_t^\top R_{t+1}^e)$  where  $R_{t+1}^e$  denotes a vector of excess returns. Let  $R_{TI,t+1} = R_{f,t} + \omega_t^{*\top} R_{t+1}^e$  denote the optimized TI portfolio return.

The TC value function is given by

$$U_t^{TC}(W_t) = \max_{C_t, \omega_t} \log C_t + \beta \mathbb{E}_t \left[ U_{t+1}^{TC}(W_{t+1}) \right], \quad (\text{A.2})$$

subject to  $W_{t+1} = (W_t - C_t)(R_{f,t} + \omega_t^\top R_{t+1}^e)$ . Again, let  $R_{TC,t+1} = R_{f,t} + \omega_t^{*\top} R_{t+1}^e$  denote the optimized TC portfolio return. It is straightforward to show that the TC value function can be expressed as

$$U_t^{TC}(W_t) = \phi^{TC} \log W_t + A_t^{TC}, \quad (\text{A.3})$$

where  $A_t^{TC}$  is not a function of current wealth and  $\phi^{TC} = \frac{1}{1-\beta}$  denotes the optimal wealth-consumption ratio of a TC agent.

Let  $\phi_t^{TI} \equiv \frac{W_t}{C_t^*}$  denote the optimized wealth-consumption ratio of a TI agent. Since  $W_{t+1} = (W_t - C_t)R_{TI,t+1} = W_t \frac{\phi_t^{TI}-1}{\phi_t^{TI}} R_{TI,t+1}$  for an agent who was TI at time  $t$ , we have

$$W_{t+1+j} = W_t \prod_{i=0}^j \frac{\phi_{t+i}^{TI}-1}{\phi_{t+i}^{TI}} R_{TI,t+1+i} \quad (\text{A.4})$$

for an agent who was TI until time  $t + j$ . Iterating (A.1) forward therefore yields

$$\begin{aligned}
U_t^{TI}(W_t, \theta_t) &= \max_{C_t, \omega_t} \log C_t + \beta \delta \mathbb{E}_t \left[ \sum_{j=0}^{\infty} \left( \prod_{i=0}^j \theta_{t+i} \right) (\beta \delta)^j \log C_{t+1+j}^* \right. \\
&\quad \left. + \sum_{j=0}^{\infty} \left( \prod_{i=0}^{j-1} \theta_{t+i} \right) (\beta \delta)^j (1 - \theta_{t+j}) U_{t+1+j}^{TC}(W_{t+1+j}) \right] \\
&= \log C_t^* + \beta \delta \mathbb{E}_t \left[ \sum_{j=0}^{\infty} \left( \prod_{i=0}^j \theta_{t+i} \right) (\beta \delta)^j \log \frac{W_{t+1+j}}{\phi_{t+1+j}^{TI}} \right. \\
&\quad \left. + \sum_{j=0}^{\infty} \left( \prod_{i=0}^{j-1} \theta_{t+i} \right) (\beta \delta)^j (1 - \theta_{t+j}) \left( \phi^{TC} \log W_{t+1+j} + A_{t+1+j}^{TC} \right) \right] \\
&= \log \frac{W_t}{\phi_t^{TI}} + \beta \delta \mathbb{E}_t \left[ \sum_{j=0}^{\infty} \left( \prod_{i=0}^j \theta_{t+i} \right) (\beta \delta)^j \log \frac{W_t \prod_{i=0}^j \frac{\phi_{t+i}^{TI}-1}{\phi_{t+i}^{TI}} R_{TI,t+1+i}}{\phi_{t+1+j}^{TI}} \right. \\
&\quad \left. + \sum_{j=0}^{\infty} \left( \prod_{i=0}^{j-1} \theta_{t+i} \right) (\beta \delta)^j (1 - \theta_{t+j}) \left( \phi^{TC} \log \left( W_t \prod_{i=0}^j \frac{\phi_{t+i}^{TI}-1}{\phi_{t+i}^{TI}} R_{TI,t+1+i} \right) + A_{t+1+j}^{TC} \right) \right] \\
&\equiv \left[ 1 + \beta \delta \mathbb{E}_t \sum_{j=0}^{\infty} (\beta \delta)^j \left( \left( \prod_{i=0}^j \theta_{t+i} \right) + \left( \prod_{i=0}^{j-1} \theta_{t+i} \right) (1 - \theta_{t+j}) \phi^{TC} \right) \right] \log W_t \\
&\quad + A_t^{TI}(\theta_t). \tag{A.5}
\end{aligned}$$

Since  $\theta_t$  is a Markov process with a  $K \times K$  transition probability matrix  $\Pi$ , the conditional expectation  $\mathbb{E}_t \prod_{i=0}^j \theta_{t+i}$  can take  $K$  values. In particular, if  $\hat{\theta}$  denotes the  $K \times 1$  vector of possible states for  $\theta_t$ , then, with a slight abuse of notation,  $\mathbb{E}(\theta_{t+1} | \theta_t = \hat{\theta}) = \Pi \hat{\theta}$  is the  $K \times 1$  vector of conditional expectations. Furthermore,  $\mathbb{E}(\theta_t \theta_{t+1} | \theta_t = \hat{\theta}) = D(\hat{\theta}) \mathbb{E}(\theta_{t+1} | \theta_t = \hat{\theta}) = D(\hat{\theta}) \Pi \hat{\theta}$  and  $\mathbb{E}(\theta_t \theta_{t+1} \theta_{t+2} | \theta_t = \hat{\theta}) = D(\hat{\theta}) \mathbb{E}(\mathbb{E}(\theta_{t+1} \theta_{t+2} | \theta_{t+1} = \hat{\theta}) | \theta_t = \hat{\theta}) = D(\hat{\theta}) \Pi \mathbb{E}(\theta_{t+1} \theta_{t+2}) = D(\hat{\theta}) \Pi D(\hat{\theta}) \Pi \hat{\theta}$ . Similarly,  $\mathbb{E}(\theta_t (1 - \theta_{t+1}) | \theta_t = \hat{\theta}) = D(\hat{\theta}) \mathbb{E}(1 - \theta_{t+1} | \theta_t = \hat{\theta}) = D(\hat{\theta}) \Pi (\mathbf{1} - \hat{\theta})$ . Thus,

$$\mathbb{E} \left( \prod_{i=0}^j \theta_{t+i} \middle| \theta_t = \hat{\theta} \right) = \left( \prod_{i=1}^j D(\hat{\theta}) \Pi \right) \hat{\theta} = \left( D(\hat{\theta}) \Pi \right)^j \hat{\theta} \tag{A.6}$$

$$\mathbb{E} \left( \prod_{i=0}^{j-1} \theta_{t+i} (1 - \theta_{t+j}) \middle| \theta_t = \hat{\theta} \right) = \left( \prod_{i=0}^{j-1} D(\hat{\theta}) \Pi \right) (\mathbf{1} - \hat{\theta}) = \left( D(\hat{\theta}) \Pi \right)^j (\mathbf{1} - \hat{\theta}), \tag{A.7}$$

