RFS Advance Access published January 4, 2011

Revisiting Asset Pricing Puzzles in an Exchange Economy

Christine A. Parlour
Haas School of Business, University of California-Berkeley

Richard Stanton
Haas School of Business, University of California-Berkeley

Johan Walden
Haas School of Business, University of California-Berkeley

We show that several well-known asset pricing puzzles are largely mitigated if we endow the representative agent with an arbitrarily small minimum consumption level. This allows us to solve the model for parameter values where the standard "Lucas tree" model is not defined. For these parameters, disasters become more important, and the market risk premium therefore higher, even though consumption is *less* risky. Our model yields reasonable risk premia, Sharpe ratios, and discount rates; excess price volatility; and a high market price-dividend ratio. We derive closed-form solutions for all variables of interest. (*JEL* G12)

When a standard one-tree consumption-based exchange economy, with Brownian log-consumption growth and a representative investor with power utility, is calibrated to data, three significant puzzles arise.[1] First is the *equity premium puzzle*, famously posed by Mehra and Prescott (1985): For reasonable values of the risk-aversion coefficient, the implied equity premium is too low. Second is the *risk-free rate puzzle* (see Weil 1989): If risk aversion is chosen to match the equity premium, then the discount rate is implausible. Third is the *excess-volatility puzzle* (see LeRoy and Porter 1981; Shiller 1981): Price volatility in the standard model is the same as dividend volatility and consumption volatility; in reality, however, price volatility is many times higher than

We thank Frederico Belo, Jonathan Berk, Harjoat Bhamra, Nicole Branger, Martijn Cremers, Nicolae Gârleanu, Julien Hugonnier, Martin Lettau, Ian Martin, Rajnish Mehra, Mark Rubinstein, and seminar participants at Columbia University, University of Lausanne, the École Polytechnique Fédérale de Lausanne, Stockholm School of Economics, Uppsala University, the University of Texas at Austin, the 2009 EFA, the 2010 WFA, and the Berkeley Lunch Seminar for helpful comments. We are also grateful to Pietro Veronesi (the editor) and an anonymous referee for helpful suggestions. The explicit characterization of the term structure presented in Section 2.2 first appeared in an unpublished Working Paper, "The Long-term Discount Rate," by the same authors. Send correspondence to Johan Walden, Haas School of Business, UC Berkeley, 545 Student Services Building #1900, Berkeley, CA 94720-1900; telephone: (510) 643-0547; fax: (510) 643-1412. E-mail: walden@haas.berkeley.edu.

[1] Though commonly referred to as a "Lucas" model, the first-order conditions for this economy and associated stochastic discount factor, $\rho \frac{U'(C_{t+1})}{U'(C_t)}$, were first derived by Rubinstein (1976) and later used by Lucas (1978).

doi:10.1093/rfs/hhq130

both consumption and dividend volatility. As summarized in LeRoy (2006), "The conclusion that appears to follow from the equity premium and price volatility puzzles is that, for whatever reason, prices of financial assets do not behave as the theory of consumption-based asset pricing predicts."

The equity premium puzzle remains perhaps the most disturbing counterfactual prediction for the standard model, mainly because such a stylized model should not be "off" by an order of magnitude. More sophisticated models inevitably build on the simple one, making its poor performance especially troubling. In this article, we show that all three of these puzzles are, in fact, extremely fragile. With calibrations as reasonable as in Mehra and Prescott (1985), a very small change to the setup leads to very different levels for the market risk premium, the risk-free rate, and the level of price volatility. The specific change we implement is to introduce an arbitrarily small risk-free consumption stream to the standard model. We call this the *minimum consumption* (MC) economy, and show that this minor modification largely mitigates all three puzzles.

Our results are based on the observation that for some parameter values, beyond what we dub the *breakpoint*, the risky tree in the standard model is so risky that the representative investor's expected utility is negative infinity, and the risk premium is therefore not well defined. With a lower bound on consumption, expected utility remains finite, though it is still strongly affected by low-consumption states in this parameter region. As a result, we obtain a much higher risk premium than in the standard model. Indeed, for low growth rates and personal discount factors, the risk premium in our model for these parameter values can approach $\gamma^2\sigma^2$ instead of the $\gamma\sigma^2$ produced by the standard model.[2] Interestingly, the consumption process in our economy, with probability 1, looks indistinguishable from the standard one-tree model in the long run. Empirically, it would therefore be impossible to distinguish the consumption process in the MC model from that in the standard model, even though the differences in asset pricing are huge. Although the effect of minimal consumption on asset prices is drastic in our model, the stochastic discount factor changes only marginally, so our approach has little to say about the Hansen-Jagannathan bounds.

We are not the first to consider the effect of extreme events in consumption-based asset pricing. For example, Barro (2005) (following Rietz 1988) shows that adding catastrophic risk, either actual or suspected, to the standard model can generate empirically reasonable equity premia. In a similar spirit, Weitzman (2007) argues that parameter uncertainty, by increasing subjective probabilities for low states, significantly increases the equity premium. However, there are two major differences between our results and these papers. First, whereas Barro (2005), Rietz (1988), and Weitzman (2007) all rely on making the lower

2 The value $\gamma^2\sigma^2$ is an upper bound for the risk premium in our model. The risk premium is given in full by $\gamma\max(\gamma-\kappa, 1)\sigma^2$, where $\kappa > 0$ depends on the parameters of the model, and can be close to zero.

tail of the consumption distribution fatter than in the standard model, we actually *reduce* the likelihood of very low states, making the lower tail of the distribution thinner than in the standard model (indeed, we impose a strict lower bound on consumption, so the lower tail has weight zero below this level). This allows us to analyze asset pricing properties for parameters that are typically ignored; we explore this idea in more detail when we consider the robustness of our model. Second, whereas prior papers have typically focused on one puzzle at a time, we show that our model is capable of substantially mitigating all three of the primary puzzles listed above with the same calibrated parameters.

For simplicity, we implement the model in a two-trees framework (see Cochrane, Longstaff, and Santa-Clara 2008), with one risky and one risk-free tree. This makes the analysis tractable, and we obtain closed-form solutions for all variables of interest. We also show that the effect of minimum consumption levels extends to broader classes of model.

In a simple calibration of the MC model, we show that to obtain a market risk premium of 5% requires a risk-aversion coefficient of only $\gamma = 12.2$, compared with the $\gamma = 31$ needed by the standard model.[3] We also show that, in stark contrast to the standard model, the long-term discount rate in our model is *independent* of risk aversion. In the calibration, we get a long rate of 2.4%, so there is no risk-free rate puzzle at the long end of the yield curve. The short rate is −2.8%, which is somewhat low, but far above the −58% implied by the standard model with the same parameters; moreover, instead of the flat term structure in the standard model, we typically get an upward-sloping term structure. Price volatility is also higher than in the standard model: Our calibration yields a price volatility of 10.3%, compared with a consumption volatility of 4%. Finally, our calibration produces a reasonable market Sharpe ratio of 0.49.

Central to our analysis is the existence of a risk-free consumption stream. There are many plausible economic frameworks that give rise to such a sector; we posit two. First, in an economy with technology shocks, if there is enough "memory" in the economy, it is natural to assume that production levels can never fall below some threshold. Similarly, a lower bound on consumption can be interpreted as subsistence farming or consumption.[4] Second, bonds may not be in zero net supply. The assumption that bonds are in zero net supply is consistent with an infinitely lived representative agent in an economy absent any frictions. In particular, any bonds that she issues, she also consumes.

[3] In the standard model, $\gamma = 12.2$ leads to a risk premium of only 2%.

[4] If a cataclysmic event such as a nuclear war occurred, a subsistence level of consumption might not exist. However, since it is also unlikely that financial assets would survive, we restrict our attention to states of the world in which no such event occurs. The only modification needed is that the representative investor has a higher effective personal discount rate in the presence of such events (similar to the increased discount rate in the portfolio problem of an investor with finite, stochastic life length, compared with an infinitely lived one).

By contrast, in a world with finitely lived investors, or with frictions, it may be possible for the current generation to borrow against the consumption of future generations, leading to a positive supply of bonds and risk-free consumption for the current generation over a significant time period. Indeed, in any economy in which Ricardian equivalence fails, government bonds can be in positive net supply.[5]

Intuitively, the existence of a minimum consumption level lowers the value to the representative consumer of claims that pay off in states when her risky consumption is low. The representative consumer weighs two factors when evaluating a claim that pays off when her other consumption is low: first, her current level of consumption, and second, the difference between current marginal utility and marginal utility when the claim pays off. The first factor is important because it affects how far into the future she will consume the claim. A higher current consumption level decreases the value of this claim by increasing the time until its payoff (because the personal discount rate is positive). However, a higher current consumption level also increases the relative difference between current marginal utility and marginal utility at payoff. In this article, we show that the relative importance of these two factors changes drastically when passing the *breakpoint*. In the region in which the standard model is defined, the first effect dominates the second, so for high consumption levels the price of a low-consumption claim is negligible. Beyond the breakpoint, however, the second effect dominates the first, and the claim becomes more and more valuable, the higher the consumption level. In the standard model, the price of such a claim is infinite, which is why the standard model is not defined beyond the breakpoint. By contrast, in the MC economy, the minimum consumption level leads to a finite, albeit high, price for the claim.

A vast literature has suggested other solutions to the classic puzzles, usually based on significant modifications of the standard model. We cannot do justice to this literature here, but we mention a few examples. To solve the equity premium puzzle, some researchers have explored preference specifications that make the stochastic discount factor (SDF) more volatile. Abel (1990) introduced catching-up-with-the-Joneses preferences, while Constantinides (1990; see also Ferson and Constantinides 1991 and Campbell and Cochrane 1999) suggested that consumers form habits. Others have investigated rational bubbles as a potential solution to the excess-volatility puzzle (see, for example, Blanchard 1979; Blanchard and Watson 1982; Froot and Obstfeld 1991). With rational bubbles, prices are highly nonlinear functions of dividends, leading to a higher price volatility. In our model, the market price of equity is a convex function of consumption, which mechanically leads to a higher risk premium and price volatility. This is similar to the price behavior in, for example, Abel

5 In the extreme case, if the representative investor does not care at all about consumption after a certain date, he will take the opportunity to transfer risk-free consumption from beyond that date, if feasible. The economy then behaves like one with a finite horizon and a minimum consumption level. (Our results also hold for long but finite horizons; see Section 3.7.)

(1990) and Froot and Obstfeld (1991). In contrast to these models, however, we make minimal modifications to the standard model; preferences are the same, and there are no bubbles in the MC economy. The only difference is the addition of an arbitrarily small additional consumption stream.

The rest of the article is structured as follows. We proceed by laying out the MC model in Section 1, and study when the differences between this and the standard economy are important. In Section 2, we address the equity premium puzzle, the risk-free rate puzzle, and the excess-volatility puzzle and present a simple calibration. We discuss robustness, how our approach is related to other approaches, and possible generalizations in Section 3. After a brief conclusion, all proofs appear in the Appendix, as does some supporting Mathematica code, which provides numerical backup for our theoretical results.

1. Model

Consider an economy that evolves between times 0 and T, in which there are two sources of the consumption good. As in the standard one-tree model, the first, risky, asset grows stochastically and pays an instantaneous dividend of $D_t\,dt$, where $D_t = D_0\,e^{y(t)}$, $y(0) = 0$, $dy = \mu\,dt + \sigma\,d\omega$, and μ and σ are constants. Here, ω is a standard Brownian motion, which generates a standard filtration, $\mathcal{F}_t$, on $t \in [0, T)$. Unlike the one-tree model, there is also a second, riskless, asset paying a dividend, $B\,dt$, where $B \geq 0$. It will be useful to consider the share of the risky asset in the overall economy, and so we define the risky share, $z(t) = \frac{D_t}{B+D_t}$. We also define $\hat{\mu} = \mu + \frac{\sigma^2}{2}$. The horizon T can be finite or infinite. We focus primarily on the case when $T = \infty$, but we show in Section 3 that the results carry over to the case with large but finite T. In Section 3, we also show how these assumptions on the growth processes can be substantially relaxed.

There is a price-taking representative investor with constant relative risk-averse (CRRA) utility, risk-aversion coefficient $\gamma > 1$, and personal discount rate $\rho > 0$. This investor consumes the total output:

$$U(t) = E_t\left[\int_t^{\infty} e^{-\rho(s-t)} u(B + D_s)\,ds\right], \tag{1}$$

where

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}. \tag{2}$$

We also write $U(t|B, D_t)$, when we want to stress the dependence on B and D_t.

In what follows, we focus our attention on the (economically interesting) case $\mu > 0$. We note that in this case (when $B > 0$), the distribution of the risky share, $z(t) \in (0, 1)$, converges in probability to one for large t, $z \rightarrow_p 1$, and the growth rate of real variables (i.e., dividends and consumption) in the economy behaves much like that in the one-tree model for large t.[6]

The market is dynamically complete, and the usual arguments imply that, in equilibrium, an asset that pays out ξ_t, where ξ_t is an $\mathcal{F}_t$-adapted process satisfying standard conditions, commands an initial price of

$$P_0 = \frac{1}{u'(B + D_0)} E_0 \left[\int_0^\infty e^{-\rho s} u'(B + D_s)\xi_s \, ds \right]. \tag{3}$$

Equation (3) is the Euler equation relating the agent's aggregate consumption, marginal utility, and valuation for all securities.

Notice that if $B = 0$, all resources are in the risky asset and the economy collapses to the standard one-tree model with constant growth and power utility. When $B > 0$, the economy is a special case of that in Cochrane, Longstaff, and Santa-Clara (2008), further generalized in Martin (2009); i.e., it is a so-called "two-trees" economy, in which one of the trees is risk free. We refer to the case $B = 0$ as the *standard model*, whereas when $B > 0$ we have the *minimum consumption* (MC) model.

As we elaborate below, providing the agent a minimal level of insurance (through the risk-free tree) provides new implications. Equivalently, we could have specified the economy as one with no riskless tree but with HARA utility, $u(c) = \frac{(B+c)^{1-\gamma}}{1-\gamma}$, or one in which there is one asset with output $B + D_t$ and a risk-free bond in zero net supply (similar to Rubinstein 1983). More generally, our results will also apply to combinations of these assumptions, such as an MC model with riskless consumption B_c, combined with HARA utility $u(c) = \frac{(B_u+c)^{1-\gamma}}{1-\gamma}$, as long as $B_c + B_u > 0$.

We define

$$\eta = \rho + (\gamma - 1)\mu - (\gamma - 1)^2 \frac{\sigma^2}{2},$$

the dividend yield in the standard model, which will be useful going forward. The properties of the standard model have been extensively analyzed and are summarized in Table 1.

1.1 The Breakpoint

In the MC model, utility and marginal utility are bounded both from above and from below, so (1) and (3) are well defined for arbitrary values of $\mu > 0$, $\sigma > 0$, $\rho > 0$, $B > 0$, and $D_0 > 0$. To clarify the differences between the MC model and the standard one, we study the expected utility of the agent in the two settings. First, observe that the homogeneity of the utility function implies that the value function, U, is scalable as $U(t|B, D_t) = (B + D_t)^{1-\gamma}$

[6] If, on the other hand, $\mu < 0$, the share converges to zero, $z \to_p 0$. In this case, real variables become almost risk free over time. If $\mu = 0$, then the share converges in probability to a two-point distribution with 50% mass at 0 and 50% mass at 1 (the convergence also holds almost surely for $\mu \neq 0$, but not for $\mu = 0$).

