Network for Studies on Pensions, Aging and Retirement

Present Bias, Asset Allocation and the Yield Curve

Jorgo Goossens, Bas Werker

NETSPAR ACADEMIC SERIES

DP 10/2020-026

Present Bias, Asset Allocation and the Yield Curve*

Jorgo T.G. Goossens[†] Bas J.M. Werker[‡]

October 5, 2020

Abstract

This paper presents a present-biased general equilibrium model that explains many features of bond behavior. Present-biased investors increase (decrease) short-term (long-term) hedge demands compared to standard preferences. Hence, present bias drives up (down) short-term bond prices (yields) and drives down (up) long-term bond prices (yields), explaining the bond premium puzzle. The model produces realistic bond behavior with a present-bias factor of $\beta = 0.35$ and a long-term annual discount factor of $\delta = 0.97$, in line with the experimental literature. Bond behavior is best explained for a present-bias interval of at most 1 year, providing an estimate for the investor's duration of the present.

Keywords: hyperbolic discounting, portfolio choice, term structure, duration present, behavioral finance

JEL Codes: G41, G11, G12

*We thank Anne Balter, Rob van den Goorbergh, Marike Knoef, Eduard Ponds, Patrick Tuijp for useful comments, and seminar participants at the Netspar Pension Day (2018, 2019), the Netspar International Pension Workshop (2020) and the AFA Ph.D. Poster Session (2021, forthcoming), as well as participants of the APG Asset Management and ABP Seminars.

[†]Tilburg University, department of Econometrics and Operations Research. Corresponding author: j.t.g.goossens@uvt.nl

[‡]Tilburg University, department of Econometrics and Operations Research, department of Finance.

Over the past three decades, researchers analyzing excess bond returns and the term structure of interest rates have uncovered an anomaly that has been termed the bond premium puzzle. For reasonable coefficients of relative risk aversion, standard representative agent general equilibrium models cannot match the sign, magnitude and variability of excess long-term bond returns nor produce an average upward-sloping term structure of interest rates as found in the data (Backus et al., 1989; Campbell and Shiller, 1991; Bansal and Coleman, 1996; Rudebusch and Swanson, 2008; Van Binsbergen et al., 2012). Excess bond returns have empirically a positive relation with maturities up to 10 years (Boudoukh et al., 1999), but standard representative agent models produce an essentially flat relationship between both. The present literature about equilibrium bond pricing needs unreasonable coefficients of risk aversion in excess of 80 to explain the bond premium puzzle (see, e.g., Van Binsbergen et al., 2012), or can only explain the short-term of the term structure of interest rates (e.g., up to maturities of 5 years, as in Wachter (2006)).

In this paper, we argue that it is possible to understand asset properties on a macro level by using experimentally observed behavior on a micro level. Specifically, we explain the bond premium puzzle by refining the way we model investor time preferences. Financial economists almost always capture impatience by assuming that people discount streams of utility over time exponentially, as introduced by Samuelson (1937). Exponential preferences have become the standard preferences and are said to be time consistent, since a person's relative preference for gratification at an earlier date over a later date is the same no matter when he is asked (O'Donoghue and Rabin, 1999). However, over the last few decades, researchers analyzing empirically the behavior of individuals have uncovered a wide range of phenomena that violate the assumptions of Samuelson's exponential discounted-utility framework (see Frederick et al. (2002) for an extensive overview).

Using experimental evidence, people were found to prefer immediate gratification over future gratification. Specifically, discount rates over a short horizon are much higher than discount rates for rewards in the distant future (Thaler, 1981). For this reason, people are biased towards the present. Present bias leads to so-called time-inconsistent behavior, because an agent will deviate from his original plans even when no now information is observed. For example, when presented a choice between one apple on April 1 versus two apples on April 2, if asked on February 1 virtually everyone would prefer the two apples on April 2. But come April 1, given the same choice, most of us prefer the one apple immediately over the two apples tomorrow April 2. In contrast, a time-consistent agent always chooses two apples on April 2, no matter when asked.

Present bias is a pillar of modern behavioral economics (Augenblick et al., 2015), having added generally to economists' understanding of the tensions involved in consumption-savings choices, task performance, temptation and self-control beyond the standard model of exponential discounting (Samuelson, 1937). Given the position of present-biased preferences in the behavioral literature, there is a clear importance constructing equilibrium asset-pricing models that incorporate present bias. A widely used way to model present bias is by a simple generalization of the standard exponential discount model. Rather than only having a long-term discount factor δ, behavioral economists' introduced an additional present-bias factor β (Strotz, 1956; Phelps and Pollak, 1968; Loewenstein and Prelec, 1992; O'Donoghue and Rabin, 1999; Laibson, 1997). This discount model is known as "beta-delta" discounting or quasi-hyperbolic discounting.

A crucial question that arises when using present bias is whether people are aware of their behavior. An agent is sophisticated when he foresees that he will deviate from his original plan and, thus, is aware of his present-biased behavior. A naive agent is unaware of his present-biased behavior and, therefore, does *not* foresee that in the future he will deviate from his original plan. More and more research suggest that models with naive agents seem to better explain observed behavior (O'Donoghue and Rabin, 2015), so we consider a naive representative agent.

In our analysis, we study optimal portfolio choice and equilibrium properties of asset returns. The representative investor has either standard time-consistent preferences, modelled by exponential discounting, or has time-inconsistent preferences in the form of present bias, modelled by quasi-hyperbolic discounting. If the representative investor has standard time-consistent preferences, then we are unable to explain the bond premium puzzle with plausible time preferences. These results replicate earlier findings in the financial economists' literature. Exponential discounting models have difficulty with explaining the behavior of bonds in equilibrium. However, if the representative investor is present biased, then we find that optimal portfolio behavior differs compared to standard time-consistent preferences. Consequently, this behavior creates an explanation for the bond premium puzzle.

In our first set of results, out of three, we show in closed form that present-biased investors have a higher demand for short-term bonds but a lower demand for long-term bonds compared to time-consistent investors. Investment demands contain speculative and hedging demands. Since we assume no predictability of Sharpe ratios, speculative demand is independent of time and the investor's horizon (Merton, 1969; Brennan and Xia, 2002), such that time preferences have no effect on this demand. Hence, differences in demands for short- and long-term bonds between time-consistent and present-biased investors are driven completely by hedging motives. We find that present-biased investors increase short-term hedge demands by 9 percentage points and decrease long-term hedge demands by 17 percentage points compared to time-consistent investors. The rationale is that present-biased investors put less value to the future and, therefore, care less about hedging opportunities for the long run and as such prefer short-term investments. This result holds for our baseline analysis, where we use a present-bias factor of $\beta = 0.35$ and a long-term annual discount factor of $\delta = 0.97$, in line with many estimates in the experimental literature.

Secondly, we find that present bias produces many of the bond characteristics as observed in the markets, while time-consistent preferences do not. Based on 4 measures, we find that present bias explains the bond premium puzzle. Firstly, in equilibrium, present-biased investment behavior matches the observed sign, magnitude and Sharpe ratio of excess (nominal) bond returns. The observed Sharpe ratio of 3 years (10 years) excess bond returns is 0.48 (0.38), the time-consistent Sharpe ratio of 3 years (10 years) excess bond returns is only 0.27 (0.26), while the present-biased Sharpe ratio of 3 years (10 years) excess bond returns equals 0.41 (0.42). Secondly, we show that present bias produces a slope of the yield curve that fits the observed actual yield curve much closer than time-consistent behavior. The observed yield spread equals 1.78%, the time-consistent yield spread is only 1.04%, while the present-biased implied yield spread is 1.88%. Our third and fourth measures are bond risk premia and coefficients of "long-rate" regressions (Campbell and Shiller, 1991), which support the conclusion that present bias fits observed data while time consistency is unable to do so.

In our third set of results, instead of traditionally asking how risk averse the representative investor has to be to explain the bond premium puzzle, we ask a different question. Present-biased preferences require a distinction between the present and the future, but what is the duration of the present? This is an open empirical question in the behavioral literature (Ericson and Laibson, 2018), and an answer allows a connection between theoretical asset pricing models and the experimental literature. We find that the bond premium puzzle is best explained if the duration of the present is at most 1 year.

To the best of our knowledge, we are the first to create and test an equilibrium asset-pricing model with present-biased preferences. We found one related paper by Zou et al. (2014), who study optimal portfolio decisions for sophisticated time-inconsistent investors. They show that the classical Merton (1969) solution is unaffected by (stochastic) hyperbolic discounting for *sophisticated* investors. The finding is similar to our discounting-independent speculative demands for a *naive* investor, but we include hedging demands as well and our result holds for any general discounting structure. Moreover, more and more research suggest that models with naive agents seem to better explain observed behavior (O'Donoghue and Rabin, 2015).

Our paper is not the only paper to address the bond premium puzzle. Early work of Backus et al. (1989) shows that a representative agent in an endowment economy with power utility cannot explain the puzzle. Rudebusch and Swanson (2008) examine the bond premium puzzle with a macro-economic dynamic stochastic general equilibrium model, but they conclude that the bond premium puzzle remains even if they include large and persistent habits, and labor market frictions. The work of Van Binsbergen et al. (2012) considers also a dynamic stochastic general equilibrium model, in which they include Epstein-Zin recursive preferences, but they need a risk aversion parameter between 40 and 80 to explain the observed bond properties. More successful is the approach of Wachter (2006), who explains the bond premium puzzle up to a maturity of 5 years by creating a consumption-based model with external habit, and calibrating it to macro-level data.

While present bias can potentially be a helpful way of thinking about the data, we emphasize that it is only a potential ingredient in an equilibrium model, and by no means a complete description of the facts. For one thing, we have not touched upon time-series properties of bonds. An equilibrium model that combines present bias with other forms of experimentally observed behavior is likely to be superior to a model that uses present bias alone. Promising directions might be reference points, loss aversion, recency effects and habit formation. Specifically, it might be promising to study asset prices with a model that combines factors leading to a preference for improving sequences with time-inconsistent preferences (Ericson and Laibson, 2018).

I. The model

Present-biased investors maximize utility over consumption to find the optimal asset allocation. The investment options are a risk-free asset, a stock, a short-term bond and a long-term bond. For now, think of these bonds as a 3-year constant maturity bond and a 10-year constant maturity bond respectively. Equilibrium prices adjust such that the supply equals demand for these assets.

A. Preferences: Present bias

Extensive experimental work suggests that present bias is an important feature of how people evaluate intertemporal decisions — see Frederick et al. (2002) for an overview. The general finding within the experimental literature is a substantial present bias. Empirical analyses have demonstrated how present bias can improve our understanding of behavior in various economic contexts (O'Donoghue and Rabin, 2015).

Present bias is an old idea, and the notion of people being susceptible to immediate gratification goes back as far as the Greeks. Psychologists working with animals in the 1960s and 1970s proposed hyperbolic discounting — a functional form of discounting that generates present bias — as a natural way to represent how animals respond to time delays, and later research extended this idea to humans.[1] Economists in the 1960s started investigating time-inconsistent behavior and used as present bias the β, δ functional form. Present-biased preferences are time inconsistent: e.g., myself at time t prefers to exercise at time $t' > t$, but at time t' I decide to take a nap at time t' (rather than the planned exercise). The now popular β, δ simplification of hyperbolic discounting is known as quasi-hyperbolic discounting (Phelps and Pollak, 1968; Laibson, 1997).

Like exponential discounting, present bias is a model of discounting. Present bias is extremely similar to standard exponential discounting and, for this reason, probably has a wide success. Suppose that intertemporal preferences from the perspective of period t can be represented by $U^t = \sum_{j=t}^{T} D(j-t)u(W_j)$, in which $u(W_j)$ is utility experienced from consuming wealth W in period $j \geq t$ and $D(x)$ reflects the (deterministic) discounting associated with delay $x \in \{0, 1, 2, ...\}$. Many variants of time inconsistency and present bias exist, but the generalized quasi-hyperbolic discount function takes the form (Harris and

[1] Ainslie, 1992 presents an overview

Laibson, 2013)

$$D(x) = \begin{cases} \delta^x & \text{if } x = [0, T_S] \\ \beta \times \delta^x & \text{if } x \in (T_S, T_L]. \end{cases} \tag{1}$$

With this functional form, $\beta = 1$ corresponds to exponential discounting, while $\beta \in (0, 1)$ reflects present bias. Figure 1 presents a comparison of the discount structure by plotting one realization of the exponential discount function, the (instantaneous) quasi-hyperbolic discount function with $T_S = 0$ (as proposed by Laibson, 1997) and a generalized hyperbolic discount function.