where for a matrix  $A$ , we mean  $A^j = AA^{j-1}$  and  $A^0 = I$ . It is then clear that

$$\begin{aligned}
K^{TI}(\hat{\theta}) &\equiv \mathbf{1} + \beta\delta\mathbb{E}_t\left[\sum_{j=0}^{\infty}(\beta\delta)^j\left(\left(\prod_{i=0}^j\theta_{t+i}\right) + \left(\prod_{i=0}^{j-1}\theta_{t+i}\right)(1-\theta_{t+j})\phi^{TC}\right)\middle|\theta_t = \hat{\theta}\right] \\
&= \mathbf{1} + \beta\delta\left[\sum_{j=0}^{\infty}(\beta\delta)^j\left(D(\hat{\theta})\Pi\right)^j\left(\hat{\theta} + (\mathbf{1} - \hat{\theta})\phi^{TC}\right)\right] \\
&= \mathbf{1} + \beta\delta\left(I - \beta\delta D(\hat{\theta})\Pi\right)^{-1}\left(\hat{\theta} + (\mathbf{1} - \hat{\theta})\phi^{TC}\right) \\
&= \left(I - \beta\delta D(\hat{\theta})\Pi\right)^{-1}\left[\left(I - \beta\delta D(\hat{\theta})\Pi\right)\mathbf{1} + \beta\delta\left(\hat{\theta} + (\mathbf{1} - \hat{\theta})\phi^{TC}\right)\right] \\
&= \left(I - \beta\delta D(\hat{\theta})\Pi\right)^{-1}\left[\mathbf{1} - \beta\delta D(\hat{\theta})\mathbf{1} + \beta\delta\left(\hat{\theta} + (\mathbf{1} - \hat{\theta})\phi^{TC}\right)\right] \\
&= \left(I - \beta\delta D(\hat{\theta})\Pi\right)^{-1}\left(\mathbf{1} + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC}\right), \tag{A.8}
\end{aligned}$$

where we used  $\Pi\mathbf{1} = \mathbf{1}$  (i.e. transition probabilities must sum to 1) and  $D(\hat{\theta})\mathbf{1} = \hat{\theta}$ . Let  $K^{TI}(\theta_t)$  denote the (scalar) value of  $K^{TI}(\hat{\theta})$  corresponding to the specific row  $k$  s.t.  $\hat{\theta}_k = \theta_t$ . We can then express (A.5) as follows

$$U_t^{TI}(W_t, \theta_t) = K^{TI}(\theta_t) \log W_t + A_t^{TI}(\theta_t). \tag{A.9}$$

### A.1.2 Optimal wealth-consumption ratios and TI SDF

To find the optimal TI wealth-consumption ratio, substitute (A.9) and (A.3) for the continuation value functions in (A.1) to get

$$\begin{aligned}
U_t^{TI}(W_t, \theta_t) &= \max_{C_t, \omega_t} \log C_t + \beta\delta\mathbb{E}_t\left[\theta_t\left(K^{TI}(\theta_{t+1}) \log W_{t+1} + A_{t+1}^{TI}(\hat{\theta})\right) + (1-\theta_t)\left(\phi^{TC} \log W_{t+1} + A_{t+1}^{TC}\right)\right] \\
&= \max_{C_t, \omega_t} \log C_t + \log(W_t - C_t)\beta\delta\mathbb{E}_t\left[\theta_t K^{TI}(\theta_{t+1}) + (1-\theta_t)\phi^{TC}\right] \\
&\quad + \beta\delta\mathbb{E}_t\left[\left(\theta_t K^{TI}(\theta_{t+1}) + (1-\theta_t)\phi^{TC}\right) \log (R_{f,t} + \omega_t^\top R_{t+1}^e)\right. \\
&\quad \left. + \theta_t A_{t+1}^{TI}(\theta_{t+1}) + (1-\theta_t)A_{t+1}^{TC}\right]. \tag{A.10}
\end{aligned}$$

The first-order condition w.r.t. consumption is

$$\phi^{TI}(\theta_t) \equiv \frac{W_t}{C_t^*(W_t, \theta_t)} = 1 + \beta\delta\mathbb{E}_t\left[\theta_t K^{TI}(\theta_{t+1}) + (1-\theta_t)\phi^{TC}\right]. \tag{A.11}$$

Stacking up the conditional wealth-consumption ratios in a vector gives us

$$\begin{aligned}
\phi^{TI}(\hat{\theta}) &= \mathbf{1} + \beta\delta D(\hat{\theta})\Pi K^{TI}(\hat{\theta}) + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC} \\
&= \mathbf{1} + \beta\delta D(\hat{\theta})\Pi \left( I - \beta\delta D(\hat{\theta})\Pi \right)^{-1} \left( \mathbf{1} + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC} \right) + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC} \\
&= \left[ I + \beta\delta D(\hat{\theta})\Pi \left( I - \beta\delta D(\hat{\theta})\Pi \right)^{-1} \right] \left( \mathbf{1} + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC} \right) \\
&= \left[ \left( I - \beta\delta D(\hat{\theta})\Pi \right) + \beta\delta D(\hat{\theta})\Pi \right] \left( I - \beta\delta D(\hat{\theta})\Pi \right)^{-1} \left( \mathbf{1} + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC} \right) \\
&= K^{TI}(\hat{\theta}).
\end{aligned} \tag{A.12}$$

Using (A.12) in (A.9) we get that

$$U_t^{TI}(W_t, \theta_t) = \phi^{TI}(\theta_t) \log W_t + A_t^{TI}(\theta_t). \tag{A.13}$$

The first-order conditions w.r.t. portfolio weights are

$$0 = \mathbb{E}_t \left[ \left( \theta_t \phi^{TI}(\theta_{t+1}) + (1 - \theta_t) \phi^{TC} \right) R_{TI,t+1}^{-1} R_{t+1}^e \right]. \tag{A.14}$$

Combining (A.11) and (A.14), we get that the SDF of a TI agent is

$$M_{t+1}^{TI} = \beta\delta \left( \theta_t \frac{\phi^{TI}(\theta_{t+1})}{\phi^{TI}(\theta_t) - 1} + (1 - \theta_t) \frac{\phi^{TC}}{\phi^{TI}(\theta_t) - 1} \right) R_{TI,t+1}^{-1}. \tag{A.15}$$

### A.1.3 Proof: TI wealth-consumption ratio is lower than TC wealth-consumption ratio

To show this, recall from (A.12) that

$$\phi^{TI}(\hat{\theta}) = \left( I - \beta\delta D(\hat{\theta})\Pi \right)^{-1} \left( \mathbf{1} + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC} \right). \tag{A.16}$$