Table 1
Properties of the standard model (the consumption model with Brownian log-consumption process and power preferences)

Variable	Value
Risk-free rate, r_s	$\rho + \gamma\left(\mu + \frac{\sigma^2}{2}\right) - \gamma(\gamma+1)\frac{\sigma^2}{2}$
Long rate, r_l	$\rho + \gamma\left(\mu + \frac{\sigma^2}{2}\right) - \gamma(\gamma+1)\frac{\sigma^2}{2}$
Market return, r_e	$\rho + \gamma\left(\mu + \frac{\sigma^2}{2}\right) - \gamma(\gamma-1)\frac{\sigma^2}{2}$
Dividend yield, $\eta \overset{\text{def}}{=} D/P$	$\rho + (\gamma-1)\mu - (\gamma-1)^2\frac{\sigma^2}{2}$
Market risk premium, $r_e - r_s$	$\gamma\sigma^2$
Consumption volatility	σ
Dividend volatility	σ
Price volatility	σ
Market Sharpe ratio	$\gamma\sigma$

$U(t|1-z, z) \overset{\text{def}}{=} (B + D_t)^{1-\gamma}\, w(z)$, where $w(z) \overset{\text{def}}{=} U(t|1-z, z)$. We call $w(z)$ the *normalized* value function at risky share z.

We define the following three variables, which will be helpful going forward:

$$q = \sqrt{\mu^2 + 2\rho\sigma^2}, \qquad \kappa = \frac{\mu + q}{\sigma^2}, \qquad \alpha = \gamma - \kappa. \tag{4}$$

We shall see later that the value of α will be extremely important for the behavior of the model. Note that it is always the case that $\alpha < \gamma$.

Our first result characterizes the normalized value function.

Proposition 1. In the MC model, the normalized value function of the representative agent, $w(z)$, is finite for all $z \in (0, 1)$. It is given by

$$w(z) = \frac{z^{-\kappa}(1-z)^{1-\gamma-\kappa}}{q(1-\gamma)}\left[V\left(\frac{1-z}{z}, \kappa, 2-\gamma\right)\right. \tag{5}$$

$$\left. + \left(\frac{1-z}{z}\right)^{\frac{2q}{\sigma^2}} V\left(\frac{z}{1-z}, \alpha + \frac{2q}{\sigma^2} - 1, 2-\gamma\right)\right]. \tag{6}$$

Here,

$$V(y, a, b) \overset{\text{def}}{=} \int_0^y t^{a-1}(1+t)^{b-1}\, dt \tag{7}$$

is defined for $a > 0$. Also, $w(0) = \frac{1}{\rho(1-\gamma)}$.

Moreover, recall that the dividend yield in the standard model (if it exists) is given by $\eta = \rho + (\gamma-1)\mu - (\gamma-1)^2\frac{\sigma^2}{2}$. Then, if $\eta > 0$, $w(1) = \frac{1}{\eta(1-\gamma)}$. If, in contrast, $\eta \leq 0$, then $w(1) = -\infty$.

The proof of this proposition is given in the Appendix. The last part of Proposition 1 is important. When $\eta > 0$, the value function in the MC model converges to that in the standard model as z approaches one. However, when $\eta \leq 0$, the two models behave completely differently. Note that in this case, while we can still calculate η, it is no longer equal to the dividend yield in the standard model (which does not exist). In this case, the value function is negative infinity in the standard model, and equilibrium is undefined. In contrast, the value function is always finite in the MC model. It is easy to check that the *breakpoint* at which the standard model becomes undefined ($\eta = 0$) occurs at the risk-aversion coefficient

$$\gamma = 1 + \kappa, \tag{8}$$

where κ is defined in (4). It is straightforward to check that $\gamma - (1 + \kappa) > 0$ is equivalent to $\alpha > 1$ and to $\eta < 0$, so above the breakpoint the dividend yield in the standard model is formally negative, as discussed above.[7] Going forward, we shall use the term "below the breakpoint" to refer to sets of parameters for which $\eta > 0$, and "above/beyond the breakpoint" to refer to sets of parameters for which $\eta < 0$. Below the breakpoint point (i.e., for lower γ), the standard model is well defined, whereas above the breakpoint it is not. Thus, although we may expect the MC model to converge to the standard model below the breakpoint, the characteristics of the MC model above the breakpoint are unclear.

To provide further intuition for the breakpoint, we note that although the true drift of the risky tree is $\hat{\mu} > 0$, the risk-adjusted drift term used by the representative investor is lower. In fact, when $B = 0$, for utility purposes the investor treats the drift of the expectation as being $\hat{\mu}' = (1 - \gamma)\hat{\mu} + (1 - \gamma)^2 \frac{\sigma^2}{2}$.[8] When $\hat{\mu}' < 0$, the investor acts as if consumption is expected to be very low for large t. Moreover, when $\rho - \hat{\mu}' < 0$, the expected utility of consumption in the far future is also very low in present-value terms. In this situation, we may expect the representative investor to be prepared to pay a lot for insurance against bad states of the world in the far future. The condition $\rho - \hat{\mu}' < 0$ is exactly the condition of being above the breakpoint.

---

[7] It is well known that expected utility is infinite beyond the breakpoint in the standard model. For example, Campbell (1986) develops a parameter restriction for general stationary processes, which is the discrete-time version of the breakpoint equation. The breakpoint condition also occurs in Martin (2009), though in a different context. Martin (2009) characterizes the prices of "small firms" below the breakpoint. We examine the properties of the market above the breakpoint.

[8] This holds in the sense that $U = \frac{D_0^{1-\gamma}}{1-\gamma} \int_0^T e^{-\rho t} E\left[\left(\frac{D_t}{D_0}\right)^{1-\gamma}\right] dt = \frac{D_0^{1-\gamma}}{1-\gamma} \int_0^T e^{-\rho t} E\left[R_t\right] dt$, where the risk-adjusted diffusion process $R_t \equiv \left(\frac{D_t}{D_0}\right)^{1-\gamma}$ satisfies $R_0 = 1$, $\frac{dR_t}{R_t} = \hat{\mu}' \, dt + (1 - \gamma)\sigma \, d\omega$ (following from standard Itô calculus), and therefore $E[R_t] = e^{\hat{\mu}' t}$.

1.1.1 Calibration. To get a sense for what the breakpoint implies for risk aversion, suppose that the consumption growth rate, volatility of growth, and personal discount rate are

$$\hat{\mu} = 0.75\%, \qquad \sigma = 4\%, \qquad \text{and} \qquad \rho = 1\%, \tag{9}$$

respectively, these values selected as follows:

Consumption Volatility Our choice of 4% is within (though at the top end of) the range used by prior authors. In particular, it is close to the 3.6% used by Mehra and Prescott (1985). Campbell (2003) reports the average annual consumption volatility for ten countries between 1970 and 2000 as 2.13%, and a value of 3.2% for annual volatility in the United States between 1891 and 1998. While our value of 4% is somewhat higher than these numbers, these previous studies almost certainly underestimate the true volatility of consumption growth. In particular, Triplett (1997) and Savov (2011) (Internet Appendix) point out that three statistical issues with the National Income and Product Accounts (NIPA) consumption data in the United States automatically lead to an artificially low volatility in measured consumption: (i) benchmarking;[9] (ii) non-reporting;[10] and (iii) the residual method used to calculate consumption.[11] In response, Savov (2011) uses garbage generation data from the Environmental Protection Agency (EPA) as a proxy for consumption and estimates consumption volatility to be around 2.5 times as high as NIPA consumption expenditures—2.9% from 1960 to 2007. He also cites an alternative survey of garbage data by the journal *Biocycle*, which estimates a volatility of 4.1% per year. In addition to these statistical issues, Parker (2001) and Gabaix and Laibson (2001) argue that another reason the usual historical measures may well be substantially too low is that consumption adjustment costs may artificially reduce measured consumption volatility. Moreover, if individual investors are adjusting consumption at infrequent, but different, points in time, aggregate consumption will be smoother than the consumption of any individual. Finally, Malloy, Moskowitz, and Vissing-Jørgensen (2009) note that asset prices are determined by those who actually hold assets. Focusing on the consumption of shareholders rather than all individuals, they estimate the annual volatility of consumption to be between 3.6% and 5.4%, depending on whether an adjustment is made for the possibility of different people being shareholders in different periods.

---

9 Benchmarking, a comprehensive measurement of consumption, occurs only once every five years. In non-benchmarking years, the Census Bureau's Retail Trade Survey is used to estimate consumption updates, but this does not include all expenditure types, so many values are interpolated or forecast based on the most recent benchmark values.

10 Savov (2011) reports that around 7% of the annual data currently suffers from this problem, down from 14% ten years ago, and probably more in the preceding decades. In addition, there is no fixed method for including new retail establishments. He suggests that it is likely that non-reporting and newly formed retailers are also those with the most volatile sales.

11 For most commodities, personal expenditure is calculated by subtracting government and business purchases from total estimated domestic supply. Business purchases are in many cases estimated.

Consumption Growth Rate and Personal Discount Rate A growth rate in the neighborhood of 1% per year is in line with observation as well as with previous theoretical studies, as is a personal discount rate of 1% per year (see, for example, Cochrane 2001 and references therein).

Implied Risk Aversion With these parameters, Equation (8) shows that the breakpoint occurs at $\gamma = 10.6$, a not unreasonably high number (Mehra and Prescott 1985 consider risk-aversion coefficients up to 10, and several studies use higher values—for example, Malloy, Moskowitz, and Vissing-Jørgensen 2009 use values of γ between 10 and 15).

Relative Sizes of Trees At this point we are not making a specific assumption about the relative size of the trees, z, but we shall be considering values close to 1. Note that $B \to 0$, $D \to \infty$, and $z \to 1$ are all equivalent, so all of our results for z close to 1 can be interpreted as results when $B \to 0$.

2. Puzzles Revisited

Without a loss of generality, we assume that $B \equiv 1$, i.e., that the risk-free part of the consumption stream is of size 1, and from (3) we define $P(D_0)$ to be the price of the total consumption output in the economy,

$$P(D_0) = E\left[\int_0^{\infty} e^{-\rho s}\left(\frac{1+D_0}{1+D_t}\right)^{\gamma}(1+D_t)\,ds\right]. \tag{10}$$

The price for general $B \neq 1$ then follows from the relation $P(B, D_0) = BP\left(\frac{D_0}{B}\right)$. We will specifically be interested in the dynamics for large D, or, equivalently, for z close to 1.

We provide an explicit characterization of the price of the market:

Proposition 2. The price function $P(D)$ is

$$P(D) = (1+D)^{\gamma}\,\frac{D^{-\kappa}}{q}\left[V(D,\kappa,2-\gamma) + D^{\frac{2q}{\sigma^2}}V\left(\frac{1}{D}, \alpha+\frac{2q}{\sigma^2}-1, 2-\gamma\right)\right], \tag{11}$$

where $V(y,a,b) \stackrel{\text{def}}{=} \int_0^y t^{a-1}(1+t)^{b-1}\,dt$, and q, κ, and α are defined as in Equation (4), i.e., $q = \sqrt{\mu^2 + 2\rho\sigma^2}$, $\kappa = \frac{\mu+q}{\sigma^2}$, $\alpha = \gamma - \kappa$.

Similar formulas are derived in Cochrane, Longstaff, and Santa-Clara (2008) (for $\gamma = 1$) and in Martin (2009), though there they are expressed in terms of hypergeometric functions.

Figure 1 shows the price-dividend ratio multiplied by $|\eta|$, for different choices of γ. This product equals 1 in the standard model, regardless of D. Recall that the breakpoint risk aversion for this set of parameters is $\gamma = 10.6$.

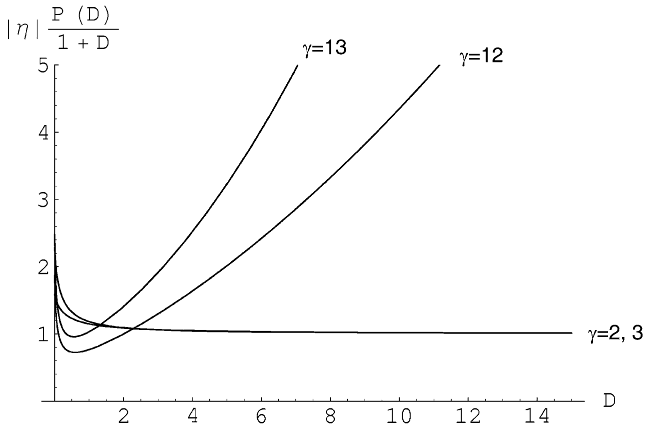

Figure 1
Scaled market price-dividend ratio in MC model as a function of D
The parameters of the model are according to (9), i.e., $\hat{\mu} = 0.75\%$, $\sigma = 4\%$, $\rho = 1\%$, with risk-aversion coefficients $\gamma = 2, 3, 12, 13$.

For $\gamma = 2$ and $\gamma = 3$ (the lower lines), the ratios quickly converge to 1 as D increases, in line with the intuition that when D is large, the economy is effectively the same as the standard model. However, for $\gamma = 12$ and $\gamma = 13$, the function quickly increases as D grows. It is clear from the figure that price dynamics above the breakpoint are quite different from those below. We now explore why.

Consider the price of a digital option that pays a very small amount (say \$1) in the event that total consumption drops to $1+\epsilon$.[12] Let $K(D_0, \epsilon)$ be the value of such an asset, starting at D_0 (where we assume that $D_0 > \epsilon$).

It follows from (3) that K is given by

$$K(D_0, \epsilon) = \left(\frac{1+D_0}{1+\epsilon}\right)^{\gamma} E_0\left[e^{-\rho\tau_f}\right],$$

where τ_f is the stopping time

$$\tau_f \stackrel{\text{def}}{=} \inf_t\{t : D_t \leq \epsilon\}. \tag{12}$$

The value of this claim is thus made up of two offsetting elements. The first element, $\left(\frac{1+D_0}{1+\epsilon}\right)^{\gamma}$, is the incremental marginal utility of the agent when he consumes, given his consumption today. The contribution of this part is heavily dependent on the agent's risk aversion. A high risk aversion implies a high difference between the marginal utility at the consumption level $1+\epsilon$ and at $1+D_0$, which has a positive effect on the price. For large D_0, the first element

12 Technically, this is an American digital cash-or-nothing put option.

behaves like D_0^{γ}, since the relative value of consumption at ϵ is higher the wealthier the economy is at the starting point. Because of the direct dependence on the risk-aversion coefficient, we call this the "risk-aversion effect."

The second element, $E_0\left[e^{-\rho\tau_f}\right]$, represents the expected discounted value of \$1 when consumption hits the boundary. We therefore call it the "discount effect." It is straightforward, using standard results for stopping-time distributions, to show that

$$E_0\left[e^{-\rho\tau_f}\right] = \left(\frac{\epsilon}{D_0}\right)^{\kappa}, \tag{13}$$

where $\kappa > 0$ is defined in (4).[13] Since $\kappa > 0$, it is always the case that this term is decreasing in D_0. This makes sense because the higher D_0, the longer it will take to reach ϵ (and the lower the chance that ϵ will ever be reached). It is easy to see that κ is increasing in μ and ρ, but decreasing in σ. All these properties are intuitive: A higher growth rate, μ, lowers the chance that ϵ will ever be reached, and thereby decreases the time value of the digital option. An increase in the volatility, σ, has the opposite effect. Finally, an increase in the personal discount rate, ρ, lowers the discounted value of the option.