Figure 2 presents a time-line of the discounting model. Present bias requires a distinction between the end of the present and the arrival of the future. T_S is the duration of the present and it indicates the short-term planning horizon.[2] After date T_S the future starts, which lasts up to the terminal long-term horizon T_L. The discount function decays exponentially at rate $-\ln\delta$ up to time T_S, drops discontinuously at T_S to a fraction β of its level just prior to T_S, and decays exponentially at rate $-\ln\delta$ thereafter.

The model's predictions depend on whether one is aware of how time preferences change over time (sophisticated), unaware of how preferences change over time (naive), or something in between (partially naive). We incorporate the idea of a naive present-biased investor. The current self is unaware of the future self's present bias, and instead believes that future self will discount exponentially. More and more research suggests that models with naive agents seem to better explain observed behavior (O'Donoghue and Rabin, 2015). An alternative would be to assume sophisticated present bias, where the current self is aware of the future self's present — a higher form of rationality. Sophistication would bring an additional layer of complexity, because the current self has to solve an equilibrium game with his future selves.

B. Financial market

The representative investor has access to an arbitrage-free complete financial market consisting of a stock, constant maturity bonds and cash. The short rate r_t is assumed to be

[2]The end of the present at time T_S might be stochastic, as shown by Harris and Laibson (2013). In our analysis, we work with the frequently used deterministic horizons. In our theoretical derivations and empirical findings, the determinstic horizons can be replaced by stochastic horizons.

affine in an N-dimensional factor of state variables $\boldsymbol{F}_t$

$$r_t = A_0 + \boldsymbol{\iota}' \boldsymbol{F}_t, \tag{2}$$

where $\boldsymbol{\iota}$ denotes an N-dimensional vector of ones.[3] The factors $\boldsymbol{F_t}$ follow an N-dimensional multivariate Ornstein-Uhlenbeck process:

$$d\boldsymbol{F_t} = \boldsymbol{\kappa} \left(\boldsymbol{\theta} - \boldsymbol{F}_t \right) dt + \boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,t}, \tag{3}$$

where $\boldsymbol{\kappa}$ is the $N \times N$ diagonal mean-reversion speed matrix, $\boldsymbol{\theta}$ is the N-dimensional column vector of long-run averages, $\boldsymbol{\sigma_F}$ is the $N \times N$ lower triangular covariance matrix with strictly positive elements on its diagonal and $\boldsymbol{Z}_{F,t}$ is the N-dimensional column vector of independent standard Brownian motions.

Empirically, not all the parameters of the affine model can be identified and certain assumptions are necessary. We assume that there are two factors driving the short rate r_t. For identification purposes, we normalize the long-run means of the factors $\boldsymbol{\theta}$ to zero (De Jong, 2000).

The investment opportunities depend on the pricing kernel in the economy, which determines the expected returns on all securities in the financial market. We assume absence of arbitrage and, thus, the existence of a stochastic discount factor process M_t with $M_0 = 1$:

$$\frac{dM_t}{M_t} = -r_t dt - \boldsymbol{\lambda}' d\boldsymbol{Z}_t, \tag{4}$$

where $\boldsymbol{\lambda} = [\lambda_S; \boldsymbol{\lambda}_F]$ and $\boldsymbol{Z}_t = [Z_{S,t}; \boldsymbol{Z}_{F,t}]$ are $(N+1)$-dimensional vectors. λ_S is the constant price-of-risk for the stock and $\boldsymbol{\lambda}_F$ is the N-dimensional vector with the constant prices-of-risk for the bonds. $Z_{S,t}$ is a standard Brownian motion representing shocks to the stock, and it is independent of the shocks $\boldsymbol{Z}_{F,t}$ to the state variables.

The dynamics of the stock price and the dynamics of the N-dimensional vector of bond

[3] Without loss of generality, see, e.g., De Jong (2000), we thus assume that the instantaneous interest rate equals a constant plus the sum of the factors.

prices follow from

$$\begin{aligned}\frac{dS_t}{S_t} &= (r_t + \boldsymbol{\sigma}'\boldsymbol{\lambda})dt + \boldsymbol{\sigma}' d\boldsymbol{Z}_t \\ \frac{d\boldsymbol{P}_t}{\boldsymbol{P}_t} &= (r_t - \mathcal{B}(\tau)'\boldsymbol{\sigma}_F \boldsymbol{\lambda}_F)dt - \mathcal{B}(\tau)'\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,t},\end{aligned} \tag{5}$$

where $\boldsymbol{\sigma} = [\sigma_S; \boldsymbol{\sigma}_{FS}]$ is an $N+1$-dimensional vector. σ_S is the volatility parameter of the stock and $\boldsymbol{\sigma}_{FS}$ is an N-dimensional vector governing the covariance between stock and bond returns.[4] $\mathcal{B}(\tau)$ follows from

$$\mathcal{B}(\tau) = [\boldsymbol{B}(\tau_1)\iota, ..., \boldsymbol{B}(\tau_N)\iota],$$

where

$$\boldsymbol{B}(t) = (\boldsymbol{I} - \exp(\boldsymbol{\kappa} t))\, \boldsymbol{\kappa}^{-1}. \tag{6}$$

in which τ_j for $j = 1, ..., N$ denotes the maturity of bond j. $\boldsymbol{B}(.)$ is an $N \times N$-dimensional matrix, so $\mathcal{B}(.)$ has the same dimensions.

C. Consumption and asset allocation

Turning now to the issue of optimal portfolio strategies for long-lived investors, we consider two classical cases. In the first, the investor solves a *terminal* consumption problem. He is concerned with maximizing the expected utility of wealth on some fixed horizon T. This problem has the merits of both simplicity and of clarifying the role of the horizon. The second case we consider is an *intermediate* consumption problem, which boils down to solving multiple *terminal* consumption problems sequentially. This problem, which is only slightly more complicated, corresponds to an investor who is concerned with maximizing the expected value of a time-additive utility function defined over lifetime consumption.

To shed some light on the second case, consider that total available wealth at each time t has to be split in an amount needed to finance the first consumption moment, an amount needed to finance the second consumption moment, and so forth. To finance each consumption moment you use (a part of) total wealth, which we call a money pot. Thus,

[4] The dynamics of the bond prices follow from the explicit bond prices in the financial market, which can be directly obtained by Lemma 1 for $\alpha = 1$.

think of total available wealth at each time t as the present value of the first money pot, the second money pot, and so forth. Essentially, optimizing the allocation for the first money pot is a stand-alone terminal wealth problem, the second money pot is a stand-alone terminal wealth problem as well, and so forth. What matters eventually, is the optimal division of total wealth over each of the money pots. This approach provides a convenient way to think about intermediate allocation and smoothing of consumption over time by sequentially solving multiple terminal wealth problems. This approach works thanks to the completeness of the market.

The representative investor maximizes CRRA utility, with risk aversion parameter γ, over consumption to find the optimal asset allocation. At each time t, the representative agent discounts utility of all future intermediate consumption W_j according to his discount structure. Think of consumption as consuming a pot of money with monetary value W_j. All derivations below hold for general deterministic discount structures $D(x)$. Eventually, we use the quasi-hyperbolic formulation in (1) as a specific case to demonstrate the empirical success of the present-biased model.

We present explicit analytical solutions for the optimal consumption and asset allocation strategies for a general discount structure. The investor has to both determine the optimal allocation of wealth over time for consumption, and decide on the optimal investment strategy for each money pot. We show that speculative demands are independent of the discount structure, but hedge demands depend on the discount structure $D(x)$. The reason is that planned consumption matters for the optimal hedge allocation through the planning horizon dependence. Namely, consumption itself depends on the discounting model. As methodology, we use the martingale method of Cox and Huang (1989) in both cases.

Terminal consumption problem — Today, the representative investor maximizes expected utility of future wealth W for investment horizon T:

$$\max_{W_T} E_0 \left[\frac{W_T^{1-\gamma}}{1-\gamma} \right], \tag{7}$$

subject to his budget constraint

$$E_0 \left[W_T M_T \right] = W_0, \tag{8}$$

where W_0 is initial total available wealth. Obviously, the terminal consumption formulation is independent of any discount structure. For this reason, discounting models do not influence the optimal consumption and the optimal asset allocation. We formalize this observation in the theorem below.

Theorem 1. *For an investor that maximizes terminal wealth:*

1. *The optimal fraction of wealth invested in constant maturity τ-year bonds at time t for investment horizon T is a vector*

$$\boldsymbol{\pi}_B^*(t,T) = \frac{1}{\gamma}\left(\frac{\lambda_S}{\sigma_S}\boldsymbol{\sigma}'_{FS} - \boldsymbol{\lambda}'_F\right)(\boldsymbol{\sigma}_F)^{-1}(\mathcal{B}(\tau)')^{-1} + (1-1/\gamma)\boldsymbol{\iota}'\boldsymbol{B}(T-t)(\mathcal{B}(\tau)')^{-1},$$

   and the optimal fraction of wealth invested in stocks is

$$\pi_S^* = \frac{\lambda_S}{\gamma\sigma_S},$$

   whereas the remainder is invested in cash.

2. *The optimal consumption path at time t for horizon T_j is W_t^*, and is explicitly given in Appendix A by equation (25).*

Theorem 1 is identical to the 1-factor model result of Brennan and Xia (2002), but our result also holds for a N-factor Vasicek model.[5] The fraction allocated to stocks is independent of time t, and equals the market price of risk divided by the risk aversion and the volatility of stock market risk. So, although the stock market returns are influenced by all of the $N+1$ sources of risk, only the stock market shocks matter for the optimal investment fraction to stocks. The result equals that of Merton (1969).

The optimal fraction of wealth allocated to bonds has two components: speculative demand and hedge demand. The first component closely resembles the form of the optimal stock market allocation. The second component depends on the investment horizon of the investor and the investor uses this component to hedge against unfavorable changes in the state variables. Note that indeed the optimal asset allocation and consumption path are independent of the discounting model.

[5]For completeness, we provide a proof in Appendix A for a N-factor Vasicek model.

Intermediate consumption problem — We now start at time t with the allocation of total available wealth to the money pots for each year. Think of the total available wealth to be used to finance the individual money pots $j = 1, ..., n$ for the years T_j for fixed n. Each money pot has its own optimal consumption strategy and investment policy during the years $t + h = t, ..., T_j$. So, we need to keep track of these 3 dates and, for this reason, we use the following notation. $W_{t,T_j,t+h}$ denotes optimal planned wealth at time $0 \leq t \leq T_j$ (first index), where the investor plans for money pot j in year $T_j \geq t$ (second index) with in between date $t + h$ (third index). Only in case the agent has time-consistent preferences $W_{t,T_j,t+h}$ will not depend on the first index t, because at every time t he makes identical decisions.

The optimization problem at each time t becomes:

$$\max_{\{W_{t,T_1,T_1},...,W_{t,T_n,T_n}\}} \mathbb{E}_t \left[\sum_{j=1}^{n} D(T_j - t) \frac{W_{t,T_j,T_j}^{1-\gamma}}{1-\gamma} \right], \tag{9}$$

subject to his budget constraint at each time t

$$\mathbb{E}_t \left[\sum_{j=1}^{n} W_{t,T_j,T_j} M_{T_j} \right] = W_t M_t \tag{10}$$

with W_t the total available wealth at time t (including portfolio returns). Clearly, the intermediate consumption formulation depends on the general discount structure $D(x)$. Because the way of discounting influences the optimal consumption path, the optimal asset allocation changes accordingly since the distribution of total available wealth over the money pots changes. We formalize this intuition in the next theorem.

Theorem 2. *For an investor, with general discount structure $D(x)$ and delay $x = T_j - t$, that maximizes intermediate consumption:*

1. *The optimal fraction of actual total wealth invested in asset $i = \{$stock, constant maturity τ-year bond, cash$\}$ at time t for investment horizon T_j is*

$$\omega_i^*(t, T_j) = \frac{\sum_{j,j>t}^{n} \pi_i^*(t, T_j) W_{t,T_j,t}^* \left(D(T_j - t)\right)}{\sum_{j,j>t}^{n} W_{t,T_j,t}^* \left(D(T_j - t)\right)} \tag{11}$$

where $\pi_i^*(t, T_j)$ follows from Theorem 1.