Consider

$$\left( I - \beta\delta D(\hat{\theta})\Pi \right) \mathbf{1} \phi^{TC} = \left( \mathbf{1} - \beta\delta \hat{\theta} \right) \phi^{TC}, \tag{A.17}$$

and subtract  $\mathbf{1} + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC}$  to obtain

$$\begin{aligned} (\mathbf{1} - \beta\delta\hat{\theta})\phi^{TC} - [\mathbf{1} + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC}] &= (\phi^{TC} - 1 - \beta\delta\phi^{TC})\mathbf{1} = \left(\frac{1}{1-\beta} - 1 - \beta\delta\frac{1}{1-\beta}\right)\mathbf{1} \\ &= \frac{\beta(1-\delta)}{1-\beta}\mathbf{1} > 0. \end{aligned} \quad (\text{A.18})$$

Thus, we have established

$$(\mathbf{I} - \beta\delta D(\hat{\theta})\Pi)\mathbf{1}\phi^{TC} > \mathbf{1} + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC}. \quad (\text{A.19})$$

Note that  $(\mathbf{I} - \beta\delta D(\hat{\theta})\Pi)^{-1} = \sum_{j=0}^{\infty} (\beta\delta D(\hat{\theta})\Pi)^j$ . Since each element of  $\beta\delta D(\hat{\theta})\Pi$  is non-negative, each element of  $(\beta\delta D(\hat{\theta})\Pi)^j$  must be non-negative as well. Thus,  $(\mathbf{I} - \beta\delta D(\hat{\theta})\Pi)^{-1} = \sum_{j=0}^{\infty} (\beta\delta D(\hat{\theta})\Pi)^j$  contains only non-negative elements. Therefore, the inequality in (A.19) is preserved by multiplying by  $(\mathbf{I} - \beta\delta D(\hat{\theta})\Pi)^{-1}$  which establishes the result:

$$\begin{aligned} \mathbf{1}\phi^{TC} &= (\mathbf{I} - \beta\delta D(\hat{\theta})\Pi)^{-1}(\mathbf{I} - \beta\delta D(\hat{\theta})\Pi)\mathbf{1}\phi^{TC} \\ &> (\mathbf{I} - \beta\delta D(\hat{\theta})\Pi)^{-1}(\mathbf{1} + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC}) = \phi^{TI}(\hat{\theta}). \end{aligned} \quad (\text{A.20})$$

#### A.1.4 Proof: TI wealth-consumption ratio decreasing in $\theta_t$

Assume:

1.  $\hat{\theta}$  is ordered from smallest to largest, i.e.  $\hat{\theta}_l - \hat{\theta}_k \geq 0$  for all  $l \geq k$ .
2.  $\sum_{i=1}^m (\Pi_{k,i} - \Pi_{l,i}) \geq 0$  for all  $m = 1, \dots, K$  and for all  $l \geq k$ .

Consider the recursion

$$x^{(1)} \equiv \mathbf{1}\phi^{TC} \quad (\text{A.21})$$

$$x^{(n+1)} \equiv \mathbf{1} + \beta\delta \left[ D(\hat{\theta})\Pi x^{(n)} + (\mathbf{1} - \hat{\theta})\phi^{TC} \right] = (1 + \beta\delta\phi^{TC})\mathbf{1} - \beta\delta D(\hat{\theta})(\mathbf{1}\phi^{TC} - \Pi x^{(n)}). \quad (\text{A.22})$$

We have that

$$\lim_{n \rightarrow \infty} x^{(n)} = \sum_{j=0}^{\infty} (\beta\delta D(\hat{\theta})\Pi)^j [\mathbf{1} + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC}] = (\mathbf{I} - \beta\delta D(\hat{\theta})\Pi)^{-1} [\mathbf{1} + \beta\delta(\mathbf{1} - \hat{\theta})\phi^{TC}] = \phi^{TI}(\hat{\theta}). \quad (\text{A.23})$$

Furthermore, for all  $n$ ,  $\mathbf{0} < x^{(n+1)} \leq \mathbf{1}\phi^{TC}$ . This is easily seen in the case of  $n = 2$  as  $x^{(2)} = (1 + \beta\delta\phi^{TC})\mathbf{1} - \beta\delta D(\hat{\theta})(\mathbf{1}\phi^{TC} - \Pi\mathbf{1}\phi^{TC}) = (1 + \beta\delta\phi^{TC})\mathbf{1} < (1 + \beta\phi^{TC})\mathbf{1} = \phi^{TC}\mathbf{1}$ . For the general pattern, suppose  $\mathbf{0} < x^{(n)} < \mathbf{1}\phi^{TC}$ . Then, since  $\Pi x^{(n)}$  is just a weighted average of the elements of  $x^{(n)}$  where  $x^{(n)} < \mathbf{1}\phi^{TC}$ , we have  $\Pi x^{(n)} \leq \mathbf{1}\phi^{TC}$ , thus  $D(\hat{\theta})\Pi x^{(n)} + (\mathbf{1} - \hat{\theta})\phi^{TC} \leq \mathbf{1}\phi^{TC}$  which in turn implies  $x^{(n+1)} = \mathbf{1} + \beta\delta[D(\hat{\theta})\Pi x^{(n)} + (\mathbf{1} - \hat{\theta})\phi^{TC}] \leq \mathbf{1} + \beta\delta\mathbf{1}\phi^{TC} < \mathbf{1} + \beta\mathbf{1}\phi^{TC} = \mathbf{1}\phi^{TC}$ . The positivity of the sequence  $x^{(n)}$  is trivial.

As we have already established, the sequence  $x^{(n)}$  converges to  $\phi^{TI}(\hat{\theta})$  as  $n \rightarrow \infty$ . Since  $x_i^{(1)}$  is non-increasing in  $i$ , i.e.  $x_k^{(1)} - x_l^{(1)} \geq 0$  for all  $l \geq k$ , our proof consists of establishing that if the elements  $x_i^{(n)}$  are such that  $x_k^{(n)} - x_l^{(n)} \geq 0$  for all  $l \geq k$ , then that implies  $x_k^{(n+1)} - x_l^{(n+1)} \geq 0$  for all  $l \geq k$ .

Let  $l \geq k$ . From (A.22) we have that

$$x_k^{(n+1)} - x_l^{(n+1)} = \beta\delta\left[\hat{\theta}_l\left(\phi^{TC} - \sum_{i=1}^K \Pi_{l,i}x_i^{(n)}\right) - \hat{\theta}_k\left(\phi^{TC} - \sum_{i=1}^K \Pi_{k,i}x_i^{(n)}\right)\right]. \quad (\text{A.24})$$

Note that

$$\begin{aligned} \phi^{TC} - \sum_{i=1}^K \Pi_{l,i}x_i^{(n)} &\geq \phi^{TC} - \sum_{i=1}^K \Pi_{k,i}x_i^{(n)} \Leftrightarrow -\sum_{i=1}^K \Pi_{l,i}x_i^{(n)} \geq -\sum_{i=1}^K \Pi_{k,i}x_i^{(n)} \Leftrightarrow \\ &\sum_{i=1}^K (\Pi_{k,i} - \Pi_{l,i})x_i^{(n)} \geq 0. \end{aligned} \quad (\text{A.25})$$