Putting the risk aversion and discount effects together, we arrive at

$$K(D_0,\epsilon) = \epsilon^{\kappa}\left(\frac{1+1/D_0}{1+\epsilon}\right)^{\gamma} D_0^{\gamma-\kappa}. \tag{14}$$

Given a fixed $\epsilon > 0$, $\epsilon^{\kappa}\left(\frac{1+1/D_0}{1+\epsilon}\right)^{\gamma}$ approaches a positive constant for large D_0. By contrast, the behavior of $D_0^{\gamma-\kappa}$ depends on $\gamma-\kappa$. Below the breakpoint (i.e., for $\gamma-\kappa < 1$), for large D_0 this asset is worth much less than D_0, i.e., $K(D_0,\epsilon)/D_0$ goes to zero as D_0 goes to infinity. In this case, the discount effect dominates the risk-aversion effect for large D_0. Above the breakpoint, on the other hand (i.e., for $\gamma-\kappa > 1$), this asset becomes very valuable for high D_0, in a nonlinear fashion. The risk-aversion effect now dominates the discount effect.

The central intuition of the article is that the trees contain this type of payout (they pay something in the bad states of the world). Therefore, above the breakpoint, the market value of these trees will also increase superlinearly with D_0. In fact, we will show that they behave just like the digital options above the breakpoint, with their value growing like $D_0^{\gamma-\kappa}$ for large D_0.

The digital option argument also provides an intuition for why the standard model does not work above the breakpoint. The single tree in the standard model also contains a collection of these types of threshold payments. The single tree does not, however, guarantee the representative agent the subsistence

[13] The expression for κ can be derived from the first-passage-time distributions (see Ingersoll 1987). It can also be derived using methods from the real-options literature. Similar to Dixit and Pindyck (1994, 142–44), the expectation can be derived as a solution (of the form (13)) to an ordinary differential equation. Here, κ is the positive root to the characteristic equation $\frac{\sigma^2}{2}\kappa^2 - \mu\kappa - \rho = 0$.

level of $B = 1$. The first term in the equation corresponding to (14) therefore contains only ϵ (not $1 + \epsilon$) in the denominator. It follows that such claims will be much more valuable in the one-tree economy because when the agent's consumption is low (ϵ low), her marginal utility will be very high and therefore the value of such claims will explode. In this way, the risky tree becomes infinitely valuable. We will return to this point in more detail in Section 3, where we show that a similar argument also holds for finite-horizon economies.

It is possible to derive the following asymptotic results for large D for the behavior of the market price-dividend ratio in the MC economy.[14]

Proposition 3. The asymptotic price-dividend ratio in the MC model depends on the parameter region. Specifically,

(i) Below the breakpoint (i.e., for $\alpha < 1$ so that the value function is finite in the one-tree model), for large D the price-dividend ratio converges to $\frac{P(D)}{1+D} = \frac{1}{\eta}$.

(ii) Above the breakpoint (i.e., for $\alpha > 1$ so that the value function is infinitely negative in the one-tree model), for large D the price-dividend ratio converges to $cD^{\alpha-1}$, for some constant $c > 0$, where α is defined in (4).

It is immediate from Proposition 3 that the exponent of the asymptotic price-dividend ratio behaves like $\max(\alpha, 1) - 1$. It is thus the "convexity parameter," $\max(\alpha, 1)$, which governs the behavior of price-dividend ratios (and prices) for large D. Figure 2 shows the convexity parameter as a function of risk aversion (γ) for some different parameter choices.

The convexity of the price function lies at the heart of our analysis of the asset pricing puzzles, to which we now turn.[15]

2.1 The Risk Premium

It is important to stress that reasonable values of the exogenous parameters are consistent with the region in which prices and price-dividend ratios are undefined in the standard model. In the MC economy, the asymptotic expected return on the market depends on the parameter values. Recall that the instantaneous expected return on the market is

$$r_e\, dt = E\left[\frac{dP}{P} + \frac{1+D}{P}\, dt\right]. \tag{15}$$

[14] Throughout the article, we study the value of the total $B + D$ consumption flows. We obtain identical asymptotic results for the purely risky part of the economy, i.e., the value of the D consumption flows.

[15] The convexity of the price function above the breakpoint (shown in Proposition 3(ii)) is crucial for the subsequent results. The convexity can also be verified numerically. We provide Mathematica code for the numerical calculations in the Appendix.

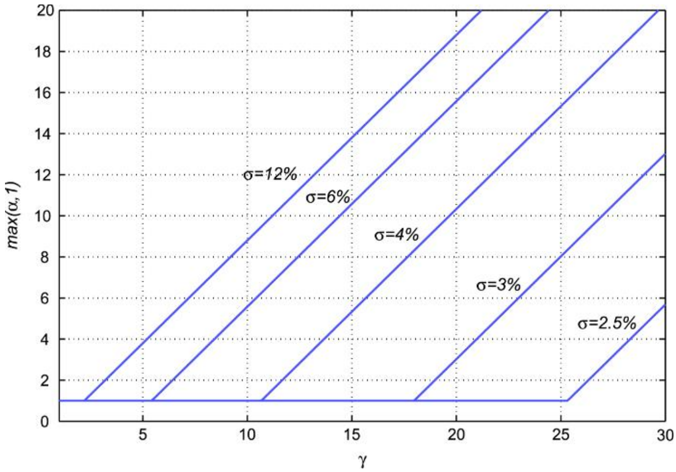

Figure 2
Convexity parameter, $\max(\alpha, 1)$**, as a function of risk aversion,** γ
Parameters: $\hat{\mu} = 0.75\%$, $\rho = 1\%$, $\sigma = 2.5\%, 3\%, 4\%, 6\%, 12\%$.

We have

Proposition 4. For z close to 1,

(i) Below the breakpoint, the expected return on the market is the same as in the standard model: $r_e = \rho + \gamma\mu - \gamma(\gamma - 2)\frac{\sigma^2}{2}$.

(ii) Above the breakpoint, the expected return on the market is $r_e = \alpha\mu + \alpha^2\frac{\sigma^2}{2}$, where α is defined in (4).

To get an intuition for the results in Proposition 4, we note that below the breakpoint, the price is essentially the same as in the standard model (as shown in Proposition 3(i)), so expected returns will essentially be the same. Above the breakpoint, however, the second term in (15) becomes small for large D (as implied by Proposition 3(ii)). Moreover, since $P(D) \sim D^\alpha$, the first term behaves like

$$E\left[\frac{dP}{P}\right] = \frac{(\mu + \frac{\sigma^2}{2})P'\,dt + \frac{\sigma^2}{2}D^2P''\,dt}{P} \approx \alpha(\mu + \frac{\sigma^2}{2})\,dt + \alpha(\alpha - 1)\frac{\sigma^2}{2}\,dt.$$

It follows that the market risk premium can also become large. In fact, it is well known that the risk premium, $r_e - r_s$, can be expressed as

$$(r_e - r_s)\,dt = -\operatorname{cov}\left(\frac{dM}{M}, \frac{dP}{P}\right), \tag{16}$$

where r_s is the short-term risk-free rate and M is the pricing kernel, which is equal to $e^{-\rho t}(1+D)^{-\gamma}$ in the MC economy (with $B = 1$). It therefore follows from standard Itô calculus that

$$r_e - r_s = \gamma \max(\alpha, 1)\sigma^2. \tag{17}$$

For $\alpha < 1$, the risk premium is thus the same as in the standard model. For $\alpha > 1$, however, it is larger, due to the convexity of P as a function of D. In this case, through α, the risk premium now depends on the economy's growth rate, μ, and the personal discount factor, ρ, and is decreasing in both of these parameters.

One immediate implication of Proposition 4 is

Corollary 1. For z close to one, for low values of μ and ρ, or high values of σ, the risk premium is close to $\gamma^2\sigma^2$.

Thus, if μ and ρ are low and/or σ is large, then κ is close to zero and the risk premium, $\gamma \max(\gamma - k, 1)\sigma^2$, is close to $\gamma^2\sigma^2$. The result emphasizes that the risk premium has a very different dependence on the parameters of the economy here, compared with the standard one-tree economy, in which the risk premium is $\gamma\sigma^2$. Since it is always the case that $\kappa > 0$, $\gamma^2\sigma^2$ is also an upper bound on the risk premium in the MC economy, regardless of parameter values.

The intuition behind the equation for the risk premium (17) is clear. As we saw in Section 2, below the breakpoint, the discount effect dominates the risk-aversion effect, so the values of digital options that pay off in bad states of the world are marginal for large D. Therefore, the pricing in the states of the world close to current D will dominate the pricing function. Since the risk-free asset is marginal in these states of the world, asset dynamics will look much like in the one-tree model. Specifically, below the breakpoint all variables of interest converge to the same values as in the one-tree economy as D becomes large—or equivalently, as $z \to 1$.

Above the breakpoint, on the other hand, the risk-aversion effect dominates the discount effect. The value of digital options that pay off in bad states of the world now increases as D^α, when D increases. Therefore, the price function is very different from the one-tree price function, even as $z \to 1$. The convexity of the price function immediately implies a higher equity premium. Specifically, from (16) it follows that there are two parts of the risk premium. The first part depends on the pricing kernel, $\frac{dM}{M}$, and the contribution of this part when z is close to 1 is $\gamma\sigma$ in both the one-tree and MC economies. The contribution of the second part, $\frac{dP}{P}$, however, is different in the two models. Whereas the contribution is σ, leading to a risk premium of $\gamma\sigma \times \sigma = \gamma\sigma^2$ when the price function is linear, it is $\alpha\sigma$ when $P(D)$ grows like D^α, $\alpha > 1$. The risk premium is therefore $\gamma\sigma \times \alpha\sigma = \gamma\alpha\sigma^2$ above the breakpoint in the MC economy for z close to 1.

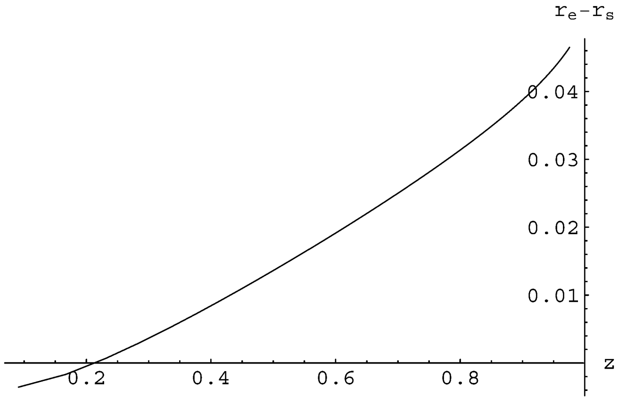

Figure 3
Market risk premium as a function of z, for fixed risk aversion
Parameters: $\hat{\mu} = 0.75\%$, $\rho = 1\%$, $\sigma = 4\%$, $\gamma = 12.25$, implying that $\alpha = 2.55$. The asymptotic risk premium is $\gamma\alpha\sigma^2 = 5.0\%$.

These are asymptotic results, for z close to 1. In Figure 3, we illustrate the market risk premium for a fixed risk aversion, $\gamma = 12.25$ (which is above the breakpoint), as we vary the risky share, z (recall that $z = \frac{D}{D+B} \in (0, 1)$), using the parameters in (9). As z approaches 1, there is indeed convergence to the asymptotic value of 5.0%. Comparing this with the risk premium implied by the standard model, $\gamma\sigma^2 = 2.0\%$, we see that the premium in the MC model is substantially higher.

The following Figure 4 displays the market risk premium for z close to 1 as we vary both risk aversion and volatility (each curve corresponds to a different volatility). Beyond the breakpoint, the risk premium increases very quickly in a convex fashion, implying that a small increase in risk aversion drastically increases the market risk premium.

2.2 The Term Structure

The term structure is also quite different in the MC economy. From (3), it follows that a zero-coupon risk-free bond with maturity date τ has the price

$$P^\tau = e^{-\rho\tau} E_0 \left[\left(\frac{B + D_0}{B + D_t} \right)^\gamma \right]. \tag{18}$$

We can rewrite this expectation in terms of the risky share, $z = \frac{D_0}{B+D_0}$,

$$P^\tau = e^{-\rho\tau} E_0 \left[\left(1 + z \left(\frac{D_t}{D_0} - 1 \right) \right)^{-\gamma} \right], \tag{19}$$

and since the distribution of $\frac{D_t}{D_0}$ does not depend on D_0, it immediately follows that the price can be written as a function of z alone, $P^\tau(z)$, given by the following proposition:

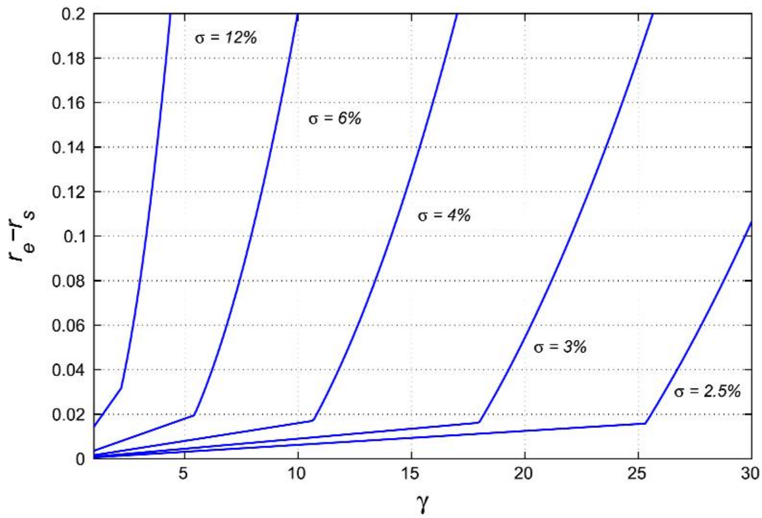

Figure 4
Market risk premium for large z close to 1, as a function of risk aversion, γ
Below the breakpoint, the risk premium is the same as in the standard model, and linear in γ, $r_e - r_s = \gamma\sigma^2$. Above the breakpoint, the risk premium is a steeply convex function. Parameters: $\hat{\mu} = 0.75\%$, $\rho = 1\%$, $\sigma =$ 2.5%, 3%, 4%, 6%, 12%.

Proposition 5. Define the log-relative size of the sectors as $d = \log(z/(1 - z))$. Then, the price of a τ-period zero-coupon bond is given by

$$P^\tau = (1 + e^d)^\gamma e^{-\rho\tau} \frac{1}{\sqrt{2\pi\sigma^2\tau}} \int_{-\infty}^{\infty} \frac{e^{-(y-d-\mu\tau)^2/(2\sigma^2\tau)}}{(1+e^y)^\gamma} dy. \qquad (20)$$

This result follows immediately from Equation (19). An equivalent expression (Equation (29)) that is more convenient for calculation appears in the Appendix.

Martin (2009) independently characterizes the term structure in an economy with many trees. His framework is more general than ours, in that it allows for general Levy processes and multiple trees, but his solution is based on Fourier transform techniques and so is different from those in Proposition 5 and in the Appendix.

In the MC economy, the term structure is no longer constant. Defining the τ-period spot rate as

$$r^\tau = -\frac{\log(P^\tau)}{\tau},$$

we use Equation (29) in the Appendix to study the yield curve with parameters chosen according to (9), $z = 70\%$, and risk-aversion coefficients between 6

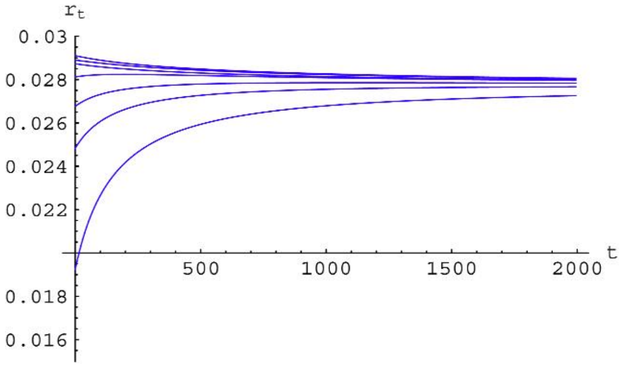

Figure 5
Term structure of interest rates in the MC model
Parameters: $z = 0.7$, γ varies between 6 (highest curve) and 12 (lowest curve). Other parameters are according to (9).

and 12. The choice of $z = 70\%$ means that the risky tree initially dominates the economy, and the risky share converges to $z = 1$ as t grows, so the consumption growth rate is fairly stable in this economy. The results are shown in Figure 5.