2. *The actual optimal consumption path at each time t for all horizons T_j is $W^*_{t,T_j,t}(D(T_j - t))$, and is explicitly given in Appendix C by equation (41).*

Theorem 2 shows that the optimal consumption path is a function of the discount structure $D(x)$ and, consequently, the optimal asset allocation depends on the discount structure through consumption. The consumption path $W^*_{t,T_j,t}(D(T_j - t))$ prescribes how much the investor at time t should optimally allocate to, and consume at, money pot T_j. The consumption rule $W^*_{t,T_j,t+h}(D(T_j - t))$, given in Appendix C by (41), prescribes for each money pot T_j, at each time t, the optimal planned consumption paths for all in-between dates $t \leq t + h \leq T_j$. Alternatively, $W^*_{t,T_j,t}(D(T_j - t))$ is the *actual* consumption path of total wealth, while $W^*_{t,T_j,t+h}(D(T_j - t))$ also includes the *planned* consumption paths.

Regarding the actual optimal asset allocation, observe that the optimal proportions $\pi^*_i(t, T_j)$, for each asset, invested for *each* separate money pot remain identical to the terminal consumption problem and, thus, independent of the discount structure. However, note that we are interested in the investment strategy of *total* wealth, i.e., the summation over all money pots j. The optimal proportions $\omega^*_i(t, T_j)$, for each asset, invested over *all* money pots become dependent on the discount structure. The reason is that the actual allocation of total wealth $W^*_{t,T_j,t}(D(T_j - t))$ over the money pots — i.e., your actual optimal consumption path — depends on the investor's time preferences through the discounting structure.[6]

More specifically, Theorem 2 shows that speculative demands (including the stock allocation) are independent of the discount structure, while hedge demands depend on the discount structure. Remember, from Theorem 1, that the stock allocation and speculative bond demands are independent of both current time and the investment horizon. In other words, the stock allocation and speculative bond demands for money pot 1 are identical to the stock allocation and speculative bond demands for money pot 2, and so forth. In mathematical terms, $\pi^*_i(t, T_j)$ is independent of every money pot j in (11). So, for the stock allocation and speculative bond demands we respectively have $\omega^*_S = \pi^*_S$ and $\omega^*_{B,spec} = \pi^*_{B,spec}$.

The hedge bonds demands in Theorem 1 depend on both current time and investment horizon. In other words, the hedge bond demands for money pot 1 could be different from the hedge bond demands for money pot 2, and so forth. In mathematical terms, $\pi^*_i(t, T_j)$

[6]Similar to planned consumption, one can also easily derive the *planned* optimal investment strategy for all in-between dates $t \leq t + h \leq T_j$, at each time t for money pot T_j. Use $W^*_{t,T_j,t+h}(D(T_j - t))$ rather than $W^*_{t,T_j,t}(D(T_j - t))$ in Theorem 2.

depends on the consumption moment of each money pot j in (11). Therefore, the optimal fraction of *total* invested wealth depends on the discount structure through the optimal consumption path.

Theorem 2 generalizes the result of Brennan and Xia (2002) and Zou et al. (2014) considerably, since (i) we include hedging demands, through a N-factor Vasicek model, and because (ii) we consider a general discount structure.

D. Equilibrium

Here, we introduce the concept of general equilibrium, such that we can compute equilibrium excess bond returns and equilibrium yields by solving for the equilibrium bond prices-of-risk $\boldsymbol{\lambda}_F$.

Definition 1. *The market is in general equilibrium if both of the following conditions are satisfied:*

1. *The representative investor solves the intermediate consumption problem (9) subject to his budget constraint (10).*

2. *Bond markets clear continuously, such that for all $t \in [0, T_n]$ we have:*

$$\boldsymbol{\omega}_B^*(t, T_j) = \hat{\boldsymbol{w}}_B(t)$$

where $\boldsymbol{\omega}_B^*(t, T_j)$ is the optimal bond demand from Theorem 2 and $\hat{\boldsymbol{w}}_B(t)$ is the exogenously given supply of bonds in the economy.

The first condition determines the demand for the stock, bonds and cash in the economy. The representative investor is infinitely lived. So, at each time t, the representative investor solves the intermediate consumption problem with discount structure $D(x)$ and terminal horizon T_n, i.e., the last pot of money.

The second condition states that the demand and supply for bonds is equal to each other in equilibrium. The market is complete, so the number of bonds with different maturities equals the number N of state variable factors. Thus, to match the supply and demand of N bonds, we need N free parameters for an exactly identified system. Standard macroeconomics uses prices to match demand and supply. We follow this approach, and we use the N prices-of-risk in the vector $\boldsymbol{\lambda}_F$ as free parameters to match demand and supply of all N bonds in equilibrium. We denote the estimated equilibrium prices-of-risk by $\hat{\boldsymbol{\lambda}}_F$.

Using the estimated equilibrium prices-of-risk, we compute equilibrium excess returns and equilibrium yields. Note that the investor may still invest in the stock market, but we do not impose equilibrium in that market. Since the optimal stock allocation is independent of the discounting model, there is no need to examine how discounting influences equilibrium stock prices (and, as a consequence, the money market).

II. Empirical findings

Given the analytical solutions derived in the previous section, in this section we can explicitly calculate optimal asset allocations, and equilibrium bond prices and yields. We consider both the exponential discounting model and the quasi-hyperbolic discounting model, reflecting time-consistent and present-biased preferences respectively.

A. Data

We use monthly zero-coupon yields with multiple maturities and monthly stock returns from 1 October 1976 to 1 January 2019. Source of the yield data is the Fed database at the St. Louis Federal Reserve and source of the stock returns is Kenneth R. French's Website.

Regarding the yield data, we take the yields with maturities 2 years, 3 years, 5 years, 7 years and 10 years from the Treasury Constant Maturity Rates, while the shorter-term yields with a maturity of 3 months and 1 year are from Treasury Bills: Secondary market rates, because these series contain less missing values. All yield data is reported in percent per annum, and annualized using a 360-day year or bank interest. As risk-free rate we use the 3 month treasury bill.

As for the stock data, we use the stock market returns from Kenneth R. French's Website. This is a value-weighted index of all CRSP firms incorporated in the U.S. and listed on the NYSE, AMEX or NASDAQ.

Table 1 gives some descriptive statistics of the yield and stock data. On average, the term-structure of interest rates is upward sloping. Moreover, the long-maturity interest rates are somewhat less volatile than the short rates. High interest rates show up in the 1980s, while the low interest rates show up at the end of our sample period. Clearly, stocks have higher returns, at the cost of higher volatility.

To determine the supply of bonds in the economy, we use monthly data from 1 October 1976 to 1 January 2019 on U.S. government debt in line with e.g. Wachter (2006) and Rudebusch and Swanson (2008). The source is Datastream's U.S. Maturity Distribution, Interest Bearing Public Debt. Table 2 provides some descriptive statistics, while Figure 3 gives a visualization of the U.S. debt development throughout time. The ratio of short-term debt — maturities smaller than 1 year, and between 1 and 5 years — to total debt declines during our sample period from roughly 84% to 69%. On the other hand, the ratio of long-term debt — maturities exceeding 5 years — to total debt increases from 15% at October 1976 to 30% at January 2019. Table 2, Panel C, shows that approximately on average 73% of the debt has a maturity lower than 5 years, while 17% of the debt has a maturity higher than 5 years.

Regarding the choice of the sample period, we did not go further back than October 1976, because yield curve estimations and debt data from earlier years contain relatively high standard errors and missing values.

B. Calibration and estimation

We calibrate our model to the monthly market data by using a standard Kalman filter and maximum likelihood estimation. Exact discretization of our economy is possible by writing the financial market processes as a multivariate Ornstein-Uhlenbeck process.

Regarding the choice of maturities for the yields, we follow De Jong (2000). For maturities over 10 years, bond data is somewhat scarce, so interpolation is less accurate. In the very short-term interest rates of one and two months, there are sometimes exceptionally large one-period changes. We feel more confident using interest rates of 3 months and longer, and of 10 years and shorter. To keep estimation feasible, we confine ourselves to four maturities — 3 months, 1 year, 5 years and 10 years.

The maximum likelihood estimation starts from multiple initial values to prevent that the optimizer finds a local optimum. We estimate 12 model parameters and 4 measurement errors for each maturity. We assume that each maturity has its own measurement error, such that the variance of the errors depends on maturity (Geyer and Pichler, 1999). Each error is drawn from a uniform distribution, and both serially and cross-sectionally uncorrelated.[7]

Table 3 shows the Kalman estimates and standard errors for our financial market. The standard errors follow from the square root of the diagonal elements of the inverted Hessian matrix $\sqrt{H_{ii}^{-1}}$ for $i = 1, .., k$ where k is the number of estimated parameters. Factor 2 exhibits stronger mean reversion than factor 1. The long-term mean of the short-rate A_0 is estimated at 3.59%. Both prices-of-risk for the two factors are negative. The negative sign of $\sigma_{F,21}$ implies that there is negative correlation between the two factors. The price-of-risk of the stock λ_S equals 0.4095, and the volatility of the stock σ_S is 15.04% per annum.

Both factors have an interpretation. Factor 1 is very highly correlated with the 10-year yield and factor 2 is very closely related to the spread between the 3-month yield and the 10-year yield. These results are in line with De Jong (2000), who identifies the factors respectively as a level factor and a slope factor.

Since our model has two factors driving the short rate, we need two constant maturity bonds (of different maturities) to assure market completeness. To determine the maturities of these two bonds, we use the U.S. government debt data. In line with Horvath et al. (2017), we assume that all debt with maturity lower than 5 years reflects a U.S. government bond with a maturity of 3 years, while all debt with maturity higher than 5 years reflects a U.S. government bond with a maturity of 10 years. So, in equilibrium, we match the supply of 3-year and 10-year U.S. government bonds with the model-implied demand for both bonds. Thus, we estimate the equilibrium prices-of-risk every October of each year by matching the demand with supply for both 3-year and 10-year bonds.[8]

C. Preferences

We compare the asset allocations and equilibrium bond prices of a representative time-consistent investor with a representative present-biased investor. We assume that the risk preferences are the same for both investors. Since our benchmark model features CRRA utility, we assume that both representative investors have a risk aversion level of $\gamma = 10$.

[7]Assuming only one error variance for all maturities is a possible and convenient simplification.

[8]Because of our monthly nature of the data, we can also compute equilibrium prices-of-risk each month of the year, but results stay nearly identical to the annual matching (differences emerge at 0.1 basispoint), which yields some computational advantage.

This value is the maximum value considered plausible by Mehra and Prescott (1985) and a frequent assumption by financial economists (see, e.g., Van Binsbergen et al., 2008).

Regarding the *time-consistent* investor, we assume time preferences that are common in the financial economists' literature. Van Binsbergen et al. (2012) try to match characteristics of bond returns and yields with a discount factor of 0.997, which is "a standard result in the literature". Barberis and Huang (2001) study individual stock returns with a value of 0.98, Campbell and Cocco (2003) study mortgage choice with a discount factor of 0.98 and Wachter (2006) studies characteristics of bond returns and yields with a discount factor of 0.98. We assume an annual long-term discount factor of $\delta = 0.97$, which implies approximately an annual discount rate of 3% per year. Our assumed factor is a bit lower than in the literature, which biases our results in favor of the time-consistent agent.

Regarding the *present-biased* investor, we assume a long-term discount factor of $\delta = 0.97$ as well. We do this for comparability reasons, such that all variation in our results comes from the present-bias factor only. We assume a naive present-biased investor with a present-bias factor of $\beta = 0.35$. The assumption of naivete implies that the investor is unaware of his present-biased behavior and, therefore, does *not* foresee that in the future he will deviate from his original plan. An alternative would be to assume a sophisticated investor, who foresees that he will deviate from his original plan and, thus, is aware of his present-biased behavior.[9] However, more and more research suggest that models with naive agents seem to better explain observed behavior (O'Donoghue and Rabin, 2015).

The present-bias factor of $\beta = 0.35$ implies substantial present bias. The estimates for time preferences, including present bias and discount factors, vary quite a lot within the experimental literature. As we focus on financial decision making, we feel most comfortable by using the estimates of Laibson et al. (2015). Their estimates are particularly helpful for our research because the authors distinguish between time-preference estimates for sophisticated and naive agents, and their model is close to ours in terms of the realm of decision making. Namely, they use a structural life-cycle model with actual data on consumption and investment decisions. Their model estimates a present-bias factor of $\beta = 0.35$ and a long-term discount factor of $\delta = 0.97$ for a naive agent.