Clearly, since  $0 < x_i^{(n)} \leq \phi^{TC}$  and  $x_i^{(n)}$  is non-increasing in  $i$ , (A.25) is satisfied by assumption 2 above. Thus

$$\begin{aligned} x_k^{(n+1)} - x_l^{(n+1)} &= \beta\delta\left[\hat{\theta}_l\left(\phi^{TC} - \sum_{i=1}^K \Pi_{l,i}x_i^{(n)}\right) - \hat{\theta}_k\left(\phi^{TC} - \sum_{i=1}^K \Pi_{k,i}x_i^{(n)}\right)\right] \\ &\geq \beta\delta(\hat{\theta}_l - \hat{\theta}_k)\left(\phi^{TC} - \sum_{i=1}^K \Pi_{k,i}x_i^{(n)}\right) \geq 0, \end{aligned} \quad (\text{A.26})$$

where we used Assumption 1:  $\hat{\theta}_l - \hat{\theta}_k \geq 0$ . To establish our result, take the limit of (A.26) to

obtain

$$\begin{aligned} \lim_{n \rightarrow \infty} x_k^{(n+1)} - x_l^{(n+1)} &= \phi^{TI}(\hat{\theta}_k) - \phi^{TI}(\hat{\theta}_l) = \beta \delta \left[ \hat{\theta}_l \left( \phi^{TC} - \sum_{i=1}^K \Pi_{l,i} \phi^{TI}(\hat{\theta}_i) \right) - \hat{\theta}_k \left( \phi^{TC} - \sum_{i=1}^K \Pi_{k,i} \phi^{TI}(\hat{\theta}_i) \right) \right] \\ &\geq \beta \delta (\hat{\theta}_l - \hat{\theta}_k) \left( \phi^{TC} - \sum_{i=1}^K \Pi_{k,i} \phi^{TI}(\hat{\theta}_i) \right) \geq 0. \end{aligned} \quad (\text{A.27})$$

## A.2 Continuum of infinitesimal (“representative”) TI agents

Suppose the economy is populated by a unit-mass continuum of identical and infinitesimal TI agents. Each agent faces the optimization problem (A.1). Furthermore, we assume that each agent is fully aware of the bias of others, i.e. he understands they will never become TC.

Since all agents are identical, individual and aggregate wealth-consumption ratios are the same, i.e.  $\phi^{TI}(\hat{\theta}) = \phi(\hat{\theta})$ , where  $\phi(\hat{\theta})$  denotes the aggregate wealth-consumption ratio. Furthermore, all agents must find it optimal to hold the same portfolio in equilibrium, which must therefore be the same as the aggregate consumption claim,  $R_{TI,t+1} = R_{C,t+1}$ . Finally, all agents must have identical consumption growth in equilibrium,  $\frac{C_{TI,t+1}}{C_{TI,t}} = \frac{C_{t+1}}{C_t}$  where the left-hand side denotes the optimized TI consumption and the right-hand side denotes the aggregate consumption endowment. We have

$$R_{TI,t+1} = R_{C,t+1} = \frac{\phi(\theta_{t+1})}{\phi(\theta_{t+1}) - 1} \frac{C_{t+1}}{C_t}. \quad (\text{A.28})$$

Since all agents are identical, all individual SDFs will be identical as well. Thus, let  $M_{t+1} = M_{t+1}^{TI}$  denote the SDF. Substituting in (A.28) into (A.15) gives us

$$\begin{aligned} M_{t+1} &= \beta \delta \left( \theta_t \frac{\phi(\theta_{t+1})}{\phi(\theta_t) - 1} + (1 - \theta_t) \frac{\phi_{TC}}{\phi(\theta_t) - 1} \right) \left( \frac{\phi(\theta_{t+1})}{\phi(\theta_{t+1}) - 1} \frac{C_{t+1}}{C_t} \right)^{-1} \\ &= \beta \delta \left( \theta_t + (1 - \theta_t) \frac{\phi_{TC}}{\phi(\theta_{t+1})} \right) \left( \frac{C_{t+1}}{C_t} \right)^{-1}. \end{aligned} \quad (\text{A.29})$$

In the special case of constant aggregate consumption, this simplifies further to

$$M_{t+1} = \beta \delta \left( \theta_t + (1 - \theta_t) \frac{\phi_{TC}}{\phi(\theta_{t+1})} \right). \quad (\text{A.30})$$

### A.3 Heterogeneous agents

Let  $s_t$  denote the share of aggregate wealth held by TI agents and let the remaining share  $1 - s_t$  be held by TC agents. By market-clearing for consumption, we have that the sum of TI and TC consumption must equal aggregate consumption, i.e.

$$C_t = C_{TI,t} + C_{TC,t} = \frac{W_{TI,t}}{\phi^{TI}(\theta_t)} + \frac{W_{TC,t}}{\phi^{TC}} = \left( \frac{s_t}{\phi^{TI}(\theta_t)} + \frac{1 - s_t}{\phi^{TC}} \right) W_t \Leftrightarrow \\ \phi(s_t, \theta_t) \equiv \frac{W_t}{C_t} = \left( \frac{s_t}{\phi^{TI}(\theta_t)} + \frac{1 - s_t}{\phi^{TC}} \right)^{-1} = \phi^{TI}(\theta_t) \phi^{TC} (s_t \phi^{TC} + (1 - s_t) \phi^{TI}(\theta_t))^{-1}. \quad (\text{A.31})$$

The pricing kernel for a TC agent takes the familiar form:

$$M_{t+1}^{TC} = \frac{1}{R_{TC,t+1}} = \left( \frac{W_{TC,t+1}}{W_{TC,t} - C_{TC,t}} \right)^{-1} = \left( \frac{(1 - s_{t+1}) W_{t+1}}{\frac{\phi^{TC}-1}{\phi^{TC}} (1 - s_t) W_t} \right)^{-1} \\ = \beta \frac{1 - s_t}{1 - s_{t+1}} \frac{W_t}{W_{t+1}} = \beta \frac{C_t}{C_{t+1}} \frac{1 - s_t}{1 - s_{t+1}} \frac{\phi_t}{\phi_{t+1}}, \quad (\text{A.32})$$

where  $R_{t+1}^{TC}$  denotes the optimized portfolio return of the TC agent and  $\frac{\phi^{TC}-1}{\phi^{TC}} = \frac{\frac{1}{1-\beta}-1}{\frac{1}{1-\beta}} = \beta$ .