We note that the yield curves in the figure can slope upward or downward, and can even be hump-shaped. The slope increases with the risk-aversion coefficient, γ, and so, in general, does the curvature. Moreover, although the short end of the curve is sensitive to γ, as in the one-tree model, there seems to be an asymptotic long-term rate that does not vary much with γ. To understand these properties of the yield curve, we analyze the short rate, r_s, and the long rate, r_l, defined to be

$$r_s = \lim_{\tau \searrow 0} r^\tau, \qquad \text{and} \qquad r_l = \lim_{\tau \to \infty} r^\tau,$$

respectively.

Proposition 6. In the MC economy, the short-term rate is

$$r_s = \rho + \gamma z\left(\mu + \frac{\sigma^2}{2}\right) - \gamma(\gamma + 1)\frac{\sigma^2}{2}z^2.$$

For $z \in (0, 1)$, if $\mu \le \gamma\sigma^2$, the long-term rate is

$$r_l = \rho + \frac{1}{2} \times \frac{\mu^2}{\sigma^2}. \tag{21}$$

If, on the other hand, $\mu > \gamma\sigma^2$, the long-term rate is

$$r_l = \rho + \gamma\left(\mu + \frac{\sigma^2}{2}\right) - \gamma(\gamma + 1)\frac{\sigma^2}{2}. \tag{22}$$

Thus, the short rate has the same structure as in the standard model and, as long as $\mu > \gamma\sigma^2$, the long rate is also the same as in the standard model. This makes intuitive sense, since the economy will almost surely be very similar to the one-tree economy in the long run. If $\mu < \gamma\sigma^2$, however, the long rate is a constant, independent of the risk-aversion parameter. Since

$$\eta = \rho + (\gamma - 1)\mu - (\gamma - 1)^2 \frac{\sigma^2}{2} > (\gamma - 1)\left(\mu - (\gamma - 1)\frac{\sigma^2}{2}\right)$$
$$> (\gamma - 1)\left(\mu - \gamma\sigma^2\right),$$

it will always be the case that the long rate is independent of risk aversion above the breakpoint (i.e., when η is negative).

In our previous numerical example, with parameters according to (9) and $\gamma = 12.25$, this implies that the long rate is $r_l = 2.4\%$. The short rate depends on z, as shown in Figure 6. For z close to unity, i.e., for large D, it becomes negative. At $z = 1$, it is -2.8%. While negative, this is nevertheless far more reasonable than the values we would obtain if we calibrated the standard model to the market risk premium. For example, a risk premium of 5% would imply a risk-free rate of -58% in the standard model. We are, of course, dealing with real variables, so a negative discount rate is obviously possible, although this value is clearly extreme.

Our focus is on the case when z is close to one. We note in passing, however, that for lower z (i.e., when the relative size of the risk-free tree is not negligible), the short rate is also positive. In our calibration, for $z \leq 0.8$, the short rate is positive. From Figure 3, we see that the risk premium is about 3% at $z = 0.8$. Finally, we note that the volatility of the short rate, $\sigma(r_S)$, is low and depends on z. In our example, $\sigma(r_S)$ varies between 0 and 0.08% and reaches its maximum at $z \approx 0.75$.

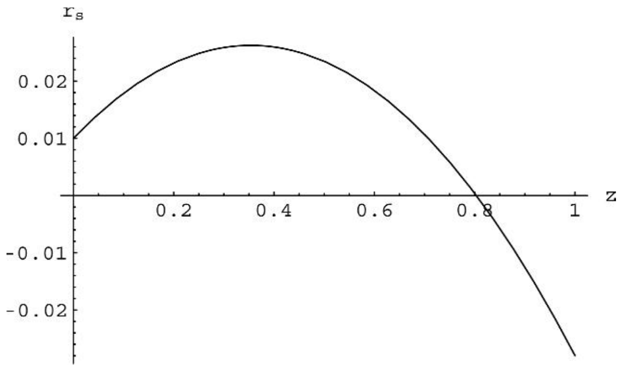

Figure 6
Short rate as a function of z
Parameters of the model are according to (9) and $\gamma = 12.25$.

This γ-independence above the breakpoint stands in stark contrast to the results in the standard model, where the interest rate is very sensitive to risk aversion. Specifically, in the MC model, the long rate is always greater than the personal discount rate, $r_l > \rho$, regardless of the aggregate risk aversion in the economy, and is therefore positive.[16,17] This γ-independence thus offers a resolution to the risk-free rate puzzle at the long end of the term structure.

The reason why risk aversion becomes unimportant for bond yields as the horizon increases, even though bond prices depend on risk aversion, is that differences between bond prices in economies indexed by different levels of risk aversion are sufficiently small, compared with the compounding inherent in the yield calculation, that the price differences become unimportant at the long end of the curve. The price of a bond is the expected discounted value of a dollar multiplied by the representative agent's marginal utility. In the MC model, the marginal utility (irrespective of risk aversion) is bounded above and below. If the agent consumes the fruit of a risk-free tree, which provides insurance, then marginal utility is always bounded above. Indeed, one can find an upper bound on the ratio of marginal utilities for two agents with the same personal discount factor but different risk-aversion coefficients independently of time horizon. Therefore, bond prices for the same maturity for any two economies that differ only in the risk aversion of their representative agents will not differ "by much." For long maturities, this will lead to similar yields.

The difference between the long rates in the standard and MC economies further underscores the fragility of the CRRA-lognormal model over longer time horizons. Regardless of how close z is to 1 in the MC model, the long-term rate is drastically different from when z is identically equal to 1. The differences between the two models are driven by the insurance the risk-free tree provides in the far-left tails. Moreover, although the long rate is always γ-independent above the breakpoint, there are also regions *below* the breakpoint in which it is γ-independent.

At a broad level, our results are reminiscent of, but distinct from, those found in Weitzman (1998, 2001). Weitzman argues that if there is parameter uncertainty, the long-term discount rate is lower than that inferred from the short- and mid-term rates. We agree with Weitzman that a careful analysis of the implicit assumptions about return distributions and utility in the tails is needed to understand the long-term discount rate. Both Weitzman's and our results are

16 A somewhat related result on the long rate is presented in Dybvig, Ingersoll, and Ross (1996), who show that long rates can never fall over time. Within our specific economy, our result is stronger than the Dybvig-Ingersoll-Ross theorem, since it states that r^l is constant over time and across risk aversion.

17 We have verified that the formula is indeed correct by numerically integrating Equation (18) directly. Mathematica code is provided in the Appendix, showing that with parameters $\rho = 1\%$, $\mu = 3.5\%$, $\sigma = 20\%$, $\gamma = 2.5$, the long rate converges to $r_l = \rho + \frac{\mu^2}{2\sigma^2} = 2.53\%$ (in line with Equation (21), since $3.5\% < 2.5 \times 20\%^2$). On the contrary, Equation (22) would, for example, give $r_l = \rho + \gamma\mu - \gamma^2\sigma^2/2 = 1\% + 2.5 \times 3.5\% - 2.5^2 \times 20\%^2/2 = -2.75\%$. By varying B, D_0, and γ, it is easily verified that r^l does not depend on these parameters.

driven by the extreme importance of the worst states in longer horizons. Unlike in Weitzman (1998, 2001), however, the long rate in our model may be higher than the short rate. This distinction is obviously important if existing market data are used to infer a maximum possible discount rate.

2.3 Excess Volatility

Above the breakpoint, prices are not linearly related to consumption, but, as we observed in Proposition 2, are a convex function of dividends. It naturally follows, then, that the volatility of prices is much higher than the volatility of the underlying dividends. In fact, it is easy to show that the price volatility is

$$vol\left(\frac{dP}{P}\right) = \max(\alpha, 1)\sigma. \tag{23}$$

In our numerical example, with parameters according to (9) and $\gamma = 12.25$, this implies a market price volatility of 10.3%, which is more than 2.5 times the dividend (and consumption) volatility of 4%. Since $C = B + D$, if we think of B as a bond and D as a stock, then the volatility of consumption will actually be somewhat lower than that of dividends. This is, of course, in line with what we see in practice. The magnitude of the difference will be small, though, as our focus is on the case where $B \ll D$. If we alternatively interpret both B and D as being (different) parts of the stock market, one riskier than the other (somewhat reminiscent of Rubinstein 1983), dividend and consumption volatility will be exactly equal. The model thus naturally leads to excess volatility, both with respect to consumption and with respect to dividends. Since $\alpha < \gamma$, an upper bound on the excess volatility is given by the risk-aversion parameter.

Table 2 summarizes the formulas and numerical results we have derived.

Table 2
Properties of the MC model for large D

Variable	Formula	Value-MC
Short rate, r_s	$\rho + \gamma\left(\mu + \frac{\sigma^2}{2}\right) - \gamma(\gamma+1)\frac{\sigma^2}{2}$	−2.8%
Long rate, r_l, when $\mu < \gamma\frac{\sigma^2}{2}$	$\rho + \frac{\mu^2}{2\sigma^2}$	2.4%
Long rate, r_l, when $\mu > \gamma\frac{\sigma^2}{2}$	$\rho + \gamma\left(\mu + \frac{\sigma^2}{2}\right) - \gamma(\gamma+1)\frac{\sigma^2}{2}$	
Market return, r_e, when $\alpha > 1$	$\alpha\mu + \alpha^2\frac{\sigma^2}{2}$	2.2%
Market return, r_e, when $\alpha < 1$	$\rho + \gamma\left(\mu + \frac{\sigma^2}{2}\right) - \gamma(\gamma-1)\frac{\sigma^2}{2}$	
Risk premium, $r_e - r_s$	$\gamma\max(\alpha, 1)\sigma^2$	5.0%
Consumption volatility	σ	4%
Dividend volatility	σ	4%
Price volatility	$\max(\alpha, 1)\sigma$	10.3%
Market Sharpe ratio	$\gamma\sigma$	0.49

An example is given with parameters according to (9), $\hat{\mu} = 0.75\%$, $\sigma = 4\%$, $\rho = 1\%$, and $\gamma = 12.25$, implying that $\alpha = 2.58$.

3. Discussion and Generalizations

3.1 Sensitivity of Standard Model

Our model shows how minor changes to the process in low-consumption states (when CRRA expected utility becomes unbounded as consumption approaches zero) drastically change the results obtained from the standard one-tree model. The fragility of expected utility when utility is unbounded has been much studied, supported by the theoretical work of Nielsen (1984, 1987), who develops an axiomatic foundation for expected utility theory that allows for unbounded utility functions. More recently, Geweke (2001) shows that the CRRA-lognormal framework is very fragile with respect to distributional assumptions in the far-left tails. For example, he shows that for $\gamma > 1$, expected utility at some future date is not finite if $\log(C)$ has a t-distribution with any number of degrees of freedom, ν, so we cannot use expected utility to make optimal choices (even though, for high values of ν, this distribution is impossible to distinguish econometrically from lognormality). Thus, even if the true distribution is normal, but the mean and variance are unknown (with standard forms for their priors), expected utility to a Bayesian updater is not finite even in the limit as the sample length goes to infinity.

Geweke (2001) notes that the extreme sensitivity of the finiteness of expected utility to assumptions about tail distributions carries over to implications we might draw about quantities such as the equity premium and the level of real interest rates. However, neither Nielsen (1984) nor Geweke (2001) provides any specific quantitative implications. Pursuing this line of thought, Barro (2005) (following Rietz 1988) generates empirically reasonable risk premia by allowing for some additional probability of extremely low consumption states. Weitzman (2007) adds additional weight to low-consumption states via parameter uncertainty.[18] The importance of very low consumption states in the CRRA-lognormal framework was also emphasized in Kogan et al. (2006), who studied the price impact of irrational traders in capital markets.

The intuition behind our model is, in spirit, somewhat similar, in that we also focus on the impact of very bad outcomes. However, whereas the papers above all fatten the lower tail of the consumption distribution, we make the lower tail *thinner*. This allows us to study regions of parameter space, invalid under the standard model, where bad events have a much larger effect on expected utility. In these regions, the risk premium is higher even though there is no "jump risk" in the MC model. In addition, whereas the other papers focus on one puzzle at a time (usually the equity premium puzzle), we show that our modification of the standard model is able at the same time to substantially mitigate the equity premium puzzle, risk-free rate puzzle, and excess volatility puzzle.

[18] Other papers making small changes to the distributional assumptions include Geweke (2000), Tsionas (2005), and Labadie (1989).

3.2 Bonds in Positive Net Supply

Ours is not the first model to provide a minimum consumption level via riskless bonds in positive net supply. In particular, Cochrane, Longstaff, and Santa-Clara (2008), in the original "two-trees" model, consider an example (Section 2.8) where one tree has a riskless dividend (though this is not the main focus of their paper). However, because they assume log utility, they are unable to consider parameter values beyond the breakpoint, so all of their economies converge to a standard one-tree economy as one of the trees gets large. Heaton and Lucas (1996) consider agents with general CRRA utility who can trade stocks and bonds and face stochastic labor income. Although they mostly assume bonds are in zero net supply, they do also consider one example with bonds in positive net supply (Section IV F). They find (p. 473) that this can have a significant impact on prices and expected returns, but their solution technique (approximating the true continuous-state model with a discrete-state Markov chain) rules out extremely low consumption states, so they are unable to address the issues studied here.

3.3 Finite Time Horizons

The standard model is not defined above the breakpoint in the infinite-horizon setting. It is, however, well defined above the breakpoint when the time horizon is finite, with the same low market risk premium, $r_e - r_s = \gamma\sigma^2$, as below the breakpoint. Similarly, it is straightforward to show that the MC economy with a long but finite horizon converges to the MC economy with infinite horizon.[19] Here, by convergence we mean that given any $z > 0$, there is a large but finite T such that the finite-horizon MC economy behaves in a manner similar to the infinite-horizon economy with the same z.

How can the results then be so different? We argue that it is the standard model that behaves strangely above the breakpoint. The price function in the finite-horizon case is $P(t, D) = \frac{1-e^{-\eta(T-t)}}{\eta} D$. Below the breakpoint, this converges to $\frac{D}{\eta}$ for large T. Above the breakpoint, on the other hand, the price explodes as time to maturity increases. The low risk premium then comes from the fact that $\frac{P_t}{P} \approx \eta$, i.e., there is a large expected price decrease at each point in time when η is negative. This decrease is driven by the low-state digital options we discussed previously. These claims are extremely valuable for long time horizons, but their value decreases very quickly over time when the terminal date approaches, since the risk that these states will ever be reached decreases rapidly. The behavior of the entire tree's value is driven in large part by the extreme behavior of these low-state digital options. Since such negative expected returns with time horizon do not seem to be present in practice, we conclude that the standard model also provides a poor characterization above the breakpoint with finite horizons.

19 The convergence follows much easier than in the standard model since, for $\gamma > 1$ and $B > 0$, the utility function and its derivative are bounded below and above for all states of the world.