Many other time preference estimates exist in the literature, but we consider them irrel-

[9]A sophisticated investor has characteristics of so-called hyper rationality, because the current self plays a dynamic game with his future selves.

evant for our research purposes because the areas of decision making are too distinct from our model. Using data on food stamps, Shapiro (2005) estimates a present bias parameter of 0.96. DellaVigna and Paserman (2005) estimate a present bias-factor of 0.89 for naive agents and 0.90 for sophisticated agents, in a job search model. Focusing on primary rewards (such as fruit juice or water) in a neuroimaging laboratory experiment, McClure et al. (2007) estimate a present-bias factor of 0.52. Paserman (2008) studies job search and cannot distinguish between naivete and sophistication; his present-bias factor estimates range from 0.40 to 0.48 for low and medium wage works. Fang and Silverman (2009) studies labor supply decisions and estimates a present-bias factor of $\beta = 0.36$ for naive agents. Augenblick et al. (2015) study real effort tasks and find a present-bias factor of 0.89 for effort-based decisions, but 0.97 for money-based decision. Using decisions over time-dated monetary rewards, Balakrishnan et al. (2017) find similar estimates approximately equal to 0.90. In an experiment on unpleasant transcription tasks, Augenblick and Rabin (2019) find a present-bias factor 0f 0.83.

Interestingly, estimated long-term discount factors for exponential discounting in the experimental literature differ substantially from the typically assumed values by financial economists. In many cases, financial economists assume discount factors that are higher than typically found in the experimental literature. In other words, they tend to underestimate the degree of actual impatience. For example, Laibson et al. (2015) estimate a long-term discount factor $\delta = 0.63$ in the exponential discounting model, while typically values of $\delta = 0.97 - 0.99$ have been used in previous financial economist' studies. Shapiro (2005) finds $\delta = 0.23$, DellaVigna and Paserman (2005) find $\delta = 0.59$ and Fang and Silverman (2009) find $\delta = 0.41$.

Regarding the *discount and utility structure*, we assume in the discounting model $D(x)$ in equation (1) that the naive investor uses a present with a duration $T_S = 1$ year and a long-term planning horizon $T_L = 10$ years. In terms of the intermediate consumption model (9), we assume that the investor has two consumption moments: one consumption moment during the present and one consumption moment during the future. We let the first consumption moment coincide with time T_S and the second consumption moment coincides with time T_L.[10]

[10]The model would qualitatively yield the same results if the consumption moments differ from the bounds: e.g., the first consumption moment takes place somewhere during the present $[0, T_S]$ and the second

Specifically, the 2-period model works as follows. The investor discounts utility of the 2 consumption moments by $D(x)$. The investor is infinitely lived, since time runs from today towards infinity. At each time t, the investor uses the discount structure $D(x)$ to discount future utility of money pot 1 and money pot 2. Money pot 1 coincides with $T_S = 1$ year and corresponds to consumption during the present. Money pot 2 coincides with $T_L = 10$ years and corresponds to consumption during the future. Thus, at each year t, the investor discounts utility of money pot 1 by δ — because the consumption moment falls within the present — and the investor discounts utility of money pot 2 by an additional scalar β — because the consumption moment falls in the future. So, at each year t, the investor decides how to optimally split wealth over the present and the future, and how to optimally invest for these two moments.

Empirically, we can also solve a consumption model with more than two consumption moments — as we did in the theoretical derivations for n money pots. But, we feel that the 2-period model is simpler, since it reveals the key mechanisms and intuition for present-biased investing which requires a distinction between the present and the future. The 2-period model cleanly groups behavior in the present and the future. We do so, because we feel that this provides a helpful discipline. Moreover, we do not allow ourselves the luxury of selecting the time-preference parameters that would fit the data best.

D. Asset allocation

We show the optimal consumption strategy and optimal investment strategy for a time-consistent investor and a present-biased investor. The optimal consumption path that is actually followed by the investor is given by $W^*_{t,T_j,t}\left(D(T_j - t)\right)$. The optimal *planned* consumption path is given by $W^*_{t,T_j,t+h}\left(D(T_j - t)\right)$. Both actual and planned consumption paths depend on the discounting model. For a time-consistent investor, the actual consumption path is identical to the planned consumption path. Since the investor is time consistent, he sticks to his plan. The investor consumes at each time $t + h$, what he planned to consume at time t given the state of the world.

Regarding the present-biased investor, the actual consumption path differs from the planned consumption path. Since the investor is time inconsistent, he deviates from his

consumption moment takes place somewhere during the future $(T_S, T_L]$. For interpretability, we let them coincide.

plans. At time t, the present-biased investor plans how much to consume at time $t + h$. But, if the investor actually arrives at time $t + h$, he decides to consume more than initially planned at time t. For this reason, a present-biased investor has a higher consumption rate at every time t than a time-consistent agent. The consumption rate equals the ratio of current consumption over current total wealth, and for both investors this ratio is concave upward sloping and reaches 1 at the terminal horizon.

Optimal investment strategies are determined by the Kalman calibrated parameter values. Table 4 summarizes the optimal asset allocations under present bias and time consistency for a 3-year bond, a 10-year bond, a stock and cash. The first observation is that speculative demands for the short-term bond, the long-term bond and the stock are independent of being present biased or time consistent. Secondly, the short-term hedge demand for a present-biased investor is 9 percentage points higher than for a time-consistent investor, while the long-term hedge demand for a present-biased investor is 17 percentage points lower than for a time-consistent investor. This is entirely driven by the higher consumption rate of a present-biased investor. The key mechanism is that a present-biased investor cares less about the future and, therefore, cares less about hedging opportunities for the long run but values short-term opportunities. The total demand across all assets sums up to 1, such that the investor invests all of his total wealth.

Table 5, Panel A, shows the optimal asset allocation for a present-biased investor for three different present-bias factors. The main observation is that a lower present-bias factor drives up the short-term hedge demands but lowers the long-term hedge demands, because the less the investor cares about future consumption opportunities. An extremely low present-bias factor of 0.05 causes the investor to go short in terms of 10-year bond hedge demands.

One might wonder, whether we could quantitatively reach similar results if we use exponential discounting model only. Table 5, Panel B, shows that a time-consistent investor with a long-term annual discount factor of $\delta = 0.86$ replicates the optimal asset allocation decisions for a present-biased investor with $\beta = 0.35, \delta = 0.97$. The hedge demands for the 3-year and 10-year bonds equal 0.49 and 0.04 of total wealth.

However, note that financial economists typically find a value of $\delta = 0.86$ too low, since it implies an annual discount rate of roughly 16%. On the other hand, the experimental literature finds $\delta = 0.86$ too high in the exponential discounting model, because their estimates range from $\delta = 0.23$ to $\delta = 0.63$.

E. Equilibrium returns and yields

The bond premium puzzle states that for reasonable coefficients of relative risk aversion, standard representative agent general equilibrium models cannot match the sign, magnitude and variability of excess long-term bond returns nor produce an average upward-sloping yield curve as found in the data (Backus et al., 1989; Campbell and Shiller, 1991; Bansal and Coleman, 1996; Rudebusch and Swanson, 2008; Van Binsbergen et al., 2012). The present literature about equilibrium bond pricing needs unreasonable coefficients of risk aversion in excess of 80 to explain the bond premium puzzle (see, e.g., Van Binsbergen et al., 2012), or can only explain the short-term of the term structure of interest rates (e.g., up to maturities of 5 years, as in Wachter (2006)).

In this section, we find that present bias leads to the observed sign, magnitude and Sharpe ratio of excess bond returns, and produces an average upward-sloping yield curve as found in the data. Present-biased behavior drives down (up) short-term bond yields (prices) but drives up (down) long-term bond yields (prices). This is the key mechanism that explains all of our findings below. Namely, a present-biased investor has higher short-term hedge demands but lower long-term hedge demands compared to a time-consistent investor. Moreover, we confirm that a standard time-consistent investor cannot match the observed behavior of bonds for reasonable parameters.

To examine the model's ability to fit the data, we use 4 measures that are common in the literature. The literature has focused on the risk premium to describe bond behavior. However, it is not directly observed in the data. Accordingly, the literature has also focused on three other empirical measures that are more easily observed: excess bond returns, the yield spread and the slope coefficients from a Campbell-Shiller (1991) "long-rate" regression.

1. Expected excess bond returns

Our main and first measure is the excess return on short-term and long-term bonds. That is, the expected instantaneous τ_j-year bond return in excess of the the short rate r_t. The mean and standard deviation of excess bond returns are a popular measure for characterizing bond behavior (Rudebusch and Swanson, 2008). For the case of a τ_j-year bond, this excess return along with its standard deviation is given by

$$r_B(\tau_j) = -(B(\tau_j)\iota)'\boldsymbol{\sigma}_F\hat{\boldsymbol{\lambda}}_F, \tag{12}$$

$$s_B(\tau_j) = \sqrt{(B(\tau_j) * \iota)'\boldsymbol{\sigma}_F\boldsymbol{\sigma}_F'(B(\tau_j) * \iota)}. \tag{13}$$

Clearly, the mean excess bond returns depend on the discounting model through the equilibrium estimated prices-of-risk $\hat{\boldsymbol{\lambda}}_F$. The standard deviation of bonds is independent of the prices-of-risk. So, in equilibrium, our model does not allow for a *direct* match on the variability of excess bond returns. However, *indirectly* we do, because in equilibrium we account for the observed variability of bonds through the Kalman estimated bond standard deviation $\hat{\boldsymbol{\sigma}}_F$. For this reason, we study the Sharpe ratio — the mean of the excess return (12) divided by the standard deviation (13) — of bonds, providing us with a comprehensive measure of excess bond return behavior.

Empirically, excess bond returns have a positive relation with maturities up to 10 years (Boudoukh et al., 1999). Table 6, Panel A, confirms that present bias supports the data, while time consistency fails to fit the data. Indeed, excess bond returns are increasing in maturity, since the mean excess return on a 3-year bond is 1.90% and the mean excess return on a 10-year bond is 4.10%, during October 1976 - January 2019. The standard representative agent model is unable to match the data in terms of mean returns and Sharpe ratios. For a time-consistent investor the equilibrium Sharpe ratios of short-term and long-term bonds are simply too low compared to the data.

On the other hand, the present-biased investor yields a consistent fit with the data in terms of sign, magnitude and Sharpe ratio. In equilibrium, the mean 3-year bond return is 1.59% and the mean 10-year bond return is 4.45%. The short- and long-term bond Sharpe ratios are both similar to the data and, thus, present bias yields a large but simple improvement over time consistency. Our present-bias model underpredicts short-term bond returns by 0.31 and overpredicts long-term bond returns by 0.35.

Note that the fit with stock returns is close, such that the equity premium puzzle is absent in our model. Unlike Rudebusch and Swanson (2008) and Van Binsbergen et al. (2012) that find distorted predictions in the stock market, we find a reasonable price-of-risk for the stock λ_S.

Table 6, Panel B, shows that the estimated equilibrium prices-of-risk $\hat{\boldsymbol{\lambda}}_F$ are negative and statistically significantly estimated. So, we conclude that the equilibrium expected excess returns are not only economically significant, but also statistically.

2. Yield spread, or slope of the yield curve

The second measure is the yield spread, also known as the slope of the yield curve. The slope of the yield curve is simply the difference between the yield to maturity on a long-term bond and the risk-free rate. The yield vector for zero-coupon bonds is a linear transformation of the factors, where the intercept and factor loadings are time-invariant functions of time to maturity vector τ

$$\boldsymbol{Y}_t(\tau) \equiv -\ln \boldsymbol{P}_t(\tau)/\tau = -A(\tau)/\tau + \boldsymbol{\iota}'\boldsymbol{B}(\tau)\boldsymbol{F}_t/\tau, \tag{14}$$

where $\boldsymbol{P}_t(\tau)$ is the price vector of zero-coupon bonds with time-to-maturity vector τ, as given by (23) in Appendix A, and a function of the estimated equilibrium prices-of-risk vector $\hat{\boldsymbol{\lambda}}_F$.

The literature uses different proxies for long term bonds. Wachter (2006) and Van Binsbergen et al. (2012) uses a 5-year nominal bond and Rudebusch and Swanson (2008) use a 10-year nominal bond. For comparability, we show both results for 5 years and for 10 years bonds. The n-year yield spread at time t in our setting equals

$$y_t(n) - y_t(3 \text{ month}), \tag{15}$$

where $n = 5$ or 10 years, and where we proxy the risk-free rate by a 3 month yield.

Table 7 shows that present bias creates the observed, on average, upward sloping yield curve, while time consistency produces a too flat relationship between yields and maturity. The average difference in yields between the 5-year bond and the 3-month bond in our sample equals 1.33%, which indeed confirms an upward-sloping yield curve. Under time consistency, the equilibrium mean 5-year yield spread is too flat and not steep enough to match the data since the difference is 0.72 percentage points. Under present bias, the slope of the yield curve (1.03%) is close to our sample observation (1.33%) and the difference equals only 0.30 percentage points.