Substituting for  $R_{TI,t+1}$  in (A.15) gives us the TI pricing kernel

$$M_{t+1}^{TI} = \beta \delta \left( \theta_t \frac{\phi^{TI}(\theta_{t+1})}{\phi_t^{TI} - 1} + (1 - \theta_t) \frac{\phi^{TC}}{\phi^{TI}(\theta_t) - 1} \right) \left( \frac{W_{TI,t+1}}{W_{TI,t} - C_{TI,t}} \right)^{-1} \\ = \beta \delta \left( \theta_t \frac{\phi^{TI}(\theta_{t+1})}{\phi^{TI}(\theta_t) - 1} + (1 - \theta_t) \frac{\phi^{TC}}{\phi^{TI}(\theta_t) - 1} \right) \left( \frac{s_{t+1} W_{t+1}}{\frac{\phi^{TI}(\theta_t)-1}{\phi^{TI}(\theta_t)} s_t W_t} \right)^{-1} \\ = \beta \delta \left( \theta_t \frac{\phi^{TI}(\theta_{t+1})}{\phi^{TI}(\theta_t)} + (1 - \theta_t) \frac{\phi^{TC}}{\phi^{TI}(\theta_t)} \right) \frac{s_t}{s_{t+1}} \frac{W_t}{W_{t+1}} \\ = \beta \delta \frac{C_t}{C_{t+1}} \left( \theta_t \frac{\phi^{TI}(\theta_{t+1})}{\phi^{TI}(\theta_t)} + (1 - \theta_t) \frac{\phi^{TC}}{\phi^{TI}(\theta_t)} \right) \frac{s_t}{s_{t+1}} \frac{\phi_t}{\phi_{t+1}}. \quad (\text{A.33})$$

The evolution of the wealth-share of TI agents is given by

$$s_{t+1} \equiv \frac{W_{TI,t+1}}{W_{t+1}} = \frac{W_{TI,t} - C_{TI,t}}{W_{t+1} - C_t} \frac{R_{TI,t+1}}{R_{C,t+1}} = \frac{\frac{\phi^{TI}(\theta_t)-1}{\phi^{TI}(\theta_t)} s_t W_t}{\frac{\phi(s_t, \theta_t)-1}{\phi(s_t, \theta_t)} W_t} \frac{R_{TI,t+1}}{R_{C,t+1}} \\ = s_t \frac{\phi^{TI}(\theta_t) - 1}{\phi^{TI}(\theta_t)} \frac{\phi(s_t, \theta_t)}{\phi(s_t, \theta_t) - 1} \frac{R_{TI,t+1}}{R_{C,t+1}}. \quad (\text{A.34})$$

### A.3.1 “Semi-complete” markets

With semi-complete markets we allow the agents to trade assets that span the state-space generated by all aggregate shocks such as shocks to  $\theta_{t+1}$  or to  $C_{t+1}$ . However, we do not allow agents to trade claims whose payoffs are contingent on an *individual* agent’s type.

In equilibrium with semi-complete markets, the two SDFs must equal state-by-state on the state-space generated by the aggregate shocks. Using (A.32) and (A.33)

$$\begin{aligned}
M_{t+1}^{TC} &= M_{t+1}^{TI} \Leftrightarrow \beta \frac{C_t}{C_{t+1}} \frac{1 - s_t}{1 - s_{t+1}} \frac{\phi(s_t, \theta_t)}{\phi(s_{t+1}, \theta_{t+1})} \\
&= \beta \delta \frac{C_t}{C_{t+1}} \left( \theta_t \frac{\phi^{TI}(\theta_{t+1})}{\phi^{TI}(\theta_t)} + (1 - \theta_t) \frac{\phi^{TC}}{\phi^{TI}(\theta_t)} \right) \frac{s_t}{s_{t+1}} \frac{\phi(s_t, \theta_t)}{\phi(s_{t+1}, \theta_{t+1})} \Leftrightarrow \\
\frac{1 - s_t}{1 - s_{t+1}} &= \delta \left( \theta_t \frac{\phi^{TI}(\theta_{t+1})}{\phi^{TI}(\theta_t)} + (1 - \theta_t) \frac{\phi^{TC}}{\phi^{TI}(\theta_t)} \right) \frac{s_t}{s_{t+1}} \Leftrightarrow \\
(1 - s_t)s_{t+1} &= \delta \left( \theta_t \frac{\phi^{TI}(\theta_{t+1})}{\phi^{TI}(\theta_t)} + (1 - \theta_t) \frac{\phi^{TC}}{\phi^{TI}(\theta_t)} \right) s_t(1 - s_{t+1}) \\
s_{t+1} &= \frac{\delta s_t (\theta_t \phi^{TI}(\theta_{t+1}) + (1 - \theta_t) \phi^{TC})}{(1 - s_t) \phi^{TI}(\theta_t) + \delta s_t (\theta_t \phi^{TI}(\theta_{t+1}) + (1 - \theta_t) \phi^{TC})}. \tag{A.35}
\end{aligned}$$

The TC agents’ wealth-share next period must therefore be given by

$$1 - s_{t+1} = \frac{(1 - s_t) \phi^{TI}(\theta_t)}{(1 - s_t) \phi^{TI}(\theta_t) + \delta s_t (\theta_t \phi^{TI}(\theta_{t+1}) + (1 - \theta_t) \phi^{TC})}. \tag{A.36}$$

We see immediately that the TC agents’ portfolio will pay off a larger fraction of aggregate wealth when TI agents’ wealth-consumption ratio is low.

Thus, substituting (A.35) and (A.36) into (A.31) at  $t + 1$ , we see that next period aggregate wealth-consumption ratio is given by:

$$\begin{aligned}
\tilde{\phi}(\theta_{t+1}; s_t, \theta_t) &\equiv \phi(s_{t+1}, \theta_{t+1}) = \phi^{TC} \phi^{TI}(\theta_{t+1}) (s_{t+1} \phi^{TC} + (1 - s_{t+1}) \phi^{TI}(\theta_{t+1}))^{-1} \\
&= \phi^{TC} \phi^{TI}(\theta_{t+1}) \left[ \frac{\delta s_t (\theta_t \phi^{TI}(\theta_{t+1}) + (1 - \theta_t) \phi^{TC}) \phi^{TC} + (1 - s_t) \phi^{TI}(\theta_t) \phi^{TI}(\theta_{t+1})}{(1 - s_t) \phi^{TI}(\theta_t) + \delta s_t (\theta_t \phi^{TI}(\theta_{t+1}) + (1 - \theta_t) \phi^{TC})} \right]^{-1} \\
&= \phi^{TC} \frac{(1 - s_t) \phi^{TI}(\theta_t) + s_t \delta (\theta_t \phi^{TI}(\theta_{t+1}) + (1 - \theta_t) \phi^{TC})}{(1 - s_t) \phi^{TI}(\theta_t) + s_t \delta \left( \theta_t + (1 - \theta_t) \frac{\phi^{TC}}{\phi^{TI}(\theta_{t+1})} \right) \phi^{TC}}, \tag{A.37}
\end{aligned}$$

which is clearly increasing in next-period TI wealth-consumption ratio. As a consequence, the realized return on the consumption claim  $\left( R_{C,t+1} = \frac{\tilde{\phi}(\theta_{t+1}; s_t, \theta_t)}{\phi(s_t, \theta_t) - 1} \frac{C_{t+1}}{C_t} \right)$  will be increasing in the TI

wealth-consumption ratio.