3.4 Relation to Literature on Bubbles

Price-dividend ratios in the MC model are nonstationary beyond the breakpoint. In fact, the convex price function is similar for large D to what occurs in the rational-bubble models that have been introduced to explain the excess-volatility puzzle (see, for example, Froot and Obstfeld 1991). In fact, our mechanism leading to excess volatility is technically similar to the intrinsic (rational) bubble mechanism used in Froot and Obstfeld (1991). The standard way of introducing rational bubbles in an infinite-horizon economy is to ignore transversality conditions (see, for example, Ingersoll 1987; Froot and Obstfeld 1991; Gilles and LeRoy 1997). Within our setting, allowing for rational bubbles would amount to changing the pricing function (3) by adding a non-zero rational-bubble term to the formula.

Without transversality conditions, there are multiple pricing functions consistent with rational pricing. As shown, e.g., in Froot and Obstfeld (1991), in a constant discount rate and investment opportunity setting, the bubble solutions take the form cD^{α} for some $\alpha > 1$, as opposed to the no-bubble solutions, which have $\alpha = 1$. Thus, these rational bubbles have the same functional form as our price function above the breakpoint, and are also nonstationary. In the MC economy, however, even though price-dividend ratios are nonstationary, there is no bubble, since the discounted cash flow formula (3) prices all assets in the economy. In fact, as noted already in Cochrane (1992), Appendix B, even with stationary distributions for consumption growth, price-dividend ratios need not be stationary. Thus, although the price functions have similar forms in the MC economy and in the rational-bubble literature, the underlying economic reason is very different.

The empirical literature that has tested for explosive stock market price dynamics has produced mixed results. For example, Diba and Grossman (1988) use a cointegration-augmented Dickey-Fuller test to conclude that prices are not explosive, a conclusion that is supported by Cochrane (1992). On the other hand, West (1987) and Froot and Obstfeld (1991) do find evidence for explosive price dynamics, findings that are also supported by Engsted (2006), who uses a cointegrated VAR method. In the MC model, the price-dividend ratios explode quite slowly and may therefore be hard to detect. In our numerical example, for example, it takes about 65 years for price-dividend ratios to double. This compares with an observed increase in the market price-dividend ratio of 3.2 times during the 65 years between 1943 and 2008.[20]

3.5 Relation to Hansen-Jagannathan Bounds

We have developed our results with respect to the market risk premium. In other words, our analysis has rested on the assumption that the equity portfolio makes up the whole market portfolio. This is the formulation developed in

[20] This calculation is based on annual price and dividend data obtained from Robert Shiller's website, http://www.econ.yale.edu/ shiller/data.htm.

Mehra and Prescott (1985) and many other papers. With that formulation, it is shown that, all else equal, the risk premium is much higher in the MC model than what seems to be implied by the standard model.

An alternative approach to the equity premium is given in Hansen and Jagannathan (1991), in which it is described as a bound on the Sharpe ratio of the equity portfolio. This bound puts restrictions on the SDF in the economy, whereas the market model puts joint restrictions on the SDF and the price function. Since the risk premium in our approach increases due to a more volatile price function, our approach therefore has less to say about the Hansen-Jagannathan bounds. It does have two implications, though. First, the interpretation of a high equity volatility differs from that in the standard model. In the standard model, a high equity volatility implies that the equity market is a highly leveraged claim on consumption. This is not the case in the MC model, in which the high volatility is introduced because of the convex price function. Second, since the price function is nonlinear, the unconditional correlation between consumption and equity returns may be low even though the two processes are instantaneously perfectly correlated. In fact, it follows from Figure 1 that for low D, equity and consumption are perfectly negatively correlated, which will decrease the unconditional correlation and, in turn, artificially make the required risk premium look higher than it actually is (for further analysis of this argument, see Berk and Walden 2009).

3.6 More General Utility

The focus of this article is on standard time-separable expected utility. One may wonder what the results would be in a model in which a more general utility specification is used. Specifically, it is well known that the standard time-separable expected utility specification jointly restricts risk aversion, γ, and the elasticity of intertemporal substitution (EIS), such that the inverse of the EIS, ψ, is equal to γ. If the representative investor has stochastic differential utility, ψ and γ may not be the same, raising the question of whether ψ or γ determines the breakpoint. It turns out that the breakpoint depends on both ψ and γ in this case. For example, under the Kreps-Porteus stochastic differential utility (Duffie and Epstein 1992), the breakpoint occurs at $\rho + (\psi - 1)\left(\mu - (\gamma - 1)\frac{\sigma^2}{2}\right)$, as follows from the analysis in Roche (2001; Equation 2.3) and Bhamra, Kuehn, and Strebulaev (2010). Both the EIS and risk-aversion parameter, therefore, contribute to the breakpoint under stochastic differential utility.

3.7 Generalizations

For simplicity, we have derived our results in a two-trees framework with one risk-free, constant-size tree. The results, however, are much more general. As long as there is a lower bound on consumption (which could grow deterministically at some rate), and the risk of ending up in these low-consumption states is bounded below by an i.i.d. growth process that, given γ, is above the breakpoint, similar results apply. The general consumption process could, for example, contain mean-reverting growth, as well as long-term i.i.d. growth. To fix ideas, we illustrate one such generalization and show how a convex price function arises under general conditions.

Proposition 7. Consider an exchange economy, with a representative agent with CRRA expected utility with risk-aversion coefficient $\gamma > 1$ and personal discount rate $\rho > 0$, in which the consumption is $C_t = f(D_t)$, where $D_t = e^{s_t}$, and where $f : \mathbb{R}_+ \to \mathbb{R}_+$ is a continuous, increasing function, such that for large d, $c_0 d \le f(d) \le c_1 d$, for some constants $0 < c_0 \le c_1 < \infty$.

For the stochastic process, $s_t \in \mathbb{R}$, define the c.d.f. $F(s, t|s_*, \mathcal{I}) = \mathbb{P}(s_t \le s|s_0 = s_*, \mathcal{I})$, where $\mathcal{I}$ captures the information known about s_t at $t = 0$. Assume that the following condition is satisfied:

$$\exists \mu, \sigma > 0,\ \bar{t} \ge 0,\ \underline{s},\ \bar{s},$$
$$\text{such that } \forall t \ge \bar{t},\ s_* \ge \bar{s}:\quad F(\underline{s}, t|s_*, \mathcal{I}) \ge \Phi\left(\frac{\underline{s} - s_* - \mu t}{\sigma\sqrt{t}}\right). \quad (24)$$

Here, Φ is the cumulative normal distribution function. Further, assume that the economy is beyond the breakpoint, i.e., that $\alpha = \gamma - \frac{\mu + \sqrt{\mu^2 + 2\rho\sigma^2}}{\sigma^2} > 1$.

Then,

(i) If $f(x) \le c_2 x$ in a neighborhood of $x = 0$, for some constant $c_2 \ge 0$, then there is no equilibrium in the economy.

(ii) If $f(0) > 0$, then in any equilibrium the price of the market satisfies $P(C_0) \ge c_3 C_0^{\alpha}$, for some constant $c_3 > 0$.

Equation (24) states that for large D_0 (i.e., for $D_0 \ge e^{\bar{s}}$) and large t (i.e., for $t \ge \bar{t}$), the risk of ending up in low-consumption states ($s_t \le \underline{s}$) is at least as high as if s were a constant coefficient Brownian motion with growth rate μ and volatility σ.[21]

Example 1. The MC economy is a special case of Proposition 7, in which $f(x) = 1 + x$, and $s_t \sim N(s_0 + \mu t, \sigma^2 t)$. It therefore satisfies (24) as an equality for all $\bar{t} > 0$, $\underline{s}$, and $\bar{s}$.

Moreover,

Example 2. Consider an MC economy with $f(x) = 1 + x$ and a mean-reverting growth process,

$$ds_t = \mu_t\, dt + \sigma\, d\omega_1,$$

[21] While Proposition 7 generalizes our results to different probability distributions for consumption, pricing depends only on the product of the p.d.f. and marginal utility at each possible consumption value, so our results could also be extended to more general utility functions.

$$d\mu_t = \beta(\mu - \mu_t)\,dt + \sigma_\mu\,d\omega_2,$$

where μ, σ, σ_μ, and β are positive constants and where $cov(d\omega_1, d\omega_2) = \rho\,dt$. Similar processes are, for example, assumed in Kim and Omberg (2002), Wachter (2002), and Bansal and Yaron (2004), and it is well known that s_t is normally distributed,

$$s_t \sim N\Bigg(s_0 + \mu t + \frac{\mu_0 - \mu}{\beta}\left(1 - e^{-\beta t}\right), \frac{\sigma_\mu^2 + 2\sigma_\mu\beta\rho\sigma + \beta^2\sigma^2}{\beta^2}t - \frac{3\sigma_\mu^2 + 4\sigma_\mu\beta\rho\sigma}{2\beta^3} + \frac{2e^{-\beta t}}{\beta^3}\left(\sigma_\mu^2 + \sigma_\mu\beta\rho\sigma\right) - \frac{e^{-2\beta t}\sigma_\mu^2}{2\beta^3}\Bigg).$$

Therefore, for large t, (24) is satisfied with $\sigma^2 \stackrel{\text{def}}{=} \frac{\sigma_\mu^2+2\sigma_\mu\beta\rho\sigma+\beta^2\sigma^2}{\beta^2}$. Here, $\sigma > 0$, as long as $\rho > -1$ or $\sigma_\mu \neq \beta\sigma$. From Proposition 7, it therefore follows that similar price dynamics occur beyond the breakpoint in the MC economy with a mean-reverting growth process.

Thus, our theory is really about minimum consumption levels in exchange economies, not about specific tree economies. In particular, referring back to the discussion after Equation (17), this result implies that if the representative investor has a low discount rate and believes that growth will slow down some time (arbitrarily far) in the future, then the effective equity premium for high D will still be approximately $\gamma^2\sigma^2$, regardless of the value of μ today (since it is only the asymptotic growth that matters). With this line of reasoning, the observed risk premium in the example we have studied throughout this article would be matched by $\gamma = \sqrt{5\%/4\%^2} = 5.6$ (instead of $\gamma = 12.25$ needed when the expected growth rate is constant). Further, if we use the numbers in Weitzman (2007)—a risk premium of 6% and consumption volatility of 2%—the risk premium is matched by $\gamma = \sqrt{6\%/2\%^2} = 12.25$ (compared with $\gamma = 6\%/2\%^2 = 150$, obtained in Weitzman 2007 under the assumption that $r_e - r_s = \gamma\sigma^2$). Thus, a high risk premium may be a sign that the economy will not be able to continue to grow fast in the long run.

4. Concluding Remarks

We have established that if risk aversion is sufficiently high, the stochastic discount factor in a simple one-tree exchange economy with minimum consumption can be a convex function of the dividend (and hence consumption) stream. This immediately leads to explosive price-dividend ratios, excess volatility, modest interest rates, and risk premia that are in line with those observed.

Intuitively, there are two main channels through which future low-consumption states affect how the representative agent values the market. The first is how the representative agent currently values these low states, and is therefore captured by the difference between marginal utility at current consumption and at the low-consumption states; the higher the current consumption, the greater the difference. Further, since marginal utilities are convex functions of consumption (when risk aversion is greater than 1), this channel also makes current market prices convex in consumption. The second channel is how likely the representative agent is to hit one of these low states; the higher her current consumption, the lower the risk that the low-consumption states will ever be reached (and the longer it will take if they are reached). Below the so-called breakpoint, the second effect outweighs the first, which means that the influence of the consumption provided in low-consumption states on the current price becomes negligible when current consumption is high. This corresponds to the standard model, in which the value of the agent's consumption stream is essentially linear in that consumption. However, when risk aversion is high enough to be above the breakpoint, the first effect dominates: The value of being able to consume in the low-consumption states *increases* convexly as current consumption grows. There is a ready analogy to this intuition in the rare-disaster literature; while there are no "disasters" in this framework, the existence of low-consumption states completely changes the properties of the model above the breakpoint.

Two immediate conclusions can be drawn from our work. First, the standard long-horizon one-tree model with a CRRA representative investor and a lognormal consumption process is highly sensitive to small perturbations, especially when risk aversion is high. In short, the framework is not robust. Second, an economically plausible assumption that is quite innocuous (subsistence consumption) renders predictions that are more in accord with empirical work. Of course, this one augmentation does not solve all puzzles; the short-term risk-free rate is still too low and consumption volatility a bit too high, as is the coefficient of risk aversion. However, we find it fascinating that such a small modification of the classic workhorse consumption model can improve the "fit" so significantly.

Finally, and more broadly, our results indicate that there is yet more to learn about the effect of the consumption process on asset prices. Because consumption (as opposed to utility) is observable, exhausting the implications of tractable models with plausible consumption streams presents a fruitful research agenda.

Proofs

Proof of Propositions 1 and 2:

Starting with Proposition 2, we have

$$P(D_0) = (1 + D_0)E\left[\int_0^{\infty} e^{-\rho s}\left(\frac{1 + D_0}{1 + D_t}\right)^{\gamma - 1} ds\right]$$

$$
\begin{aligned}
&= (1+D_0)^{\gamma} \int_{-\infty}^{\infty}\int_0^{\infty} e^{-\rho s}\left(\frac{1}{1+D_0e^y}\right)^{\gamma-1} \frac{e^{-\frac{(y-\mu s)^2}{2\sigma^2 s}}}{\sqrt{2\pi\sigma^2 s}}\,ds\,dy \\
&= (1+D_0)^{\gamma} \int_{-\infty}^{\infty}\left(\frac{1}{1+D_0e^y}\right)^{\gamma-1} \frac{e^{\frac{\mu y-q|y|}{\sigma^2}}}{q}\,dy \\
&= (1+D_0)^{\gamma}\left[\int_{-\infty}^{0}\left(\frac{1}{1+D_0e^y}\right)^{\gamma-1}\frac{e^{y\kappa}}{q}\,dy\right. \\
&\qquad\left. + \int_0^{\infty}\left(\frac{1}{1+D_0e^y}\right)^{\gamma-1}\frac{e^{y(\kappa-2q/\sigma^2)}}{q}\,dy\right] \\
&= (1+D_0)^{\gamma}\frac{D_0^{-\kappa}}{q}\left[V(D_0,\kappa,2-\gamma)\right. \\
&\qquad\left. + D_0^{\frac{2q}{\sigma^2}} V\left(\frac{1}{D_0},\alpha+\frac{2q}{\sigma^2}-1,2-\gamma\right)\right],
\end{aligned}
$$

where

$$V(y,a,b) \stackrel{\text{def}}{=} \int_0^y t^{a-1}(1+t)^{b-1}\,dt \tag{25}$$

is defined for $a > 0$.

In the last step, we used the transformation $t = D_0e^y$ for the first integral. For the second integral, we rewrote $\left(\frac{1}{1+D_0e^y}\right)^{\gamma-1} = \left(\frac{D_0^{-1}e^{-y}}{D_0^{-1}e^{-y}+1}\right)^{\gamma-1}$ and then used the transformation $t = D_0^{-1}e^{-y}$ to get the expression. The function V is related to the incomplete Beta function, $B(x,a,b) \stackrel{\text{def}}{=} \int_0^x t^{a-1}(1-t)^{b-1}\,dt$ (see Gradshteyn and Ryzhik 2000), via the relation $V(x,a,b) = (-1)^a B(-x,a,b)$. However, the Beta function is complex valued for negative values, so we prefer using the real-valued function V. Also, since the Beta function and the hypergeometric function satisfy the relationship $B(x,a,b) = {}_2F_1(1-b,a,a+1,x)$, we could equivalently have expressed the formula in terms of hypergeometric functions.