In the case of a longer long-term bond, our present-biased model performs even better. The observed 10-year yield spread is 1.78%. The time-consistent preferences produce an equilibrium 10-year yield spread of 1.04%, indicating a *downward* mismatch of 0.74 percentage points. Present-biased preferences produce an equilibrium 10-year yield spread of 1.88%, indicating an *upward* mismatch of only 0.10%. Again, our model by construction, does fit the standard deviations reasonably well in both discounting cases.

The main observation is that, indeed, equilibrium yield curves under exponential discounting provide an essentially flat relationship between yields and maturity, while the equilibrium term structure of interest rates under present bias provides the upward sloping positive relationship as observed in the data (Boudoukh et al., 1999). So, present bias is a potential explanation for why the term structure of interest rates is upward sloping.

3. Bond risk premia

Arguably the cleanest conceptual measure of long-term bond risk is the term premium, also known as the bond risk premium. The risk premium or term premium is typically expressed as the difference between the yield on the bond and the unobserved risk-neutral yield for that same bond (Rudebusch and Swanson, 2008). However, the term premium is not directly observed in the data and must be inferred using term-structure models (or other methods). The term premium in our setting equals

$$y_t(n) - \tilde{y}_t(n), \tag{16}$$

where $\tilde{y}_t(n)$ is the risk-neutral yield for an n-year bond at time t. We compute the risk-neutral yields at every time t by setting both the prices-of-risk for the stock λ_S and for the factors $\boldsymbol{\lambda}_F$ to zero.

As can be seen from Table 8, the average term premium under time consistency is 0.88% for a 5-year bond and 1.55% for a 10-year bond, respectively 0.52 and 0.57 percentage points smaller than the data. The simulated results are much closer to the data for present-biased preferences. For 5-year bonds the equilibrium term premium 1.32% is nearly identical to our sample 1.40%. The 10-year bond premium under present bias also matches the data, contrary to time consistency. By construction, our model is unable to fit the standard deviations and, therefore, we do not show these.

4. Campbell-Shiller (1991) "long-rate" slope coefficients

Our final and fourth measure of long-term bond risk is based on the "long-rate" regressions performed by Campbell and Shiller (1991) to test the hypothesis of constant risk premia on bonds, also known as the generalized expectations hypothesis. The regression equation is

$$y_{t+1}(n-1) - y_t(n) = \text{constant}_n + \alpha_n \frac{1}{n-1}(y_t(n) - y_t(3 \text{ month})) + \varepsilon_{n,t+1}, \tag{17}$$

where the dependent variable is the change in the n-period zero-coupon yield from period t to $t+1$, and the independent variable is the slope of the yield curve at time t divided by $n-1$. The intercept 'constant$_n$' and slope coefficient α_n are maturity specific. If bond risk premia are constant, and hence excess returns on long-term bonds are unpredictable, α_n should be equal to one. Instead, Campbell and Shiller (1991) find a coefficient that is negative at all maturities, and significantly different from one. Time variation in the term premium pushes α_n away from unity. Moreover, the higher the maturity, the lower is α_n.

Table 9 shows the coefficients α_n when we run the regression equation (17) on the sample, and under time consistency and present bias. The estimated coefficients from the sample are -1.87 and -3.46. The 10-year slope coefficient is similar to Rudebusch and Swanson (2008), who find a slope coefficient α_{10} of -3.49. We confirm the failure of the expectations hypothesis, because the coefficients are negative at both the 5-year and 10-year maturities, and are decreasing in maturity. This is equivalent to the statement that excess returns on long-term bonds are predictable (in sample).

The time consistency model goes a long way in explaining the failure of the expectations hypothesis, but the present bias model comes closer to the data. Under both discounting models, the signs of both slope coefficients are correct and negative. However, present-biased preferences provide more support for the data in terms of magnitudes.

III. Duration of the present

What is the duration of the present? This is an open empirical question in the behavioral literature by Ericson and Laibson (2018). Using a discounting model based on present bias requires distinguishing between "now" and later. Experimental evidence over consumption such as juice, water and effort finds a duration of the present ranging from a few minutes to a few weeks (McClure et al., 2007; Augenblick et al., 2015). However, a different picture emerges if we study structural models with annual periods. These models treat consumption anytime "this year" as immediate (Angeletos et al., 2001). The paper closest to ours is the myopic loss aversion paper Benartzi and Thaler (1995), who study investment behavior under myopic loss aversion. They report an investment evaluation period of 1 year for decision making. To the extent of our knowledge, there are no papers measuring the duration of the present for investment decisions and we try to fill this gap here.

We repeat our analysis from above, however we estimate equilibrium excess returns for a short-term planning horizon of $T_S = 3$ months and $T_S = 3$ years, rather than 1 year. So, for $T_S = 3$ months (years), the first consumption is during the present and coincides with $T_S = 3$ months (years), while the second consumption moment is during the future and coincides with $T_L = 10$ years. Since a short-term planning horizon of 3 months makes the investor by definition more myopic — the first consumption moment arrives earlier — we presume that we need a higher present-bias factor (i.e., more patience). A short-term planning horizon of 3 years makes the investor more far-sighted — the first consumption moment arrives later — we presume that we need a lower present-bias factor (i.e., more impatient). We keep the long-term annual discount factor still equal to $\delta = 0.97$.

Table 10 shows the equilibrium properties of the excess returns for the 3 different horizons. The key result is that a shorter duration of the present, i.e 3 months, leads to a more consistent fit with the data than a longer duration of the present, i.e 3 years. Specifically, a duration of 3 months with $\beta = 0.7$ provides even a better fit with the data than our benchmark duration of 1 year with $\beta = 0.35$. If the duration of the present equals 3 years, then we are unable to solve the bond premium puzzle even though we adopt an unreasonable low present-bias parameter $\beta = 0.05$. The value of the present-bias factor $\beta = 0.7$ is in line with estimated present-bias factors in the experimental literature. $\beta = 0.05$ seems unreasonable, but Laibson et al. (2015) do estimate such low values of present bias for a sophisticated agent.

For this reason, we conclude that the present has a duration of at most 1 year, since for reasonable preference specifications we can match the excess bond returns.[11] This value is in line with the finding of Benartzi and Thaler (1995). They are successful in explaining the

[11] We do not report yield spreads, risk premia and "long-rate" slope coefficients, because these measures lead to similar conclusions.

equity premium puzzle by means of an evaluation period of 1 year. It may not come as a surprise, and actually be very plausible, that investors use the same duration in stock and bond markets. For this reason, we conclude that on average investors use a duration of the present no longer than 1 year.

A. Do organizations display this behavior?

A possible objection to our model might be that it is based on an individual representative investor, while the majority of assets we are concerned with are held by organizations such as pension funds (Benartzi and Thaler, 1995). We argue that a principal-agent problem due to decentralized investment might be an interpretation for our 2-period present-biased model, which groups behavior in the present and in the future.

Pension funds have essentially an infinite time horizon, as our representative investor in the general equilibrium. At each time t, the pension fund has to decide how to allocate the pot of wealth over the 2 'consumption' moments at dates $T_S = 1$ year and $T_L = 10$ years. While the pension fund is indeed likely to exist as long as the company remains in business, the pension fund manager — for example a Chief Investment Officer (CIO) — does not expect to be in this job forever. The investment horizon for a CIO is typically somewhat long and might equal a long-term planning horizon up to 10 years (Van Binsbergen et al., 2008).

But, the CIO typically delegates portfolio decisions to his asset managers, who are usually compensated on an annual basis. So, the investment horizons of the CIO and the asset managers differ. Due to the annual compensation, the investment horizons of asset managers are generally short and run no further than 1 year. Gains and losses after the 1-year duration of the present matter less for the asset managers than for the CIO. The CIO has the incentive to invest for a longer term of 10 years. Therefore, the infinitely-lived pension fund is subject to a long-term planning horizon of 10 years — by means of the incentives of the CIO — while the fund is directly subject to present bias — by means of the incentives of the asset managers.

In conclusion, a principal-agent problem due to decentralized investment might produce the present-biased behavior with a present duration of at most 1 year, and a long-term planning horizon of 10 years. The duration of the present has a similar interpretation and identical magnitude as the evaluation period of 1 year reported by Benartzi and Thaler (1995). Their evaluation period, similar to our duration of the present, should not, in any way, be confused with the long-term planning horizon of the investor.

B. Impatient time-consistent investor

One might wonder which value of the long-term discount factor δ in the exponential discounting model is required to explain observed bond characteristics. In other words, can we find a time-consistent investor, who requires only one parameter δ, that supports the data? This is a fair question, and we address it here.

We know from the empirical optimal asset allocation that a time-consistent investor with $\beta = 1, \delta = 0.86$ replicates the hedge and speculative demands of a present-biased investor with $\beta = 0.35, \delta = 0.97$ for a 1-year duration of the present. We repeat our equilibrium analysis, regarding excess bond returns, for an impatient time-consistent investor with a discount factor of 0.86. The value $\delta = 0.86$ implies severe impatience, because the time-consistent investor uses an unreasonable discount rate of approximately 16% per year. Typically, financial economists find such a value too low, since they tend to use annual discount factor close to 1. On the other hand, the experimental literature finds the factor $\delta = 0.86$ too high in the exponential discounting model, because their discount factor estimates range from $\delta = 0.23$ to $\delta = 0.63$.

Table 11 shows the excess returns in the data and in equilibrium, for an impatient time-consistent investor ($\beta = 1, \delta = 0.86$) and the present-biased investor ($\beta = 0.35, \delta = 0.97$) with $T_S = 1$ year. The impatient time-consistent investor replicates the equilibrium findings of the present-biased investor nearly identically, such that an impatient time-consistent investor supports the data as well. Also the factor prices-of-risk $\boldsymbol{\lambda}_F$ are closely replicated.

So, if one is willing to assume unreasonable discount factors in the exponential discounting model, then bond characteristics can be explained as well. However, it appears from our model that it is necessary to introduce a distinction between the present and the future by means of present bias, for example. Our finding is in line with many experimental evidence that people are subject to present bias.

IV. Conclusion

In this paper, we offer an explanation to the bond premium puzzle based on the time preferences of a representative investor. Our solution to the puzzle is a simple introduction of present-biased preferences into the standard exponential expected utility framework (Phelps and Pollak, 1968; Laibson, 1997). We formulate a dynamic investment problem with a general discount function in a market for bonds and stocks, where interest rates are driven by a 2-factor Gaussian affine term structure model (such as Vasicek, 1977).

We give explicit analytical solutions by using the martingale method (Cox and Huang, 1989) for the representative investor's optimal consumption path and optimal portfolio choice for general deterministic time preferences, including present bias. These optimal solutions imply equilibrium bond returns and yields, which show that a present-biased investor — a special case of our discounting model — is able to explain many characteristics of bond behavior using preference parameters as previously estimated in the experimental literature. Specifically, present-biased behavior explains expected excess bond returns, the yield spread, bond risk premia and long-rate slope coefficients while time-consistent behavior is unable to do so.

In summary, a substantial body of experimental literature suggests that present bias plays an important role in making intertemporal decisions (Frederick et al., 2002). Present-biased investors, modelled by quasi-hyperbolic discounting, value the present more than the future and, therefore, care less about hedging opportunities for the long run. For this reason, present-biased investors have a lower long-term hedge demand and demand a premium to hold long-term bonds. As a result, the term structure of interest rates is on average upward sloping and matches the data. We find a realistic premium if the present is at most 1 year, providing an indication for the duration of the present.

Appendix

A. Proof of Theorem 1

The result is fairly standard and present for completeness only. We solve firstly for the optimal consumption path and, then, for the optimal investment strategy. The representative investor maximizes expected CRRA utility of wealth at a terminal horizon. He solves (7) subject to his budget constraint (8).

Standard calculations using the Lagrange method, lead to implicit optimal consumption

$$W_t^* = \frac{1}{M_t} E_t \left[W_T^* M_T \right] = \frac{W_0}{M_t} \frac{E_t \left[M_T^{1-1/\gamma} \right]}{E_0 \left[M_T^{1-1/\gamma} \right]}, \tag{18}$$

where we use the first fundamental theorem of asset pricing.