Substituting (A.31), (A.35), and (A.37) into (A.33), we get the following expression of the equilibrium SDF:

$$M_{t+1} = \beta \frac{C_t}{C_{t+1}} \frac{(1 - s_t)\phi^{TI}(\theta_t) + s_t \delta \left( \theta_t + (1 - \theta_t) \frac{\phi^{TC}}{\phi^{TI}(\theta_{t+1})} \right) \phi^{TC}}{s_t \phi^{TC} + (1 - s_t) \phi^{TI}(\theta_t)}, \quad (\text{A.38})$$

which under the assumption of constant aggregate consumption simplifies to

$$M_{t+1} = \beta \frac{(1 - s_t)\phi^{TI}(\theta_t) + s_t \delta \left( \theta_t + (1 - \theta_t) \frac{\phi^{TC}}{\phi^{TI}(\theta_{t+1})} \right) \phi^{TC}}{s_t \phi^{TC} + (1 - s_t) \phi^{TI}(\theta_t)}. \quad (\text{A.39})$$

## B Model solution: Epstein-Zin utility

In this section we present the model with Epstein-Zin utility and overlapping generations (OLG). All individual agents  $i$  of a given type TC or TI have identical preferences and face identical risks. In particular, there are no differences between older and younger agents, except the wealth they enter the period with. Therefore, we will solve the problem for some arbitrary individual belonging to the group of TI or TC agents respectively, without explicitly indicating their birth cohort.

The problem of an individual TC or TI agent  $i$  can be written

$$U_{TC,t}(W_{i,t}) = \max_{C_{i,t}, \omega_{i,t}} \left[ C_{i,t}^\rho + (1 - \lambda)\beta \mathbb{E}_t \left[ U_{TC,t+1}(W_{i,t+1})^\alpha \right]^{\frac{\rho}{\alpha}} \right]^{\frac{1}{\rho}} \quad (\text{B.1})$$

$$U_{TI,t}(W_{i,t}) = \max_{C_{i,t}, \omega_{i,t}} \left[ C_{i,t}^\rho + (1 - \lambda)\beta \delta \mathbb{E}_t \left[ \theta_t U_{TI,t+1}(W_{i,t+1})^\alpha + (1 - \theta_t) U_{TC,t+1}(W_{i,t+1})^\alpha \right]^{\frac{\rho}{\alpha}} \right]^{\frac{1}{\rho}} \quad (\text{B.2})$$

subject to the budget constraint  $W_{i,t+1} = (W_{i,t} - C_{i,t})(R_{f,t} + \omega_{i,t}^\top R_{t+1}^e)$  where  $R_{t+1}^e$  denotes the vector of excess returns.

Since all agents of a given type are the same (except for the level of their wealth), the wealth-consumption ratio of any individual agent who is TC will be  $\phi_{TC,t}$  and for any individual agent who is TI it will be  $\phi_{TI,t}$ . Similarly, all agents of the same type will hold the same portfolio, i.e.  $R_{TC,t+1}$  and  $R_{TI,t+1}$ , respectively.

Let us guess the following form of the value functions:

$$U_{TC,t+1}(W_{i,t+1}) = \phi_{TC,t+1}^{\frac{1-\rho}{\rho}} W_{i,t+1} \quad (\text{B.3})$$

$$U_{TI,t+1}(W_{i,t+1}) = \phi_{TI,t+1}^{\frac{1-\rho}{\rho}} W_{i,t+1}. \quad (\text{B.4})$$

Using (B.3) and (B.4) along with the budget constraint, we can write the objective functions as

follows

$$U_{TC,t}(W_{i,t}) = \max \left[ C_{i,t}^\rho + (1 - \lambda)\beta(W_{i,t} - C_{i,t})^\rho \mathbb{E}_t \left[ \phi_{TC,t+1}^{\frac{(1-\rho)\alpha}{\rho}} R_{TC,t+1}^\alpha \right]^{\frac{1}{\alpha}} \right]^{\frac{1}{\rho}} \quad (\text{B.5})$$

$$U_{TI,t}(W_{i,t}) = \max \left[ C_{i,t}^\rho + (1 - \lambda)\beta\delta(W_{i,t} - C_{i,t})^\rho \mathbb{E}_t \left[ \left( \theta_t \phi_{TI,t+1}^{\frac{(1-\rho)\alpha}{\rho}} + (1 - \theta_t) \phi_{TC,t+1}^{\frac{(1-\rho)\alpha}{\rho}} \right) R_{TI,t+1}^\alpha \right]^{\frac{1}{\alpha}} \right]^{\frac{1}{\rho}}. \quad (\text{B.6})$$

We need to show that the value functions at time  $t$  take the same forms as (B.3) and (B.4).

The first-order condition w.r.t. time  $t$  consumption for a TC agent is

$$\begin{aligned} 0 &= C_{i,t}^{\rho-1} - (1 - \lambda)\beta(W_{i,t} - C_{i,t})^{\rho-1} \mathbb{E}_t \left[ \phi_{TC,t+1}^{\frac{(1-\rho)\alpha}{\rho}} R_{TC,t+1}^\alpha \right]^{\frac{1}{\alpha}} \Leftrightarrow \\ \left( \frac{W_{i,t} - C_{i,t}}{C_{i,t}} \right)^{1-\rho} &= (1 - \lambda)\beta \mathbb{E}_t \left[ \phi_{TC,t+1}^{\frac{(1-\rho)\alpha}{\rho}} R_{TC,t+1}^\alpha \right]^{\frac{1}{\alpha}} \Leftrightarrow \\ (\phi_{TC,t} - 1)^{1-\rho} &= (1 - \lambda)\beta \mathbb{E}_t \left[ \phi_{TC,t+1}^{\frac{(1-\rho)\alpha}{\rho}} R_{TC,t+1}^\alpha \right]^{\frac{1}{\alpha}}. \end{aligned} \quad (\text{B.7})$$

Similarly, for a TI agent we get

$$(\phi_{TI,t} - 1)^{1-\rho} = (1 - \lambda)\beta\delta \mathbb{E}_t \left[ \left( \theta_t \phi_{TI,t+1}^{\frac{(1-\rho)\alpha}{\rho}} + (1 - \theta_t) \phi_{TC,t+1}^{\frac{(1-\rho)\alpha}{\rho}} \right) R_{TI,t+1}^\alpha \right]^{\frac{1}{\alpha}}. \quad (\text{B.8})$$