This proves Proposition 2. For Proposition 1, for $z = D_0/(B+D_0)$, we have

$$
\begin{aligned}
w(z) = (B+D_0)^{\gamma-1}U(0|B,D_0) &= \frac{1}{1-\gamma}\frac{P(B,D)}{B+D} \\
&= \frac{1}{1-\gamma}\frac{B}{B+D}P\left(\frac{D_0}{B}\right) = \frac{1}{1-\gamma}zP\left(\frac{z}{1-z}\right),
\end{aligned}
$$

where the second equality holds for the CRRA utility, which follows from (3). Therefore, from Equation (11), we immediately have that for $z \in (0,1)$,

$$w(z) = \frac{z^{-\kappa}(1-z)^{1-\gamma-\kappa}}{q(1-\gamma)}\Big[V\left(\frac{1-z}{z}, \kappa, 2-\gamma\right) \tag{26}$$

$$+\left(\frac{1-z}{z}\right)^{\frac{2q}{\sigma^2}} V\left(\frac{z}{1-z}, \alpha+\frac{2q}{\sigma^2}-1, 2-\gamma\right)\Big]. \tag{27}$$

For $w(0)$ and $w(1)$, we define $\hat{w}_T(z) = E\left[\int_0^T e^{-\rho t} u\left(1-z+ze^{y_t}\right) dt\right]$, where $y_t = \log(D_t/D_0)$. Thus, $w(z) = \hat{w}_\infty(z)$. It follows immediately that $w(1) = \hat{w}_\infty(1) = \int_0^\infty \frac{e^{-\rho t}}{1-\gamma} dt = \frac{1}{\rho(1-\gamma)}$. Moreover, $\hat{w}_T(0) = \int_0^T \frac{e^{-\eta t}}{1-\gamma} dt = \frac{1-e^{-\eta T}}{\eta(1-\gamma)}$, so for $\eta > 0$, $w(0) = \hat{w}_\infty(0) = \frac{1}{\eta(1-\gamma)}$, whereas for $\eta < 0$, $\lim_{T\to\infty} \hat{w}_T(0) = -\infty$. The proposition is proved.

We note that although $\hat{w}_\infty(0) = \lim_{T\to\infty} \lim_{z\to 0} \hat{w}_T(z) = -\infty$ when $\eta < 0$, it does not immediately follow that $\lim_{z\to 0} w(z) = \lim_{z\to 0} \lim_{T\to\infty} \hat{w}_T(z)$ is equal to $-\infty$ (for example, if $\hat{w}_T(z) = -\frac{1}{zT}$, then the former expression is infinite, whereas the second is zero). However, the latter result follows, since $\hat{w}_T(z)$ is decreasing in T for arbitrary $z \in [0,1]$, and $\hat{w}_T(z)$ is continuous in z for arbitrary finite T. Specifically, for an arbitrary constant, $k > 0$, it follows that for T^* large enough, $\hat{w}_{T^*}(0) \le -2k$, and because of the continuity in z, $\hat{w}_{T^*}(z) \le -k$ for all $z \le z^*$, for some $z^* > 0$. Therefore, $\hat{w}_\infty(z) \le \hat{w}_{T^*}(z) \le -k$ for all $z \le z^*$, and since k was arbitrary, it is indeed the case that $\lim_{z\to 0} w(z) = \lim_{z\to 0} \hat{w}_\infty(z) = -\infty$. ■

Proof of Proposition 3: We first study the case when $\alpha > 1$. We look at $\frac{P(D)}{(1+D)^\alpha}$, for large D. From (11), it follows that

$$\frac{P(D)}{(1+D)^{\gamma-\kappa}} = \left(\frac{1+D}{D}\right)^{\kappa} \frac{1}{q}\Big[V(D,\kappa,2-\gamma) + D^{\frac{2q}{\sigma^2}} V\left(\frac{1}{D}, \alpha+\frac{2q}{\sigma^2}-1, 2-\gamma\right)\Big]$$

$$= \frac{1+o(1)}{q}\left[\int_0^D t^{\kappa-1}(1+t)^{1-\gamma}\,dt + D^{\frac{2q}{\sigma^2}} \int_0^{1/D} t^{\alpha+\frac{2q}{\sigma^2}-2}(1+t)^{1-\gamma}\,dt\right]. \tag{28}$$

Here, $\lim_{D\to\infty} o(1) = 0$. Since $\kappa > 0$ and $\gamma - \kappa > 1$, $\lim_{D\to\infty} \int_0^D t^{\kappa-1}(1+t)^{1-\gamma}\,dt = c_1$, where $0 < c_1 < \infty$. Moreover,

$$D^{\frac{2q}{\sigma^2}} \int_0^{1/D} t^{\alpha+\frac{2q}{\sigma^2}-2}(1+t)^{1-\gamma}\,dt = D^{\frac{2q}{\sigma^2}} \int_0^{1/D} t^{\alpha+\frac{2q}{\sigma^2}-2}\,dt$$

$$= c_2 D^{\frac{2q}{\sigma^2}} D^{-\left(\alpha+\frac{2q}{\sigma^2}-1\right)} = c_2 D^{1-\alpha},$$

which converges to zero for large D. The finiteness of the integral is ensured, since $\alpha + \frac{2q}{\sigma^2} - 2 > -1$.

Thus, for large D, the expression converges to $\frac{c_1}{q}$.

For $\alpha < 1$, we use that

$$\begin{aligned}\frac{P(D)}{1+D} &= (1+D)^{\gamma-1}\frac{D^{-\kappa}}{q}\Bigg[V(D,\kappa,2-\gamma) \\ &\qquad + D^{\frac{2q}{\sigma^2}}V\left(\frac{1}{D},\alpha+\frac{2q}{\sigma^2}-1,2-\gamma\right)\Bigg] \\ &= \frac{1+o(1)}{q}D^{\gamma-1-\kappa}\Bigg[\int_0^D t^{\kappa-1}(1+t)^{1-\gamma}\,dt \\ &\qquad + D^{\frac{2q}{\sigma^2}}\int_0^{1/D} t^{\alpha+\frac{2q}{\sigma^2}-2}(1+t)^{1-\gamma}\,dt\Bigg].\end{aligned}$$

For the first term, we note that

$$\begin{aligned}\int_0^D t^{\kappa-1}(1+t)^{1-\gamma}\,dt &= \frac{D^\kappa}{\kappa}\,{}_2F_1(\gamma-1,\kappa,1+\kappa,-D) \\ &= \frac{D^{\kappa-\gamma+1}}{\kappa}\,{}_2F_1\left(\gamma-1,1,1+\kappa,\frac{D}{D+1}\right).\end{aligned}$$

For large D, the first term therefore converges to

$$\begin{aligned}\frac{1}{q\kappa}\,{}_2F_1(\gamma-1,1,1+\kappa,1) &= \frac{\Gamma(\kappa+1)\Gamma(1+\kappa-\gamma)}{\Gamma(2+\kappa-\gamma)\Gamma(\kappa)} = \frac{1}{q(1-\gamma+\kappa)} \\ &= -\frac{\sigma^2}{q}\times\frac{1}{(\gamma-1)\sigma^2-\mu-q}.\end{aligned}$$

For the second term, we note that

$$\begin{aligned}&\int_0^{1/D} t^{\alpha+\frac{2q}{\sigma^2}-2}(1+t)^{1-\gamma}\,dt \\ &\quad = \frac{D^{1-\alpha-\frac{2q}{\sigma^2}}\sigma^2}{2q+(\alpha-1)\sigma^2}\,{}_2F_1\left(\gamma-1,\alpha+\frac{2q}{\sigma^2}-1,\alpha+\frac{2q}{\sigma^2},-\frac{1}{D}\right).\end{aligned}$$

Since ${}_2F_1\left(\gamma-1,\alpha+\frac{2q}{\sigma^2}-1,\alpha+\frac{2q}{\sigma^2},0\right) = 1$, and $\alpha = \gamma-\kappa$, the second term therefore converges to

$$\frac{\sigma^2}{q(2q+(\alpha-1)\sigma^2)} = \frac{\sigma^2}{q}\times\frac{1}{(\gamma-1)\sigma^2-\mu+q}.$$

Thus,

$$
\begin{aligned}
\lim_{D\to\infty}\frac{P(D)}{1+D} &= \frac{\sigma^2}{q}\times\left(\frac{1}{(\gamma-1)\sigma^2-\mu+q}-\frac{1}{(\gamma-1)\sigma^2-\mu-q}\right)\\
&= \frac{\sigma^2}{q}\times\frac{2q}{((\gamma-1)\sigma^2-\mu)^2-q^2}\\
&= \frac{1}{\rho+\mu(\gamma-1)-(\gamma-1)^2\frac{\sigma^2}{2}}\\
&= \frac{1}{\eta}. \qquad \blacksquare
\end{aligned}
$$

Proof of Proposition 4. It is easy to see from (28) of Proposition 3 that for large D, when $\alpha > 1$, $\frac{d}{dD}\left[\frac{P(D)}{(1+D)^\alpha}\right]$ converges to 0, as does $\frac{d^2}{dD^2}\left[\frac{P(D)}{(1+D)^\alpha}\right]$. Therefore, in this case, $P' = \alpha(1+o(1))c_2 D^{\alpha-1}$, and $P'' = \alpha(\alpha-1)(1+o(1))$ $c_2 D^{\alpha-2}$ for large D, and it follows that $\frac{P'(D)D}{P(D)}$ converges to α and $\frac{P''(D)D^2}{P(D)}$ converges to $\alpha(\alpha-1)$. (ii) then follows from standard Itô calculus.

For $\alpha < 1$, an identical argument for $\frac{P(D)}{D}$ proves (i). $\blacksquare$

Proof of Proposition 5. Defining $F(x) = e^{x^2}\mathrm{Erfc}(x)$, where Erfc is the error function, $\mathrm{Erfc}(x) = (\sqrt{\pi})^{-1}\int_z^\infty e^{-t^2}\,dt$, we show that Equation (20) can be expressed in the following form:

$$
\begin{aligned}
P^\tau = \frac{(1+e^d)^\gamma e^{-\rho\tau-(d+\mu\tau)^2/(2\sigma^2\tau)}}{2}\times\lim_{\epsilon\searrow 0}\\
\sum_{n=0}^{\infty}(-1)^n e^{-\epsilon n}a_n\left(F\left(\frac{\epsilon+d+\mu\tau+n\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right)\right.\\
\left.+F\left(\frac{\epsilon-d-\mu\tau+(n+\gamma)\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right)\right).
\end{aligned}
\tag{29}
$$

Here,

$$a_n = \frac{\Gamma(\gamma+n)}{\Gamma(\gamma)\Gamma(n+1)},$$

where $\Gamma(x) = \int_0^\infty e^{-t}t^{x-1}\,dt$, which reduces to $a_n = \binom{n+\gamma-1}{\gamma}$ when γ is integer-valued.

(i) The function $\frac{1}{(1+z)^\gamma}$ is analytic in the complex plane, $|z| < 1$, and can therefore be expanded in the power expansion

$$\frac{1}{(1+z)^\gamma} = \sum_{n=0}^{\infty}(-1)^n a_n z^n.$$

For $y < 0$, we use this expansion to get $1/\left(1+e^y\right)^\gamma = \sum_{n=0}^{\infty}(-1)^n a_n e^{ny}$, and for $y > 0$, we get a similar expansion $1/\left(1+e^y\right)^\gamma = e^{-\gamma y} \sum_{n=0}^{\infty}(-1)^n a_n e^{-ny}$.

Now, from Equation (18), it follows that

$$
\begin{aligned}
\sqrt{2\pi\sigma^2\tau}\frac{P^\tau}{\left(1+e^d\right)^\gamma e^{-\rho\tau}} &= \int_{-\infty}^{\infty}\frac{e^{-(y-d-\mu\tau)^2/\left(2\sigma^2\tau\right)}}{\left(1+e^y\right)^\gamma}\,dy \\
&= \left(\int_{-\infty}^{-\epsilon}+\int_{\epsilon}^{\infty}+\int_{-\epsilon}^{\epsilon}\right)\frac{e^{-(y-d-\mu\tau)^2/\left(2\sigma^2\tau\right)}}{\left(1+e^y\right)^\gamma}\,dy \\
&= \int_{-\infty}^{-\epsilon}\frac{e^{-(y-d-\mu\tau)^2/\left(2\sigma^2\tau\right)}}{\left(1+e^y\right)^\gamma}\,dy \\
&\quad +\int_{\epsilon}^{\infty}\frac{e^{-(y-d-\mu\tau)^2/\left(2\sigma^2\tau\right)}}{\left(1+e^y\right)^\gamma}\,dy + O(\epsilon) \\
&= \int_{-\infty}^{0}\frac{e^{-(y-\epsilon-d-\mu\tau)^2/\left(2\sigma^2\tau\right)}}{\left(1+e^{y-\epsilon}\right)^\gamma}\,dy \\
&\quad +\int_{0}^{\infty}\frac{e^{-(y+\epsilon-d-\mu\tau)^2/\left(2\sigma^2\tau\right)}}{\left(1+e^{y+\epsilon}\right)^\gamma}\,dy \\
&\quad +O(\epsilon)
\end{aligned}
\tag{30}
$$

for all $\epsilon > 0$ and $y < 0$. However, since

$$
\begin{aligned}
\frac{e^{-(y-d-\epsilon-\mu\tau)^2/\left(2\sigma^2\tau\right)}}{\left(1+e^{y-\epsilon}\right)^\gamma} &= \sum_{n=0}^{\infty}(-1)^n a_n e^{-(y-\epsilon-d-\mu\tau)^2/\left(2\sigma^2\right)+n(y-\epsilon)} \\
&= \sum_{n=0}^{\infty}(-1)^n a_n e^{-\frac{\epsilon}{2}n} e^{(y-\epsilon-d-\mu\tau)^2/\left(2\sigma^2\right)\tau+n\left(y-\frac{\epsilon}{2}\right)},
\end{aligned}
$$

the first term is equal to

$$
\int_{-\infty}^{0}\sum_{n=0}^{\infty}(-1)^n a_n e^{-\frac{\epsilon}{2}n} e^{-(y-\epsilon-d-\mu\tau)^2/\left(2\sigma^2\tau\right)+n\left(y-\frac{\epsilon}{2}\right)}\,dy. \tag{31}
$$

Now, define $g_{M,\epsilon}(y) = \sum_{n=0}^{M} a_n(-1)^n e^{-\epsilon n/2} e^{-(y-\epsilon-d-\mu\tau)^2/\left(2\sigma^2\tau\right)+n\left(y-\frac{\epsilon}{2}\right)}$, $y < 0$, $M \in \mathbb{N}$, and $h_\epsilon(y) = e^{-(y-\epsilon-d-\mu\tau)^2/\left(2\sigma^2\tau\right)}$. Then, since $a_n \sim Cn^\gamma$ for large n, it is clear that $\sup_{n\geq 0} a_n e^{-\epsilon n/2} = C < \infty$. Therefore,

$$
\left|g_{M,\epsilon}(y)\right| \leq C\sum_{n=0}^{M} e^{-(y-\epsilon-d-\mu\tau)^2/\left(2\sigma^2\tau\right)+n\left(y-\frac{\epsilon}{2}\right)}
$$

$$
\begin{aligned}
&\leq C\sum_{n=0}^{\infty} e^{-(y-\epsilon-d-\mu\tau)^2/(2\sigma^2\tau)+n(y-\frac{\epsilon}{2})} \\
&= C\frac{e^{-(y-\epsilon-d-\mu\tau)^2/(2\sigma^2\tau)}}{1-e^{-\epsilon/2}e^{y}} \\
&\leq C'_\epsilon e^{-(y-\epsilon-d-\mu\tau)^2/(2\sigma^2\tau)} = C'_\epsilon h_\epsilon(y).
\end{aligned}
$$