Lemma 1. *For positive α, the conditional expectation at time t of the stochastic discount factor follows from*

$$E_t \left[\left(\frac{M_T}{M_t} \right)^{\alpha} \right] = \exp \left(\alpha m(\boldsymbol{F}_t, T-t) + \frac{1}{2} \alpha^2 v^2 (T-t) \right) \tag{19}$$

$$= e^{\frac{1}{2}(\alpha-1)\alpha v^2 (T-t)} P_t (T-t)^{\alpha} \tag{20}$$

where ($\tau = T - t$)

$$m(\boldsymbol{F}_t, \tau) = -\left(A_0 + \boldsymbol{\iota}' \boldsymbol{\theta} \right)(\tau) - \boldsymbol{\iota}' \boldsymbol{B}(\tau) \left(\boldsymbol{F}_t - \boldsymbol{\theta} \right) - \frac{1}{2} \boldsymbol{\lambda}' \boldsymbol{\lambda}(\tau) \tag{21}$$

$$v^2(\tau) = \int_0^{\tau} ||\boldsymbol{\iota}' \boldsymbol{B}(\tau - v) \boldsymbol{\sigma}_F||^2 dv + \boldsymbol{\lambda}' \boldsymbol{\lambda}(\tau) + 2 \int_0^{\tau} \langle \boldsymbol{\iota}' \boldsymbol{B}(\tau - v) \boldsymbol{\sigma}_F, \boldsymbol{\lambda}_F \rangle dv, \tag{22}$$

and

$$P_t(T-t) = \exp \left(A(T-t) - \boldsymbol{\iota}' \boldsymbol{B}(T-t) \boldsymbol{F}_t \right) \tag{23}$$

with deterministic functions

$$
\begin{aligned}
A(T-t) = & -\left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta}\right)(T-t) + \boldsymbol{\iota}'\boldsymbol{B}(T-t)\boldsymbol{\theta} \\
& + \frac{1}{2}\int_t^T ||\boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F||^2 dv + \int_t^T \langle \boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F, \boldsymbol{\lambda}_F' \rangle dv
\end{aligned} \tag{24}
$$

Appendix provides the proof of Lemma 1. Now, the explicit optimal consumption path follows directly from (18) by plugging in the result of Lemma 1. Mathematically,

$$
\begin{aligned}
W_t^* = & W_0 \exp\left(-\int_0^t r_s ds - \int_0^t \boldsymbol{\lambda}' d\boldsymbol{Z}_s - \frac{1}{2}\int_0^t \boldsymbol{\lambda}'\boldsymbol{\lambda} ds\right)^{\alpha-1} \\
& \times \exp\left(\frac{1}{2}(\alpha-1)\alpha v^2(T-t) - \frac{1}{2}(\alpha-1)\alpha v^2(T)\right) \\
& \times \{\exp[-\left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta}\right)(T-t) - \boldsymbol{\iota}'\boldsymbol{B}(T-t)\left(\boldsymbol{F}_t - \boldsymbol{\theta}\right) \\
& + \frac{1}{2}\int_t^T ||\boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F||^2 dv + \int_t^T \langle \boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F, \boldsymbol{\lambda}_F' \rangle dv]^{\alpha} \\
& - \exp[-\left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta}\right)(T) - \boldsymbol{\iota}'\boldsymbol{B}(T)\left(\boldsymbol{F}_0 - \boldsymbol{\theta}\right) \\
& + \frac{1}{2}\int_0^T ||\boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F||^2 dv + \int_0^T \langle \boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F, \boldsymbol{\lambda}_F' \rangle dv]^{\alpha}\}.
\end{aligned} \tag{25}
$$

Note that the stochasticity in (25) comes from

$$
-(\alpha-1)\left(\boldsymbol{\lambda}_F'\boldsymbol{Z}_{F,t} + \lambda_S Z_{S,t}\right)
$$

and from

$$
-\alpha\boldsymbol{\iota}'\boldsymbol{B}(T-t)\boldsymbol{F}_t.
$$

So, rewriting (25) into a stochastic differential equation, with $\alpha = 1 - 1/\gamma$ due to CRRA utility, yields

$$
d\log W_t^* = g_1(r, T-t)dt + \frac{\lambda_S}{\gamma} dZ_{S,t} + \left(\frac{\boldsymbol{\lambda}_F'}{\gamma} - (1-1/\gamma)\boldsymbol{\iota}'\boldsymbol{B}(T-t)\boldsymbol{\sigma}_F\right) d\boldsymbol{Z}_{F,t}, \tag{26}
$$

where the drift term $g_1(r, T-t)$ is a function of the interest rate r_t and the remaining investment horizon $\tau = T - t$.

Finally, consider the return on a portfolio that consists of the three available assets. Let $\boldsymbol{\pi}^* = (\pi_S^*(t), \boldsymbol{\pi}_{\boldsymbol{B}}^*(t,\tau), \pi_M^*(t))$ be the optimal proportion of wealth invested at time t in a stock, in a vector of bonds with maturities τ and in cash. If the individual invests in such a portfolio, then total wealth $A(t)$ evolves according to

$$\frac{dA_t}{A_t} = \left(\pi_S(t)\frac{dS_t}{S_t} + \boldsymbol{\pi}_{\boldsymbol{B}}(t,\tau)'\frac{d\boldsymbol{P}_t}{\boldsymbol{P}_t} + \pi_M(t)\frac{dB_t}{B_t}\right), \tag{27}$$

where B_t is the process for the risk-free asset, or cash. Substituting the dynamics of the assets and taking the log, yields

$$d\log A_t = g_2(.)dt + \pi_S(t)\sigma_S dZ_{S,t} + \left(\pi_S(t)\boldsymbol{\sigma}'_{FS} - \boldsymbol{\pi}_{\boldsymbol{B}}(t,\tau)'\mathcal{B}(\tau)'\boldsymbol{\sigma}_F\right) d\boldsymbol{Z}_{F,t}, \tag{28}$$

where $g_2(.)$ is the drift term.

Then, the optimal investment demands $\boldsymbol{\pi}^*$ follow by simply equating the coefficients of the diffusion terms in (26) and (28). □

B. Proof of Lemma 1

Using (2) and (3), we have

$$\begin{aligned} r_t &= A_0 + \boldsymbol{\iota}'\left[\boldsymbol{\theta} + \exp(-\boldsymbol{\kappa}t)(\boldsymbol{F}_0 - \boldsymbol{\theta})\right] + \int_0^t \boldsymbol{\iota}'\exp(-\boldsymbol{\kappa}(t-s))\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,s} \\ &= A_0 + \boldsymbol{\iota}'\boldsymbol{\theta} + \boldsymbol{\iota}'\exp(-\boldsymbol{\kappa}t)(\boldsymbol{F}_0 - \boldsymbol{\theta}) + \int_0^t \boldsymbol{\iota}'\exp(-\boldsymbol{\kappa}(t-s))\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,s}, \end{aligned} \tag{29}$$

which follows directly from the solution of the stochastic differential equation to the Orstein-Uhlenbeck process (see, e.g., Chin et al. (2014)).

From the above, we find

$$\int_0^t r_v dv = \int_0^t \left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta} + \boldsymbol{\iota}' \exp(-\boldsymbol{\kappa} v)(\boldsymbol{F}_0 - \boldsymbol{\theta}) + \int_0^v \boldsymbol{\iota}' \exp(-\boldsymbol{\kappa}(v-s))\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,s} \right) dv$$
$$= (A_0 + \boldsymbol{\iota}'\boldsymbol{\theta})\, t + \boldsymbol{\iota}'\boldsymbol{B}(t)\,(\boldsymbol{F}_0 - \boldsymbol{\theta}) + \int_0^t \left(\int_0^v \boldsymbol{\iota}' \exp(-\boldsymbol{\kappa}(v-s))\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,s} \right) dv$$
$$= (A_0 + \boldsymbol{\iota}'\boldsymbol{\theta})\, t + \boldsymbol{\iota}'\boldsymbol{B}(t)\,(\boldsymbol{F}_0 - \boldsymbol{\theta}) + \int_0^t \boldsymbol{\iota}'\boldsymbol{B}(t-v)\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,v} \tag{30}$$

where in the second equality we use (6) and in the third equality we use stochastic integration by parts. More general

$$\int_t^T r_v dv = (A_0 + \boldsymbol{\iota}'\boldsymbol{\theta})\,(T-t) + \boldsymbol{\iota}'\boldsymbol{B}(T-t)\,(\boldsymbol{F}_t - \boldsymbol{\theta}) + \int_t^T \boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,v}, \tag{31}$$

which has a normal distribution with (conditional) mean and variance:

$$E_t\left[\int_s^T r_v dv\right] = (A_0 + \boldsymbol{\iota}'\boldsymbol{\theta})\,(T-s) + \boldsymbol{\iota}'\boldsymbol{B}(T-s)\,(\boldsymbol{F}_s - \boldsymbol{\theta})\,, \tag{32}$$
$$V_t\left[\int_s^T r_v dv\right] = \int_t^T ||\boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F||^2 dv, \tag{33}$$

where we use Itô isometry to obtain the variance.[12] Hence, the stochastic discount factor follows a log-normal distribution with mean

$$m(r_t, T-t) = E_t\left[\ln \frac{M_T}{M_t}\right]$$
$$= E_t\left[-\int_s^T r_v dv - \int_s^T \boldsymbol{\lambda}' d\boldsymbol{Z}_v - \frac{1}{2}\int_s^T \boldsymbol{\lambda}'\boldsymbol{\lambda} dv\right]$$
$$= -\,(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta})\,(T-s) - \boldsymbol{\iota}'\boldsymbol{B}(T-s)\,(\boldsymbol{F}_s - \boldsymbol{\theta}) - \frac{1}{2}\boldsymbol{\lambda}'\boldsymbol{\lambda}(T-s), \tag{34}$$

[12] $||.||$ denotes the Euclidean norm of a vector (i.e., the square root of inner product of the vector and itself).

and variance

$$
\begin{aligned}
v^2(T-t) &= V_t\left[\ln \frac{M_T}{M_t}\right] \\
&= V_t\left[-\int_s^T r_v dv - \int_s^T \boldsymbol{\lambda}' d\boldsymbol{Z}_v - \frac{1}{2}\int_s^T \boldsymbol{\lambda}'\boldsymbol{\lambda} dv\right] \\
&= \int_s^T ||\boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F||^2 dv + \boldsymbol{\lambda}'\boldsymbol{\lambda}(T-s) + 2E_t\left[\int_s^T r_v dv \int_s^T \boldsymbol{\lambda}_F' d\boldsymbol{Z}_{F,v}\right] \\
&= \int_s^T ||\boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F||^2 dv + \boldsymbol{\lambda}'\boldsymbol{\lambda}(T-s) + 2\int_s^T \langle \boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F, \boldsymbol{\lambda}_F'\rangle dv. \qquad (35)
\end{aligned}
$$

We use that the prices-of-risk are constant, Z_S is independent of $\boldsymbol{Z}_F$ and Itô isometry. $\square$

C. Proof of Theorem 2

We first solve for the optimal consumption path and, then, for the optimal investment strategy. The representative investor maximizes CRRA utility of intermediate consumption. He solves (9) subject to his budget constraint (10).