Using the expressions for the wealth-consumption ratios (B.7) and (B.8) in the objective functions (B.5) and (B.6) yields what we want to show:

$$\begin{aligned} U_{TC,t}(W_{i,t}) &= \left[ C_{i,t}^\rho + (W_{i,t} - C_{i,t})^\rho (\phi_{TC,t} - 1)^{1-\rho} \right]^{\frac{1}{\rho}} \\ &= \left[ \left( \frac{W_{i,t}}{\phi_{TC,t}} \right)^\rho + \left( \frac{W_{i,t}}{\phi_{TC,t}} \right)^\rho (\phi_{TC,t} - 1)^\rho (\phi_{TC,t} - 1)^{1-\rho} \right]^{\frac{1}{\rho}} \\ &= [1 + (\phi_{TC,t} - 1)]^{\frac{1}{\rho}} \frac{W_{i,t}}{\phi_{TC,t}} = \phi_{TC,t}^{\frac{1}{\rho}-1} W_{i,t} \\ &= \phi_{TC,t}^{\frac{1-\rho}{\rho}} W_{i,t} \end{aligned} \quad (\text{B.9})$$

$$\begin{aligned} U_{TI,t}(W_{i,t}) &= \left[ C_{i,t}^\rho + (W_{i,t} - C_{i,t})^\rho (\phi_{TI,t} - 1)^{1-\rho} \right]^{\frac{1}{\rho}} \\ &= \phi_{TI,t}^{\frac{1-\rho}{\rho}} W_{i,t}. \end{aligned} \quad (\text{B.10})$$

The first-order conditions w.r.t. portfolio choice are

$$0 = \mathbb{E}_t \left[ \phi_{TC,t+1}^{\frac{(1-\rho)\alpha}{\rho}} R_{TC,t+1}^{\alpha-1} R_{t+1}^e \right] \quad (\text{B.11})$$

$$0 = \mathbb{E}_t \left[ \left( \theta_t \phi_{TI,t+1}^{\frac{(1-\rho)\alpha}{\rho}} + (1 - \theta_t) \phi_{TC,t+1}^{\frac{(1-\rho)\alpha}{\rho}} \right) R_{TI,t+1}^{\alpha-1} R_{t+1}^e \right]. \quad (\text{B.12})$$

Combining the first-order conditions w.r.t. consumption, (B.7) and (B.8), with the first-order conditions w.r.t. portfolio choice, (B.11) and (B.12), we see that

$$M_{TC,t+1} = [(1 - \lambda)\beta]^{\frac{\alpha}{\rho}} \left( \frac{\phi_{TC,t+1}}{\phi_{TC,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} R_{TC,t+1}^{\alpha-1} \quad (\text{B.13})$$

$$M_{TI,t+1} = [(1 - \lambda)\beta\delta]^{\frac{\alpha}{\rho}} \left[ \theta_t \left( \frac{\phi_{TI,t+1}}{\phi_{TI,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} + (1 - \theta_t) \left( \frac{\phi_{TC,t+1}}{\phi_{TI,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} \right] R_{TI,t+1}^{\alpha-1} \quad (\text{B.14})$$

are the pricing kernels for any TC and TI agent respectively.

## B.1 Semi-complete markets

Assume that traded assets span the state-space generated by aggregate shocks, i.e. shocks to aggregate consumption and  $\theta_{t+1}$ . In this case, we have  $M_{TC,t+1} = M_{TI,t+1} = M_{t+1}$  state-by-state in the state-space generated by  $(\theta_{t+1}, \varepsilon_{t+1})$ .<sup>12</sup> It is easy to see that it is sufficient if for any current state  $(s_t, \theta_t)$ , agents can trade contracts with payoffs  $\chi_{\theta_{t+1}=\hat{\theta}_k} W_{t+1}$ , where  $\chi_{\theta_{t+1}=\hat{\theta}_k}$  is an indicator function that takes the value 1 if  $\theta_{t+1} = \hat{\theta}_k$  and 0 otherwise.

Let  $g_i(\hat{\theta}_k; s_t, \theta_t)$  denote the fraction of the contract with payoff  $\chi_{\theta_{t+1}=\hat{\theta}_k} W_{t+1}$  agent  $i$  buys if current state is  $(s_t, \theta_t)$ . Then, the wealth of agent  $i$  contingent on being alive at  $t + 1$  will be  $W_{i,t+1} = g_i(\theta_{t+1}; s_t, \theta_t) W_{t+1}$ . Let

$$g_{TI}(\theta_{t+1}; s_t, \theta_t) \equiv \int_{i \in TI} g_i(\theta_{t+1}; s_t, \theta_t) di \quad (\text{B.15})$$

$$g_{TC}(\theta_{t+1}; s_t, \theta_t) \equiv \int_{i \in TC} g_i(\theta_{t+1}; s_t, \theta_t) di \quad (\text{B.16})$$

$$W_{TI,t} \equiv \int_{i \in TI} W_{i,t} di \quad (\text{B.17})$$

$$W_{TC,t} \equiv \int_{i \in TC} W_{i,t} di. \quad (\text{B.18})$$

---

<sup>12</sup>We do not allow contracts contingent on whether a given TI agent becomes a TC agent next period, as this is a positive probability event for the TI agent in question, but a zero probability event for all other agents.

Since all agents of a given type (TI or TC) are identical, we must have  $\frac{g_i(\theta_{t+1}; s_t, \theta_t)}{W_{i,t}} = \frac{g_{TI}(\theta_{t+1}; s_t, \theta_t)}{W_{TI,t}}$  for all agents  $i$  who are TI and  $\frac{g_i(\theta_{t+1}; s_t, \theta_t)}{W_{i,t}} = \frac{g_{TC}(\theta_{t+1}; s_t, \theta_t)}{W_{TC,t}}$  for all agents  $i$  who are TC.

It is clear that market clearing in the asset markets require

$$g_{TC}(\theta_{t+1}; s_t, \theta_t) = 1 - g_{TI}(\theta_{t+1}; s_t, \theta_t). \quad (\text{B.19})$$

By assumption, the wealth of old agents who died between  $t$  and  $t+1$  are redistributed equally among new-born agents. Thus, since a fraction  $\lambda$  of agents die each period, a total wealth  $\lambda W_{t+1}$  needs to be redistributed among new-born agents at the beginning of  $t+1$ . Since a fraction  $\zeta_{TI}$  of new-born agents are TI, new-born TI agents get a total wealth  $\zeta_{TI}\lambda W_{t+1}$  at the beginning of  $t+1$ . Furthermore, the old TI agents at time  $t$  had bought claims for a total wealth of  $g_{TI}(\theta_{t+1}; s_t, \theta_t)W_{t+1}$  at the beginning of  $t+1$ , but only a fraction  $1-\lambda$  actually survives. Therefore, it follows that the TI wealth-share evolves as follows:

$$s_{t+1} = (1-\lambda)g_{TI}(\theta_{t+1}; s_t, \theta_t) + \lambda\zeta_{TI} \equiv s(\theta_{t+1}; s_t, \theta_t) \quad (\text{B.20})$$

Importantly, the wealth-share of TI agents next period does not depend on aggregate consumption shocks - it only depends on shocks to  $\theta_{t+1}$ . Thus, with a slight abuse of notation, we write  $\phi_{TC,t+1} = \phi_{TC}(s_{t+1}, \theta_{t+1}) = \phi_{TC}(\theta_{t+1}; s_t, \theta_t)$ ,  $\phi_{TI,t+1} = \phi_{TI}(s_{t+1}, \theta_{t+1}) = \phi_{TI}(\theta_{t+1}; s_t, \theta_t)$ , and  $\phi_{t+1} = \phi(s_{t+1}, \theta_{t+1}) = \phi(\theta_{t+1}; s_t, \theta_t)$ .