Clearly, $\int_{-\infty}^{0} C'_\epsilon h_\epsilon(y)\,dy < \infty$, and therefore the dominated convergence theorem implies that $\int_{-\infty}^{0} \lim_{n\to\infty} g_{M,\epsilon}(y)\,dy = \lim_{n\to\infty}\int_{-\infty}^{0} g_{M,\epsilon}(y)\,dy$, i.e.,

$$
\begin{aligned}
&\int_{-\infty}^{0} \frac{e^{-(y-\epsilon-d-\mu\tau)^2/(2\sigma^2\tau)}}{(1+e^{y-\epsilon})^{\gamma}}\,dy \\
&\quad= \sum_{n=0}^{\infty}\int_{-\infty}^{0} (-1)^n a_n e^{-\frac{\epsilon}{2}n} e^{-(y-\epsilon-d-\mu\tau)^2/(2\sigma^2\tau)+n(y-\frac{\epsilon}{2})}\,dy \\
&\quad= \sum_{n=0}^{\infty}(-1)^n a_n e^{-\frac{\epsilon}{2}n}\int_{-\infty}^{0} e^{-(y-\epsilon-d-\mu\tau)^2/(2\sigma^2\tau)+n(y-\frac{\epsilon}{2})}\,dy.
\end{aligned}
$$

Define $F(x) = e^{x^2}\mathrm{Erfc}(x)$, where Erfc is the error function $\mathrm{Erfc}(x) = (\sqrt{\pi})^{-1}\int_z^{\infty} e^{-t^2}\,dt$ (see Abramowitz and Stegun 1964). Then, since

$$
\begin{aligned}
&\frac{1}{\sqrt{2\pi\sigma^2\tau}}\int_{-\infty}^{0} e^{-(y-\epsilon d-\mu\tau)^2/(2\sigma^2\tau)+n(y-\epsilon/2)}\,dy \\
&\quad= \frac{1}{2}e^{n(\epsilon/2+d+\mu\tau)+n^2\tau\sigma^2/2}\mathrm{Erfc}\left(\frac{\epsilon+d+\mu\tau+n\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right) \\
&\quad= \frac{e^{-n\frac{\epsilon}{2}}e^{-(\epsilon+d+\mu\tau)^2/(2\sigma^2\tau)}}{2}F\left(\frac{\epsilon+d+\mu\tau+n\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right),
\end{aligned}
$$

it follows that

$$
\begin{aligned}
&\frac{1}{\sqrt{2\pi\sigma^2\tau}}\int_{-\infty}^{0} \frac{e^{-(y-\epsilon-d-\mu\tau)^2/(2\sigma^2\tau)}}{(1+e^{y-\epsilon})^{\gamma}}\,dy \\
&= \frac{1}{2}\sum_{n=0}^{\infty}(-1)^n a_n e^{-\epsilon n} e^{-(\epsilon+d+\mu\tau)^2/(2\sigma^2\tau)} F\left(\frac{\epsilon+d+\mu\tau+n\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right) \\
&= (1+O(\epsilon))\frac{1}{2}e^{-(d+\mu\tau)^2/(2\sigma^2\tau)} \\
&\quad\times\sum_{n=0}^{\infty}(-1)^n a_n e^{-\epsilon n} F\left(\frac{\epsilon+d+\mu\tau+n\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right).
\end{aligned}
$$

An identical argument for the $\int_0^\infty \frac{e^{-(y+\epsilon-d-\mu\tau)^2/(2\sigma^2\tau)}}{(1+e^{y+\epsilon})^\gamma}\,dy$ term leads to

$$
\begin{aligned}
\frac{1}{\sqrt{2\pi\sigma^2\tau}} \int_0^\infty \frac{e^{-(y+\epsilon-d-\mu\tau)^2/(2\sigma^2\tau)}}{(1+e^{y+\epsilon})^\gamma}\,dy
&= \frac{1}{2}\sum_{n=0}^{\infty}(-1)^n a_n e^{-\epsilon n} e^{-(\epsilon+d+\mu\tau)^2/(2\sigma^2\tau)} \\
&\quad \times F\left(\frac{\epsilon-d-\mu\tau+(n+\gamma)\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right) \\
&= (1+O(\epsilon))e^{-(d+\mu\tau)^2/(2\sigma^2\tau)}\sum_{n=0}^{\infty}(-1)^n a_n e^{-\epsilon n} \\
&\quad \times F\left(\frac{\epsilon-d-\mu\tau+(n+\gamma)\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right).
\end{aligned}
$$

Putting it all together in Equation (30), we get

$$
\begin{aligned}
P^\tau &= \frac{(1+e^d)^\gamma e^{-\rho\tau}}{\sqrt{2\pi\sigma^2\tau}}\left(\int_{-\infty}^0 \frac{e^{-(y-\epsilon-d-\mu\tau)^2/(2\sigma^2\tau)}}{(1+e^{y-\epsilon})^\gamma}\,dy\right. \\
&\qquad\qquad \left. + \int_0^\infty \frac{e^{-(y+\epsilon-d-\mu\tau)^2/(2\sigma^2\tau)}}{(1+e^{y+\epsilon})^\gamma}\,dy + O(\epsilon)\right) \\
&= O(\epsilon) + \frac{(1+e^d)^\gamma e^{-\rho\tau-(d+\mu\tau)^2/(2\sigma^2\tau)}}{2} \\
&\quad \times \sum_{n=0}^{\infty}(-1)^n e^{-\epsilon n} a_n \left(F\left(\frac{\epsilon+d+\mu\tau+n\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right)\right. \\
&\qquad\qquad \left. + F\left(\frac{\epsilon-d-\mu\tau+(n+\gamma)\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right)\right),
\end{aligned}
$$

and thus, as $\epsilon \searrow 0$, we get convergence to Equation (29).
The formula is straightforward to use, since $F(x) \sim 1/x$ for large x. An error analysis implies that if n terms are used in the expansion, $\epsilon \sim \log(n)/n$ should be chosen.

(ii) When $\gamma = 1$, $a_n = 1$ for all n, and we can choose $\epsilon = 0$ and still apply the dominated convergence theorem in Equation (31) to get

$$
P^\tau = \frac{(1+e^d)^\gamma e^{-\rho\tau-(d+\mu\tau)^2/(2\sigma^2\tau)}}{2}
$$

$$\times \sum_{n=0}^{\infty}(-1)^n a_n \left(F\left(\frac{d+\mu\tau+n\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right) \right.$$
$$\left. + F\left(\frac{-d-\mu\tau+(n+\gamma)\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right)\right). \quad (32)$$

(iii) For

$$\frac{d+\mu\tau}{\sigma^2\tau} = m \in \mathbb{N},$$

Equation (32) reduces to a case for which closed-form expressions exist, so

$$P^\tau = \frac{(1+e^d)e^{-\rho\tau-m^2\sigma^2\tau/2}}{2}\left(1+2\sum_{n=1}^{m-1}(-1)^n e^{n^2\sigma^2\tau/2}\right).$$

Finally, we note that since $P^\tau = e^{-r(\tau)\tau}$, where $r(\tau)$ is the time-τ spot rate, we have

$$r(\tau) = \rho + \frac{\mu^2}{2\sigma^2} + \frac{1}{\tau}\left(\log\left(-\frac{(1+e^d)^\gamma}{2}\right) + \frac{d^2}{2\sigma^2\tau} + \frac{d\mu}{\sigma^2} + \log(z)\right),$$

where $z = \lim_{\epsilon\searrow 0}\sum_{n=0}^{\infty}(-1)^n e^{-\epsilon n} a_n \left(F\left(\frac{\epsilon+d+\mu\tau+n\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right) + F\left(\frac{\epsilon-d-\mu\tau+(n+\gamma)\tau\sigma^2}{\sqrt{2\sigma^2\tau}}\right)\right)$. ■

Proof of Proposition 6. The result for r_s is standard. Using Feynman-Kac, we know that

$$P_t^\tau + \frac{1}{2}\sigma^2 z^2(1-z)^2 P_{zz}^\tau + \left[-\widehat{\mu} z(1-z) + 2\sigma^2 z(1-z)^2\right] P_z^\tau$$
$$-\left[\rho + \gamma\widehat{\mu}(1-z) - \frac{1}{2}\gamma(\gamma+1)\sigma^2(1-z)^2\right] P^\tau = 0,$$

and since $P^\tau(\tau, z) = 1$, it is clear that $P(0,z) = 1 - \left[\rho + \gamma\widehat{\mu}(1-z) - \frac{1}{2}\gamma(\gamma+1)\sigma^2(1-z)^2\right]\tau + o(\tau)$, for small τ. Since $-\log(1-s) = s + O(s^2)$ for small s, it is clear that $r_s = \lim_{\tau\searrow 0} -\frac{\log(P^\tau)}{\tau} = \rho + \gamma\hat{\mu}(1-z) - \frac{1}{2}\gamma(\gamma+1)\sigma^2(1-z)^2$.

For r_l, we proceed as follows: We have

$$P^\tau = (1+e^d)^\gamma e^{-\rho\tau}\frac{1}{\sqrt{2\pi\sigma^2\tau}}\int_{-\infty}^{\infty}\frac{e^{-(y-\mu\tau)^2/(2\sigma^2\tau)}}{(1+e^{d+y})^\gamma}\,dy$$
$$= (1+e^d)^\gamma e^{-\rho\tau}\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\frac{e^{-x^2/2}}{(1+e^d e^{x\sigma\sqrt{\tau}+\mu\tau})^\gamma}\,dx.$$

We study the behavior of $\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\frac{e^{-x^2/2}}{(1+e^{d}e^{x\sigma\sqrt{\tau}+\mu\tau})^{\gamma}}\,dx$ for large τ. We decompose

$$\begin{aligned}\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}&\frac{e^{-x^2/2}}{(1+e^{d}e^{x\sigma\sqrt{\tau}+\mu\tau})^{\gamma}}\,dx\\&=\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}\frac{e^{-x^2/2}}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^{\gamma}}\,dx\\&\quad+\frac{1}{\sqrt{2\pi}}\int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty}\frac{e^{-x^2/2}}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^{\gamma}}\,dx.\end{aligned}\tag{33}$$

We prove the results for r_l by studying the first and second terms in Equation (33) separately for the two cases $\mu\le\gamma\sigma^2$ and $\mu>\gamma\sigma^2$, respectively. By showing that the first term behaves like $e^{-\frac{\mu}{2\sigma^2}\tau}$ for large τ for all μ, whereas the second term behaves like $e^{-\frac{\mu}{2\sigma^2}\tau}$ when $\mu\le\gamma\sigma^2$ and like $e^{-(\gamma\mu-\gamma^2\sigma^2/2)\tau}$ when $\mu>\gamma\sigma^2$, the result will follow.

Since $0<e^{x\sigma\sqrt{\tau}+\mu\tau+d}\le 1$ for $x\le-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}$, we have

$$\begin{aligned}\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}&\frac{e^{-x^2/2}}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^{\gamma}}\,dx\\&=C\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}e^{-x^2/2}\,dx=C\times N\left(-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}\right),\end{aligned}$$

for some $C\in[1/2^{\gamma},1]$, where $N(\cdot)$ is the cumulative normal distribution function, $N(v)\overset{\text{def}}{=}\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{v}e^{-y^2/2}\,dy$. Now, we use

$$N(-v)=C_2\frac{e^{-v^2/2}}{v},\qquad C_2\in\frac{1}{\sqrt{2\pi}}\left[\frac{v^2}{1+v^2},1\right],\tag{34}$$

which is valid for $v\gg 0$, to get

$$\begin{aligned}\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}&\frac{e^{-x^2/2}}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^{\gamma}}\,dx\\&=C\times C_2\frac{e^{-q^2/2}}{q}=C_3\frac{e^{-\frac{\mu^2}{2\sigma^2}\tau-\frac{\mu d}{\sigma^2}-\frac{d^2}{2\sigma^2\tau}}}{q},\end{aligned}$$

where

$$C_3\in\frac{1}{\sqrt{2\pi}}\left[\frac{1}{2^{\gamma+1}},1\right],\qquad\text{and}\qquad q=\frac{\mu\tau+d}{\sigma\sqrt{\tau}}.$$

We next study the second term in Equation (33), when $\mu < \gamma\sigma^2$. First, we note that $\mu < \gamma\sigma^2$ implies that $\gamma\sigma - \frac{\mu}{\sigma} > 0$. Obviously, $\frac{1}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^\gamma} \leq e^{-\gamma(x\sigma\sqrt{\tau}+\mu\tau+d)}$, so

$$
\begin{aligned}
0 &\leq \frac{1}{\sqrt{2\pi}} \int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty} \frac{e^{-x^2/2}}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^\gamma}\, dx \\
&\leq \frac{1}{\sqrt{2\pi}} \int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty} e^{-(x^2+2x\gamma\sigma\sqrt{\tau})/2-\gamma\mu\tau-\gamma d}\, dx \\
&= \frac{1}{\sqrt{2\pi}} \int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty} e^{-(x+\gamma\sigma\sqrt{\tau})^2/2+\frac{\gamma^2\sigma^2\tau}{2}-\gamma\mu\tau-d\gamma}\, dx \\
&= e^{-\gamma d} e^{-\tau(\gamma\mu-\gamma^2\sigma^2/2)} \frac{1}{\sqrt{2\pi}} \int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}+\gamma\sigma\sqrt{\tau}}^{\infty} e^{-x^2/2}\, dx \\
&= e^{-\gamma d} e^{-\tau(\gamma\mu-\gamma^2\sigma^2/2)} \frac{1}{\sqrt{2\pi}} \int_{(\gamma\sigma-\frac{\mu}{\sigma})\sqrt{\tau}-\frac{d}{\sigma\sqrt{\tau}}}^{\infty} e^{-x^2/2}\, dx \\
&= e^{-\gamma d} e^{-\tau(\gamma\mu-\gamma^2\sigma^2/2)} N\left(-\left(\gamma\sigma-\frac{\mu}{\sigma}\right)\sqrt{\tau}+\frac{d}{\sigma\sqrt{\tau}}\right) \\
&\leq e^{-\gamma d} e^{-\tau(\gamma\mu-\gamma^2\sigma^2/2)} \frac{1}{\sqrt{2\pi}} \frac{e^{-q_2^2/2}}{q_2} \\
&= e^{-\gamma d} e^{-\tau(\gamma\mu-\gamma^2\sigma^2/2)} \frac{1}{\sqrt{2\pi}} \frac{e^{-\frac{d^2}{2\sigma^2\tau}+\gamma d-\frac{d\mu}{\sigma^2}-(\gamma^2\frac{\sigma^2}{2}-\gamma\mu+\frac{\mu^2}{2\sigma^2})\tau}}{q_2} \\
&= e^{-\frac{d^2}{2\sigma^2\tau}-\frac{d\mu}{\sigma^2}} \times \frac{1}{\sqrt{2\pi}} \frac{e^{\frac{-\mu^2}{2\sigma^2}\tau}}{q_2},
\end{aligned}
$$

where $q_2 = \left(\gamma\sigma - \frac{\mu}{\sigma}\right)\sqrt{\tau} - \frac{d}{\sigma\sqrt{\tau}}$, and we used that $\frac{1}{\sqrt{2\pi}}\int_v^\infty e^{-y^2/2}\,dy = N(-v)$, and Equation (34). Thus,

$$
\frac{1}{\sqrt{2\pi}} \int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty} \frac{e^{-x^2/2}}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^\gamma}\, dx = C_4 e^{-\frac{d^2}{2\sigma^2\tau}-\frac{d\mu}{\sigma^2}} \times \frac{e^{\frac{-\mu^2}{2\sigma^2}\tau}}{q_2},
$$

where $C_4 \in \left[0, \frac{1}{\sqrt{2\pi}}\right]$. Putting it all together, for large τ we get

$$
\begin{aligned}
P^\tau = (1+e^d)^\gamma e^{-\rho t} \Bigg(& \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}} \frac{e^{-x^2/2}}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^\gamma}\, dx \\
& + \frac{1}{\sqrt{2\pi}} \int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty} \frac{e^{-x^2/2}}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^\gamma}\, dx \Bigg)
\end{aligned}
$$

$$= (1+e^d)^\gamma e^{-\rho t} \left(C_3 \frac{e^{-\frac{\mu^2}{2\sigma^2}\tau - \frac{\mu d}{\sigma^2} - \frac{d^2}{2\sigma^2 \tau}}}{q} + C_4 e^{-\frac{d^2}{2\sigma^2\tau} - \frac{d\mu}{\sigma^2}} \times \frac{e^{\frac{-\mu^2}{2\sigma^2}\tau}}{q_2} \right)$$

$$= e^{-\left(\rho + \frac{\mu^2}{2\sigma^2}\right)\tau} (1+e^d)^\gamma e^{-\frac{\mu d}{\sigma^2} - \frac{d^2}{2\sigma^2\tau}} \left(\frac{C_3}{q} + \frac{C_4}{q_2} \right).$$