Using η as Lagrange multiplier for the budget constraint, we find

$$
\mathcal{L} = \mathbb{E}_t\left[\sum_{j=1,t\leq T_j}^{n} D(T_j-t)\frac{W_{t,T_j,T_j}^{1-\gamma}}{1-\gamma}\right] + \eta\left(W_t M_t - \mathbb{E}_t\left[\sum_{j=1,t\leq T_j}^{n} W_{t,T_j,T_j} M_{T_j}\right]\right). \qquad (36)
$$

Taking the first order conditions with respect to terminal wealth for every state of the world, yields

$$
\frac{\partial \mathcal{L}}{\partial W_{t,T_j,T_j}} = D(T_j-t)W_{t,T_j,T_j}^{-\gamma} - \eta M_{T_j} = 0. \qquad (37)
$$

This implies *implicit* optimal *terminal* wealth

$$
W_{t,T_j,T_j}^{*} = D^{1/\gamma}(T_j-t)\left(\eta M_{T_j}\right)^{-1/\gamma}. \qquad (38)
$$

and isolating the Lagrange multiplier with the budget constraint yields

$$\eta^{-1/\gamma} = \frac{W_t M_t}{\mathbb{E}_t\left[\sum_{j=1,t\leq T_j}^{n} D^{1/\gamma}(T_j - t) M_{T_j}^{1-1/\gamma}\right]}. \tag{39}$$

Substituting the expression for the Lagrange multiplier in implicit optimal terminal wealth yields *explicit* optimal *terminal* wealth

$$W^*_{t,T_j,T_j} = D^{1/\gamma}(T_j - t) M_t M_{T_j}^{-1/\gamma} \frac{W_t}{\sum_{j=1,t\leq T_j}^{n} D^{1/\gamma}(T_j - t)\mathbb{E}_t\left[M_{T_j}^{1-1/\gamma}\right]}. \tag{40}$$

Now, using the first the fundamental theorem of asset pricing, the *actual* optimal consumption *path* at time t with horizon $T_j \geq t$ for in-between dates $t + h = t, ..., T_j$ for every money pot $j = 1, ..., n$ equals

$$\begin{aligned}
&W^*_{t,T_j,t+h}\left(D(T_j - t)\right) = \frac{1}{M_{t+h}}\mathbb{E}_{t+h}\left[W^*_{t,T_j,T_j-t} M_{T_j}\right] \\
&= D(T_j - t)^{1/\gamma} W_t \frac{M_t}{M_{t+h}} \frac{\mathbb{E}_{t+h}\left[M_{T_j}^{1-1/\gamma}\right]}{\sum_{j=1,t\leq T_j}^{n} D(T_j - t)^{1/\gamma}\mathbb{E}_t\left[M_{T_j}^{1-1/\gamma}\right]} \\
&= D(T_j - t)^{1/\gamma} W_t M_t M_{t+h}^{-1/\gamma} \frac{\exp\left(\frac{1}{2}\frac{1-\gamma}{\gamma^2} v^2 (T_j - (t+h))\right) P_{t+h}(T_j - (t+h))^{1-1/\gamma}}{\sum_{j=1,t\leq T_j}^{n} D(T_j - t)^{1/\gamma}\mathbb{E}_t\left[M_{T_j}^{1-1/\gamma}\right]} \\
&= D(T_j - t)^{1/\gamma} W_t M_t^{1-1/\gamma} \exp\left(-\int_t^{t+h} r_s ds - \boldsymbol{\lambda}'\left(\boldsymbol{Z}_{t+h} - \boldsymbol{Z}_t\right) - \frac{1}{2}\boldsymbol{\lambda}'\boldsymbol{\lambda}((t+h) - t)\right) \\
&\cdot \frac{\exp\left(\frac{1}{2}\frac{1-\gamma}{\gamma^2} v^2 (T_j - (t+h))\right) P_{t+h}(T_j - (t+h))^{1-1/\gamma}}{\sum_{j=1,t\leq T_j}^{n} D(T_j - t)^{1/\gamma}\mathbb{E}_t\left[M_{T_j}^{1-1/\gamma}\right]},
\end{aligned} \tag{41}$$

where we use Lemma 1 in the third equation, and in the fourth equation we use the explicit expression of the stochastic discount factor.

The optimal investment strategy $\pi_i^*(t, T_j)$ for asset $i = \{\text{stock, constant maturity } \tau\text{-year bonds, cash}\}$ at time t for investment horizon T_j follows from Theorem 1. So, the optimal fraction of actual total invested wealth at time t for each money pot j follows from the investor's *planned* optimal consumption path $W^*_{t,T_j,t+h}$ with $h = 0$, leading to the investor's *actual* optimal consumption path $W^*_{t,T_j,t}$. Both consumption paths are a function of the discount structure $D(T_j - t)$. This proves Theorem 2.[13] $\square$

[13] Note that $j > t$ because investments are for future dates, i.e., the investor does not invest for a payment at a payment date

References

Ainslie, G. (1992). *Picoeconomics: The Strategic Interaction of Successive Motivational States within the Person.* Cambridge, UK: Cambridge University Press.

Angeletos, G., D. Laibson, A. Repetto, J. Tobacman, and S. Weinberg (2001). “The hyperbolic consumption model: Calibration, simulation and empirical evaluation”. In: *Journal of Economic Perspectives* 15.3, pp. 47–68.

Augenblick, N., M. Niederle, and C. Sprenger (2015). “Working over Time: Dynamic Inconsistency in Real Effort Tasks”. In: *Quarterly Journal of Economics*, pp. 1067–1115.

Augenblick, N. and M. Rabin (2019). “An Experiment on Time Preference and Misprediction in Unpleasant Tasks”. In: *Review of Economic Studies* 86, pp. 941–975.

Backus, D., A. Gregory, and S. Zin (1989). “Risk Premiums in the Term Structure: Evidence from Artificial Economies”. In: *Journal of Monetary Economics* 24, pp. 371–399.

Balakrishnan, U., J. Haushofer, and P. Jakiela (2017). “How Soon Is Now? Evidence of Present Bias from Convex Time Budget Experiments”. In: *NBER Working Paper of Labour Economics*.

Bansal, R. and W. Coleman (1996). “A Monetary Explanation of the Equity Premium, Term Premium, and Risk-Free Rate Puzzles”. In: *Journal of Political Economy* 104.6, pp. 1135–1171.

Barberis, N. and M. Huang (2001). “Mental Accounting, Loss Aversion, and Individual Stock Returns”. In: *Journal of Finance* 56.4, pp. 1247–1292.

Benartzi, S. and R. Thaler (1995). “Myopic Loss Aversion and the Equity Premium Puzzle”. In: *Quarterly Journal of Economics* 110.1, pp. 73–92.

Boudoukh, J., M. Richardson, T. Smith, and R. Whitelaw (1999). “Ex Ante Bond Returns and the Liquidity Preference Hypothesis”. In: *Journal of Finance* 104.54, pp. 1153–1167.

Brennan, M. and Y. Xia (2002). “Dynamic Asset Allocation under Inflation”. In: *Journal of Finance* 57.3, pp. 1201–1238.

Campbell, J. and J. Cocco (2003). “Household Risk Management and Optimal Mortgage Choice”. In: *Quarterly Journal of Economics* 118.4, pp. 1449–1494.

Campbell, J. and R. Shiller (1991). “Yield Spreads and Interest Rate Movement: A Bird’s Eye View”. In: *Review of Economic Studies* 58, pp. 495–514.

Chin, E., D. Nel, and S. Ólafsson (2014). *Problems and Solutions in Mathematical Finance. Volume 1: Stochastic Calculus*. Chichester: John Wiley & Sons, Ltd.

Cox, J. and C. Huang (1989). "Optimal Consumption and Portfolio Policies when Asset Prices Follow a Diffusion Process". In: *Journal of Economic Theory* 49, pp. 33–83.

De Jong, F. (2000). "Time Series and Cross-section Information in Affine Term-Structure Models". In: *Journal of Business & Economic Statistics* 18.3, pp. 300–314.

DellaVigna, S. and M. Paserman (2005). "Job Search and Impatience". In: *Journal of Labour Economics* 23.3, pp. 527–588.

Ericson, K. and D. Laibson (2018). "Intertemporal choice". In: *NBER working paper series*.

Fang, H. and D. Silverman (2009). "Time-consistency and Welfare Program Participation: Evidence from the NLSY". In: *International Economic Review* 50.4, pp. 1043–1077.

Frederick, S., G. Loewenstein, and T. O'Donoghue (2002). "Time Discounting and Time Preference: A Critical Review". In: *Journal of Economic Literature* 40, pp. 351–401.

Geyer, A. and S. Pichler (1999). "A State-Space Approach to Estimate and Test Multifactor Cox-Ingersoll-Ross models of the Term Structure". In: *Journal of Financial Research* 22.1, pp. 107–130.

Harris, C. and D. Laibson (2013). "Instantaneous Gratification". In: *Quarterly Journal of Economics* 128.1, pp. 205–248.

Horvath, F., F. De Jong, and B. Werker (2017). "Robust Pricing of Fixed Income Securities". In: *Netspar Working Paper*.

Laibson, D. (1997). "Golden Eggs and Hyperbolic Discounting". In: *Quarterly Journal of Economics* 112.2, pp. 443–477.

Laibson, D., P. Maxted, and A. Repetto (Oct. 2015). "Estimating Discount Functions with Consumption Choices over the Lifecycle". Mimeo, Harvard University.

Loewenstein, G. and D. Prelec (1992). "Anomalies in Intertemporal Choice: Evidence and an Interpretation". In: *Quarterly Journal of Economics* 107.2, pp. 573–597.

McClure, S., K. Ericson, D. Laibson, G. Loewenstein, and J. Cohen (2007). "Time Discounting for Primary Reward". In: *Journal of Neuroscience* 27.21, pp. 5796–5804.

Mehra, R. and E. Prescott (1985). "The Equity Premium: A Puzzle". In: *Journal of Monetary Economics* 15, pp. 145–161.

Merton, R. (1969). "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case". In: *Review of Economics and Statistics* 51.3, pp. 247–257.

O'Donoghue, T. and M. Rabin (1999). "Doing It Now or Later". In: *American Economic Review* 98.1, pp. 103–124.

— (2015). "Present Bias: Lessons Learned and To Be Learned". In: *American Economic Review* 105.5, pp. 273–279.

Paserman, M. (2008). "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation". In: *Economic Journal* 118, pp. 1418–1452.

Phelps, E. and R. Pollak (1968). "On Second-Best National Saving and Game-Equilibrium Growth". In: *Review of Economic Studies* 35.2, pp. 185–199.

Rudebusch, G. and E. Swanson (2008). "Examining the bond premium puzzle with a DSGE mode". In: *Journal of Monetary Economics* 55, S111–S126.

Samuelson, P. (1937). "A Note on Measurement of Utility". In: *Review of Economic Studies* 4, pp. 155–161.

Shapiro, J. (2005). "Is there a daily discount rate? Evidence fromthe food stamp nutrition cycle". In: *Journal of Public Economics* 89.2-3, pp. 303–325.

Strotz, R. (1956). "Myopia and Inconsistency in Dynamic Utility Maximisation". In: *Review of Economic Studies* 23.3, pp. 165–180.

Thaler, R. (1981). "Some empirical evidence on dynamic inconsistency". In: *Economics Letters* 8.3, pp. 201–207.

Van Binsbergen, J., M. Brandt, and R. Koijen (2008). "Optimal Decentralized Investment Management". In: *Journal of Finance* 63.4, pp. 1849–1895.

Van Binsbergen, J., J. Fernández-Villaverde, R. Koijen, and J. Rubio-Ramírez (2012). "The termstructure of interest rates in a DSGE model with recursive preferences". In: *Journal of Monetary Economics* 59, pp. 634–648.

Vasicek, O. (1977). "An equilibrium characterization of the term structure". In: *Journal of Financial Economics* 5, pp. 177–188.

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

Zou, Z., S. Chen, and L. Wedge (2014). "Finite horizon consumption and portfolio decisions with stochastic hyperbolic discounting". In: *Journal of Mathematical Economics* 52, pp. 70–80.

Tables

Table 1: **Descriptive Statistics of Yields and Stock Returns.** The data consists of monthly U.S. interest rates and stock returns from 1 October 1976 to 1 January 2019. All values are annualized.

	Mean	Standard deviation	Minimum	Maximum
3-month yield	4.52	3.59	0.01	16.30
1-year yield	5.03	3.83	0.10	16.72
5-year yield	5.85	3.50	0.62	15.93
10-year yield	6.29	3.23	1.50	15.32
Stock return	11.97	15.18	-	-

Table 2: **U.S. Government Debt by Maturity.** Panel A of this table presents the composition of U.S. government debt at 1 October 1976. Panel B of this table presents the composition of U.S. government debt at 1 January 2019. Panel C shows the mean debt composition during the observation period 1 October 1976 to 1 January 2019.

	Debt outstanding (million USD)	Fraction of total debt outstanding
Panel A: Composition at 1 October 1976		
Total	294,595	1.00
< 1 year	153,302	0.52
1-5 years	94,845	0.32
5-10 years	31,247	0.11
10-20 years	7,939	0.03
> 20 years	7,262	0.02
Panel B: Composition at 1 January 2019		
	Debt outstanding (million USD)	Fraction of total debt outstanding
Total	13,385,359	1.00
< 1 year	3,927,279	0.29
1-5 years	5,426,079	0.41
5-10 years	2,524,238	0.19
10-20 years	113,097	0.01
> 20 years	1,394,666	0.10
Panel C: Composition from 1 October 1976 to 1 January 2019		
	Debt outstanding (million USD)	Fraction of total debt outstanding
Total	3,798,442	1.00
< 1 year	1,263,723	0.37
1-5 years	1,460,573	0.36
5-10 years	623,445	0.14
10-20 years	147,308	0.05
> 20 years	303,393	0.08

Table 3: **Estimates of Model Parameters.** This table reports maximum likelihood estimates and standard errors for the parameters of the financial market with monthly observations. At each time, we observed 4 points on the U.S. zero-coupon yield curve, corresponding with maturities of 3 months, 1 year, 5 years and 10 years. And, we observed the (value-weighted) stock market index. The observation period is 1 October 1976 to 1 January 2019

Parameter	Estimate	Standard error
$\hat{\kappa}_1$	0.0398	0.0015
$\hat{\kappa}_2$	0.4623	0.0119
$\hat{A}_0$	0.0359	0.0015
$\hat{\lambda}_{F,1}$	-0.1378	0.0044
$\hat{\lambda}_{F,2}$	-0.4785	0.0383
$\hat{\lambda}_S$	0.4095	0.1381
$\hat{\sigma}_{F,11}$	0.0139	0.0004
$\hat{\sigma}_{F,21}$	-0.0065	0.0007
$\hat{\sigma}_{F,22}$	0.0166	0.0005
$\hat{\sigma}_{FS,1}$	-0.0264	0.0259
$\hat{\sigma}_{FS,2}$	-0.0171	0.0101
$\hat{\sigma}_S$	0.1504	0.0020

Table 4: **Optimal asset allocation under present bias and time consistency.** This table reports the the optimal fraction of total wealth invested in a 3-year bond, a 10-year bond, a stock and cash. Panel A presents the optimal demands for the present-biased investor with preferences $\beta = 0.35, \delta = 0.97$. Panel B presents the optimal demands for the time-consistent investor with preferences $\beta = 1, \delta = 0.97$.