The returns on the equilibrium TI and TC portfolios (conditional on surviving) can therefore be written

$$R_{TI,t+1} = \frac{g_{TI}(\theta_{t+1}; s_t, \theta_t)W_{t+1}}{W_{TI,t} - C_{TI,t}} = \frac{\phi_{TI,t}}{\phi_{TI,t} - 1} \frac{g_{TI}(\theta_{t+1}; s_t, \theta_t)}{s_t} \frac{W_{t+1}}{W_t} \quad (\text{B.21})$$

$$R_{TC,t+1} = \frac{g_{TC}(\theta_{t+1}; s_t, \theta_t)W_{t+1}}{W_{TC,t} - C_{TC,t}} = \frac{\phi_{TC,t}}{\phi_{TC,t} - 1} \frac{1 - g_{TI}(\theta_{t+1}; s_t, \theta_t)}{1 - s_t} \frac{W_{t+1}}{W_t}. \quad (\text{B.22})$$

Substituting (B.21) and (B.22) into (B.13) and (B.14) gives us

$$M_{TC,t+1} = [(1 - \lambda)\beta]^{\frac{\alpha}{\rho}} \left( \frac{\phi_{TC,t+1}}{\phi_{TC,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} \left( \frac{\phi_{TC,t}}{\phi_{TC,t} - 1} \frac{1 - g_{TI}(\theta_{t+1}; s_t, \theta_t)}{1 - s_t} \frac{\phi_{t+1}}{\phi_t} \right)^{\alpha-1} \left( \frac{C_{t+1}}{C_t} \right)^{\alpha-1} \quad (\text{B.23})$$

$$M_{TI,t+1} = [(1 - \lambda)\beta\delta]^{\frac{\alpha}{\rho}} \left[ \theta_t \left( \frac{\phi_{TI,t+1}}{\phi_{TI,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} + (1 - \theta_t) \left( \frac{\phi_{TC,t+1}}{\phi_{TI,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} \right] \times \left( \frac{\phi_{TI,t}}{\phi_{TI,t} - 1} \frac{g_{TI}(\theta_{t+1}; s_t, \theta_t)}{s_t} \frac{\phi_{t+1}}{\phi_t} \right)^{\alpha-1} \left( \frac{C_{t+1}}{C_t} \right)^{\alpha-1}. \quad (\text{B.24})$$

Only the part  $\left( \frac{C_{t+1}}{C_t} \right)^{\alpha-1}$  depends on aggregate consumption growth. The remaining parts of the SDFs are functions of  $\theta_{t+1}$  conditional on  $(s_t, \theta_t)$ .

Equalizing (B.23) and (B.24) yields:

$$\begin{aligned} M_{TC,t+1} &= M_{TI,t+1} \\ &\Leftrightarrow \left( \frac{\phi_{TC,t+1}}{\phi_{TC,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} \left( \frac{\phi_{TC,t}}{\phi_{TC,t} - 1} \right)^{\alpha-1} \left( \frac{1 - g_{TI}(\theta_{t+1}; s_t, \theta_t)}{1 - s_t} \right)^{\alpha-1} \\ &= \delta^{\frac{\alpha}{\rho}} \left[ \theta_t \left( \frac{\phi_{TI,t+1}}{\phi_{TI,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} + (1 - \theta_t) \left( \frac{\phi_{TC,t+1}}{\phi_{TI,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} \right] \left( \frac{\phi_{TI,t}}{\phi_{TI,t} - 1} \right)^{\alpha-1} \left( \frac{g_{TI}(\theta_{t+1}; s_t, \theta_t)}{s_t} \right)^{\alpha-1} \\ &\Leftrightarrow \left( \frac{\phi_{TC,t+1}}{\phi_{TC,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho(\alpha-1)}} \frac{\phi_{TC,t}}{\phi_{TC,t} - 1} \frac{1 - g_{TI}(\theta_{t+1}; s_t, \theta_t)}{1 - s_t} \\ &= \delta^{\frac{\alpha}{(\alpha-1)\rho}} \left[ \theta_t \left( \frac{\phi_{TI,t+1}}{\phi_{TI,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} + (1 - \theta_t) \left( \frac{\phi_{TC,t+1}}{\phi_{TI,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} \right]^{\frac{1}{\alpha-1}} \frac{\phi_{TI,t}}{\phi_{TI,t} - 1} \frac{g_{TI}(\theta_{t+1}; s_t, \theta_t)}{s_t} \\ &\Updownarrow \\ g_{TI}(\theta_{t+1}; s_t, \theta_t) &= s_t \left( \frac{\phi_{TC,t+1}}{\phi_{TC,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho(\alpha-1)}} \frac{\phi_{TC,t}}{\phi_{TC,t} - 1} \times \left[ s_t \left( \frac{\phi_{TC,t+1}}{\phi_{TC,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho(\alpha-1)}} \frac{\phi_{TC,t}}{\phi_{TC,t} - 1} \right. \\ &\quad \left. + (1 - s_t) \delta^{\frac{\alpha}{(\alpha-1)\rho}} \left( \theta_t \left( \frac{\phi_{TI,t+1}}{\phi_{TI,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} + (1 - \theta_t) \left( \frac{\phi_{TC,t+1}}{\phi_{TI,t} - 1} \right)^{\frac{(1-\rho)\alpha}{\rho}} \right)^{\frac{1}{\alpha-1}} \frac{\phi_{TI,t}}{\phi_{TI,t} - 1} \right]^{-1}. \quad (\text{B.25}) \end{aligned}$$

The conditions in (B.7), (B.8), (B.20), (B.25) along with the market clearing condition for current

consumption:

$$\begin{aligned} C_t &= C_{TI,t} + C_{TC,t} = \frac{W_{TI,t}}{\phi_{TI,t}} + \frac{W_{TC,t}}{\phi_{TC,t}} = \left( \frac{s_t}{\phi_{TI,t}} + \frac{1-s_t}{\phi_{TC,t}} \right) W_t \Leftrightarrow \\ \phi_t &= \left( \frac{s_t}{\phi_{TI,t}} + \frac{1-s_t}{\phi_{TC,t}} \right)^{-1} \end{aligned} \quad (\text{B.26})$$

make up the equilibrium conditions in the economy.

Our solution method is backward induction on an  $S \times K$  grid for current states  $(s_t, \theta_t)$ . For each current state, we take as given the  $S \times K$  grids of next period TI and TC wealth-consumption ratios and use interpolation over TI wealth share to get TI and TC wealth-consumption ratios as continuous functions of  $s_{t+1}$ , and solve for current TI and TC wealth-consumption ratios. We initialize both the TI and TC wealth-consumption ratios at the constant wealth-consumption ratio that would have been optimal in a representative TC economy with the same parameters.