	3-year bond	10-year bond	Stock	Cash
Panel A: Present-biased investor				
Hedge demand	0.48	0.04		
Speculative demand	2.44	-0.63	0.27	
Total demand	2.92	-0.59	0.27	-1.61
Panel B: Time-consistent investor				
Hedge demand	0.39	0.21		
Speculative demand	2.44	-0.63	0.27	
Total demand	2.83	-0.42	0.27	-1.68

Table 5: **Sensitivity optimal asset allocation under present bias and time consistency.** For several discount models, this table reports the optimal fraction of total wealth invested in a 3-year bond, a 10-year bond, a stock and cash. Panel A presents the optimal demands for the present-biased investor with varying present-bias factors β. Panel B presents the optimal demands for the time-consistent investor.

	3-year bond	10-year bond	Stock	Cash
	Panel A: Present-biased investor			
$\beta = 0.05, \delta = 0.97$				
Hedge demand	0.54	-0.06	-	-
Speculative demand	2.44	-0.63	0.27	-
Total demand	2.98	-0.69	0.27	-1.56
$\beta = 0.5, \delta = 0.97$				
Hedge demand	0.46	0.09	-	-
Speculative demand	2.44	-0.63	0.27	-
Total demand	2.90	-0.54	0.27	-1.63
$\beta = 0.8, \delta = 0.97$				
Hedge demand	0.42	0.16	-	-
Speculative demand	2.44	-0.63	0.27	-
Total demand	2.85	-0.47	0.27	-1.66
	Panel B: Time-consistent investor			
$\beta = 1, \delta = 0.86$				
Hedge demand	0.49	0.04	-	-
Speculative demand	2.44	-0.63	0.27	-
Total demand	2.92	-0.59	0.27	-1.60

Table 6: **Properties of asset returns in the data and in the model.** Data runs from October 1976 to January 2019. The time-consistent investor has preferences $\beta = 1, \delta = 0.97$, and the present-biased investor has preferences $\beta = 0.35, \delta = 0.97$. Panel A presents 3-year bond, 10-year bond and stock returns per annum, in excess of the risk-free rate. The Sharpe ratio is the mean of the excess return divided by the standard deviation. Panel B reports the estimated coefficients and standard errors for the equilibrium prices-of-risk $\hat{\boldsymbol{\lambda}}_F$.

	Data	Time consistency	Present bias
	Panel A: Excess returns		
3-year bond			
Mean	1.90	1.06	1.59
Sharpe	0.48	0.27	0.41
10-year bond			
Mean	4.10	2.79	4.45
Sharpe	0.38	0.26	0.42
Stock			
Mean	7.27	7.01	7.48
Sharpe	0.48	0.46	0.49
	Panel B: Prices-of-risk		
Factor 1: $\hat{\lambda}_{F,1}$			
Coefficient	-0.1304	-0.2221	-0.3741
Standard error	0.0039	0.0480	0.0335
Factor 2: $\hat{\lambda}_{F,2}$			
Coefficient	-0.5059	-0.1578	-0.1945
Standard error	0.0288	0.0109	0.0056

Table 7: **Yield spreads in the data and in the model.** This table reports the mean and standard deviation of the yield spread, also known as the slope of the yield curve. The yield spread is the difference in yields between the long-term n-year bond and the 3-month bond. Data runs from October 1976 to January 2019. The time-consistent investor has preferences $\beta = 1, \delta = 0.97$, and the present-biased investor has preferences $\beta = 0.35, \delta = 0.97$.

Maturity n	Data	Time consistency	Present bias
5 years			
Mean	1.33	0.61	1.03
Standard deviation	0.97	0.96	0.97
10 years			
Mean	1.78	1.04	1.88
Standard deviation	1.22	1.28	1.30

Table 8: **Bond risk premia in the data and in the model.** This table reports the mean bond risk premium. The bond risk premium, or term premium, is the difference between the yield on the long-term n-year bond and the unobserved risk-neutral yield for that same bond. Data runs from October 1976 to January 2019. The time-consistent investor has preferences $\beta = 1, \delta = 0.97$, and the present-biased investor has preferences $\beta = 0.35, \delta = 0.97$.

Maturity n	Data	Time consistency	Present bias
5 years			
Mean	1.40	0.88	1.32
10 years			
Mean	2.12	1.55	2.41

Table 9: **"Long-rate" regressions in the data and in the model.** Coefficients α_n from the regression $y_{t+1}(n-1) - y_t(n) = \text{constant}_n + \alpha_n \frac{1}{n-1}(y_t(n) - y_t(3 \text{ month})) + \varepsilon_{n,t+1}$ using actual and model data on bond yields with annual periods. According to the expectations hypothesis, this coefficient should equal one. Constant terms (not shown) are included in all regressions. Robust standard errors are below estimated coefficients. Data runs from October 1976 to January 2019. The time-consistent investor has preferences $\beta = 1, \delta = 0.97$, and the present-biased investor has preferences $\beta = 0.35, \delta = 0.97$.

Maturity n	Data	Time consistency	Present bias
5 years			
Coefficient	-1.87	-2.35	-2.23
10 years			
Coefficient	-3.46	-4.31	-3.91

Table 10: **Properties of asset returns in the data, and for present bias with multiple durations of the present.** Data runs from October 1976 to January 2019. The present-biased investor has varying pairs of durations of the present and present-bias factors (T_S, β): $(3 \text{ months}, 0.7)$, $(1 \text{ year}, 0.35)$ and $(3 \text{ years}, 0.05)$. The long-term discount factor equals $\delta = 0.97$ in all cases. The table presents 3-year bond, 10-year bond and stock returns per annum, in excess of the risk-free rate. The Sharpe ratio is the mean of the excess return divided by the standard deviation.

		Duration present		
	Data	3 months	1 year	3 years
3-year bond				
Mean	1.90	1.66	1.59	1.05
Sharpe	0.48	0.42	0.41	0.27
10-year bond				
Mean	4.10	4.39	4.45	3.11
Sharpe	0.38	0.41	0.42	0.29
Stock				
Mean	7.27	7.50	7.48	7.05
Sharpe	0.48	0.49	0.49	0.46

Table 11: **Properties of asset returns in the data, and for impatient time consistency and present bias.** Data runs from October 1976 to January 2019. The impatient time-consistent investor has preferences $\beta = 1, \delta = 0.86$, and the present-biased investor has preferences $\beta = 0.35, \delta = 0.97$. Panel A presents 3-year bond, 10-year bond and stock returns per annum, in excess of the risk-free rate. The Sharpe ratio is the mean of the excess return divided by the standard deviation. Panel B reports the estimated coefficients and standard errors for the equilibrium prices-of-risk $\hat{\boldsymbol{\lambda}}_F$.

	Data	Impatient time consistency	Present bias
	Panel A: Excess returns		
3-year bond			
Mean	1.90	1.61	1.59
Sharpe	0.48	0.41	0.41
10-year bond			
Mean	4.10	4.48	4.45
Sharpe	0.38	0.42	0.42
Stock			
Mean	7.27	7.49	7.48
Sharpe	0.48	0.49	0.49
	Panel B: Prices-of-risk		
Factor 1: $\hat{\lambda}_{F,1}$			
Coefficient	-	-0.3772	-0.3741
Standard error	-	0.0332	0.0335
Factor 2: $\hat{\lambda}_{F,2}$			
Coefficient	-	-0.1953	-0.1945
Standard error	-	0.0055	0.0056

Figures

Figure 1: **Three discount functions.** This figure graphs the exponential discount function for $\delta = 0.97$, the instantaneous quasi-hyperbolic discount function for $\beta = 0.6$ and $\delta = 0.99$ and a generalised hyperbolic discount function ($D(x) = (1 + \alpha x)^{-\beta/\alpha}, \quad \alpha, \beta > 0$, with $\alpha = 10^5$ and $\beta = 5 \cdot 10^3$).

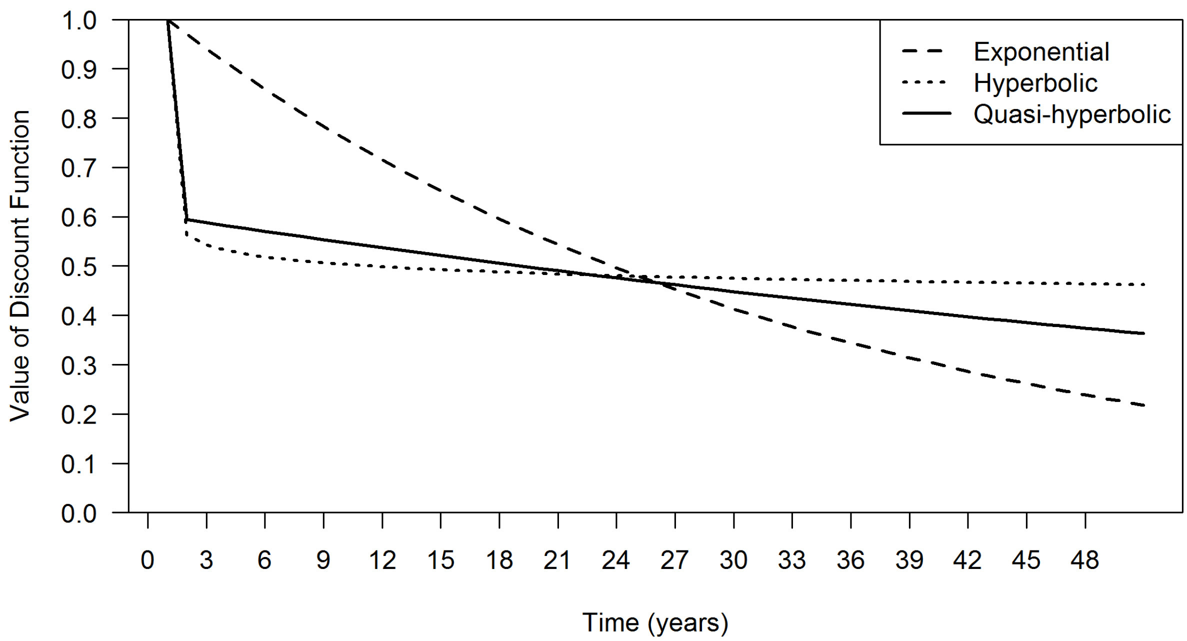

Figure 2: **Quasi-hyperbolic versus exponential Discounting.** This timeline presents the discount structure for an investor with exponential discounting (E) and quasi-hyperbolic discounting (H), from the perspective of time $t = 0$. The quasi-hyperbolic discount function discretely drops with size β when the present ends and the future begins. The present-to-future transition occurs at $T_S = 1$ year and the future ends at year $T_L = 10$, respectively the short-term and long-term planning horizons.

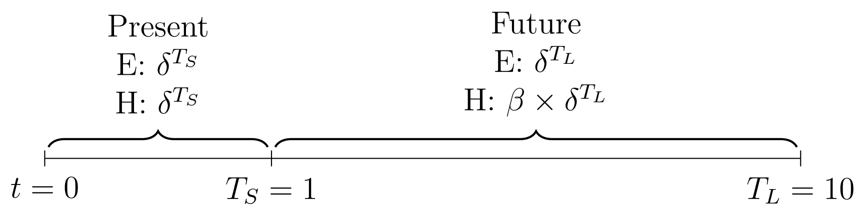

Figure 3: **Ratio of U.S. government debt by maturity to total U.S. government debt outstanding.** This figure shows per maturity the U.S. government debt as a fraction of total U.S. government debt outstanding throughout the period 1 October 1976 to 1 January 2019.