Therefore,

$$-\frac{\log(P^\tau)}{\tau} = \rho + \frac{\mu^2}{2\sigma^2} + \frac{Q(\tau)}{\tau}, \qquad \text{where}$$

$$Q(\tau) = \log\left((1+e^d)^\gamma e^{-\frac{\mu d}{\sigma^2} - \frac{d^2}{2\sigma^2\tau}} \left(\frac{C_3}{q} + \frac{C_4}{q_2} \right) \right).$$

Now, $Q(\tau) = \log\left((1+e^d)^\gamma\right) - \frac{\mu d}{\sigma^2} - \frac{d^2}{2\sigma^2\tau} + \log\left(\frac{C_3}{q} + \frac{C_4}{q_2}\right)$, and since $C_3 \in \frac{1}{\sqrt{2\pi}}\left[\frac{1}{2^{\gamma+1}}, 1\right]$, $C_4 \in \left[0, \frac{1}{\sqrt{2\pi}}\right]$, $q = \frac{\mu\tau + d}{\sigma\sqrt{\tau}}$, and $q_2 = \left(\gamma\sigma - \frac{\mu}{\sigma}\right)\sqrt{\tau} - \frac{d}{\sigma\sqrt{\tau}}$, it follows that $|Q(\tau)| = o(\tau)$ for large τ, i.e., that $\lim_{\tau\to\infty} \frac{|Q(\tau)|}{\tau} = 0$. From this, it immediately follows that $\lim_{\tau\to\infty} -\frac{\log(P^\tau)}{\tau} = \rho + \frac{\mu^2}{2\sigma^2}$.

We now consider the case when $\mu > \gamma\sigma^2$ and define $\upsilon = \mu/\sigma - \gamma\sigma > 0$. We first note that $\frac{\mu^2}{2\sigma^2} \geq \gamma\mu - \gamma^2\sigma^2/2$, since $\mu^2/(2\sigma^2) - \gamma\mu + \gamma^2\sigma^2/2 = \frac{1}{2\sigma^2}(\mu - \gamma\sigma^2)^2 \geq 0$. Thus, since the $\int_{-\infty}^{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}} \frac{e^{-x^2/2}}{(1+e^d e^{x\sigma\sqrt{\tau}+\mu\tau})^\gamma}\, dx$-term in Equation (33) behaves like $e^{-\tau\times\mu^2/(2\sigma^2)}$ for large τ, if the $\int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty} \frac{e^{-x^2/2}}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^\gamma}\, dx \sim e^{-\tau(\mu\gamma - \gamma^2\sigma^2/2)}$ for large τ, then the result we wish to prove follows, since it is always the case that $c_1 e^{-\alpha_1\tau} + c_2 e^{-\alpha_2\tau} \sim e^{-\min(\alpha_1,\alpha_2)\tau}$ for large τ, for arbitrary $c_1 > 0$, $c_2 > 0$, $\alpha_1 > 0$, $\alpha_2 > 0$.

We have

$$\begin{aligned}
&\frac{1}{\sqrt{2\pi}} \int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty} \frac{e^{-x^2/2}}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^\gamma}\, dx \\
&\quad \leq \frac{1}{\sqrt{2\pi}} \int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty} e^{-(x^2+2x\gamma\sigma\sqrt{\tau})/2 - \gamma\mu\tau - \gamma d}\, dx \\
&\quad = \frac{1}{\sqrt{2\pi}} \int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty} e^{-(x+\gamma\sigma\sqrt{\tau})^2/2 + \frac{\gamma^2\sigma^2\tau}{2} - \gamma\mu\tau - d\gamma}\, dx \\
&\quad = e^{-\gamma d} e^{-\tau(\gamma\mu - \gamma^2\sigma^2/2)} \frac{1}{\sqrt{2\pi}} \int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}+\gamma\sigma\sqrt{\tau}}^{\infty} e^{-x^2/2}\, dx \\
&\quad = e^{-\gamma d} e^{-\tau(\gamma\mu - \gamma^2\sigma^2/2)} N\left(\upsilon\sqrt{\tau} + \frac{d}{\sigma\sqrt{\tau}}\right) \\
&\quad = e^{-\gamma d} e^{-\tau(\gamma\mu - \gamma^2\sigma^2/2)} (1 - O(e^{-\upsilon\tau})).
\end{aligned}$$

Also, since $1+e^{x\sigma\sqrt{\tau}+\mu\tau+d} \le 2e^{x\sigma\sqrt{\tau}+\mu\tau+d}$,

$$\begin{aligned}
&\frac{1}{\sqrt{2\pi}}\int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty}\frac{e^{-x^2/2}}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^{\gamma}}\,dx\\
&\quad\ge \frac{1}{2^{\gamma}}\frac{1}{\sqrt{2\pi}}\int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty} e^{-(x^2+2x\gamma\sigma\sqrt{\tau})/2-\gamma\mu\tau-\gamma d}\,dx\\
&\quad= \frac{1}{2^{\gamma}}\frac{1}{\sqrt{2\pi}}\int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty} e^{-(x+\gamma\sigma\sqrt{\tau})^2/2+\frac{\gamma^2\sigma^2\tau}{2}-\gamma\mu\tau-d\gamma}\,dx\\
&\quad= \frac{1}{2^{\gamma}}e^{-\gamma d}e^{-\tau(\gamma\mu-\gamma^2\sigma^2/2)}\frac{1}{\sqrt{2\pi}}\int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}+\gamma\sigma\sqrt{\tau}}^{\infty} e^{-x^2/2}\,dx\\
&\quad= \frac{1}{2^{\gamma}}e^{-\gamma d}e^{-\tau(\gamma\mu-\gamma^2\sigma^2/2)}N\left(\upsilon\sqrt{\tau}+\frac{d}{\sigma\sqrt{\tau}}\right)\\
&\quad= \frac{1}{2^{\gamma}}e^{-\gamma d}e^{-\tau(\gamma\mu-\gamma^2\sigma^2/2)}(1-O(e^{-\upsilon\tau})).
\end{aligned}$$

Thus, it is the case that

$$\frac{1}{\sqrt{2\pi}}\int_{-\frac{\mu\tau+d}{\sigma\sqrt{\tau}}}^{\infty}\frac{e^{-x^2/2}}{(1+e^{x\sigma\sqrt{\tau}+\mu\tau+d})^{\gamma}}\,dx = C_5 e^{-\tau(\gamma\mu-\gamma^2\sigma^2/2)},$$

where $C_5 \in \left[\frac{e^{-\gamma d}}{2^{\gamma}}-\epsilon, e^{-\gamma d}+\epsilon\right]$, for arbitrary $\epsilon > 0$, for large enough τ.

We therefore get

$$\begin{aligned}
-\frac{\log(P^{\tau})}{\tau} = -\frac{1}{\tau}\log\Bigg(&(1+e^d)^{\gamma}e^{-\rho\tau}\\
&\times\left(e^{-\tau\frac{\mu^2}{2\sigma^2}}e^{-\frac{\mu d}{\sigma^2}-\frac{d^2}{2\sigma^2\tau}}\frac{C_3}{q}+C_5e^{-\tau(\gamma\mu-\gamma^2\sigma^2/2)}\right)\Bigg).
\end{aligned}$$

Now, since $\frac{\mu^2}{2\sigma^2} \ge \gamma\mu-\gamma^2\sigma^2/2$, the second term within the log expression dominates the first, so we get

$$\begin{aligned}
-\frac{\log(P^{\tau})}{\tau} &= -\frac{1}{\tau}\left(\log\left((1+e^d)^{\gamma}e^{-\rho\tau}C_5e^{-\tau(\gamma\mu-\gamma^2\sigma^2/2)}\right)+o(\tau)\right)\\
&= \frac{(\rho+\gamma\mu-\gamma^2\sigma^2/2)\tau+o(\tau)}{\tau},
\end{aligned}$$

so indeed $\lim_{\tau\to\infty} -\frac{\log(P^{\tau})}{\tau} = \rho+\gamma\mu-\gamma^2\sigma^2/2 = \rho+\gamma(\mu+\sigma^2/2)-\gamma(\gamma+1)\sigma^2/2$. ■

Proof of Proposition 7.

Without loss of generality, we assume that $\underline{s} \le \overline{s}$, since the whole proof otherwise goes through by replacing $\underline{s}$ with $\overline{s}$.

We begin with (ii): It is easy to show the following inequality, which is valid for an arbitrary constant, $x \leq 0$:

$$\int_{\bar{t}}^{\infty} e^{-\rho s} \frac{e^{-\frac{(x-\mu s)^2}{2\sigma^2 s}}}{\sqrt{2\pi\sigma^2 s}} \, ds \geq \frac{e^{-(\rho+\mu\bar{t})}}{q} e^{\kappa x}, \tag{35}$$

where κ and q are defined in (4).

Now,

$$\begin{aligned}
P(C_0) &= E\left[\int_0^{\infty} e^{-\rho t} \left(\frac{f(D_0)}{f(D_t)}\right)^{\gamma} f(D_t)\, dt\right] \\
&\geq f(D_0)^{\gamma} E\left[\int_{\bar{t}}^{\infty} e^{-\rho t} f(D_t)^{1-\gamma}\, dt\right] \\
&\geq c_0^{\gamma} D_0^{\gamma} \int_{\bar{t}}^{\infty} e^{-\rho t} E\left[f(D_t)^{1-\gamma}\right] dt \\
&\geq c_0^{\gamma} D_0^{\gamma} \int_{\bar{t}}^{\infty} e^{-\rho t} E\left[f(D_t)^{1-\gamma} I_{s_t \leq \underline{s}}\right] dt \\
&\geq c_0^{\gamma} D_0^{\gamma} f(e^{\underline{s}})^{1-\gamma} \int_{\bar{t}}^{\infty} e^{-\rho t} E\left[I_{s_t \leq \underline{s}}\right] dt \\
&\geq c_0^{\gamma} D_0^{\gamma} f(e^{\underline{s}})^{1-\gamma} \int_{\bar{t}}^{\infty} e^{-\rho t} \Phi\left(\frac{\underline{s} - s_0 - \mu t}{\sigma\sqrt{t}}\right) dt \\
&= c_0^{\gamma} D_0^{\gamma} f(e^{\underline{s}})^{1-\gamma} \int_{\bar{t}}^{\infty} \int_{-\infty}^{\underline{s}-s_0} e^{-\rho t} \frac{e^{-\frac{(x-\mu t)^2}{2\sigma^2 t}}}{\sqrt{2\pi\sigma^2 t}}\, dx\, dt \\
&= c_0^{\gamma} D_0^{\gamma} f(e^{\underline{s}})^{1-\gamma} \int_{-\infty}^{\underline{s}-s_0} \int_{\bar{t}}^{\infty} e^{-\rho t} \frac{e^{-\frac{(x-\mu t)^2}{2\sigma^2 t}}}{\sqrt{2\pi\sigma^2 t}}\, dt\, dx \\
&\geq c_0^{\gamma} D_0^{\gamma} f(e^{\underline{s}})^{1-\gamma} \frac{e^{-(\rho+\mu\bar{t})}}{q} \int_{-\infty}^{\underline{s}-s_0} e^{\kappa x}\, dx \\
&= c_0^{\gamma} D_0^{\gamma} f(e^{\underline{s}})^{1-\gamma} \frac{e^{-(\rho+\mu\bar{t})}}{q} \frac{e^{\kappa\underline{s}}}{\kappa} \times e^{-\kappa s_0} \\
&= c_0^{\gamma} D_0^{\gamma} f(e^{\underline{s}})^{1-\gamma} \frac{e^{-(\rho+\mu\bar{t})}}{q} \frac{e^{\kappa\underline{s}}}{\kappa} \times D_0^{-\kappa} \\
&= c_0^{\gamma} f(e^{\underline{s}})^{1-\gamma} \frac{e^{-(\rho+\mu\bar{t})}}{q} \frac{e^{\kappa\underline{s}}}{\kappa} \times D_0^{\alpha} \\
&\geq c_0^{\gamma} f(e^{\underline{s}})^{1-\gamma} \frac{e^{-(\rho+\mu\bar{t})}}{q} \frac{e^{\kappa\underline{s}}}{\kappa} c_1^{-\alpha} \times f(D_0)^{\alpha} \\
&= c_3 C_0^{\alpha}.
\end{aligned}$$

For (i), we note that when $f(\epsilon) < c_2\epsilon$, we can choose an arbitrary $m > \max\{0, -\underline{s}\}$, to bound

$$\begin{aligned}
P(C_0) &= E\left[\int_0^\infty e^{-\rho t}\left(\frac{f(D_0)}{f(D_t)}\right)^\gamma f(D_t)\,dt\right] \\
&\geq f(D_0)^\gamma E\left[\int_{\bar{t}}^\infty e^{-\rho t} f(D_t)^{1-\gamma}\,dt\right] \\
&\geq c_0^\gamma D_0^\gamma \int_{\bar{t}}^\infty e^{-\rho t} E\left[f(D_t)^{1-\gamma}\right] dt \\
&\geq c_0^\gamma D_0^\gamma \int_{\bar{t}}^\infty e^{-\rho t} E\left[f(D_t)^{1-\gamma} I_{s_t \leq -m}\right] dt \\
&\geq c_0^\gamma D_0^\gamma f(e^{-m})^{1-\gamma} \int_{\bar{t}}^\infty e^{-\rho t} E\left[I_{s_t \leq -m}\right] dt \\
&\geq c_0^\gamma D_0^\gamma f(e^{-m})^{1-\gamma} \int_{\bar{t}}^\infty e^{-\rho t} \Phi\left(\frac{-m - s_0 - \mu t}{\sigma\sqrt{t}}\right) dt \\
&= c_0^\gamma D_0^\gamma f(e^{-m})^{1-\gamma} \int_{\bar{t}}^\infty \int_{-\infty}^{-m-s_0} e^{-\rho t} \frac{e^{-\frac{(x-\mu t)^2}{2\sigma^2 t}}}{\sqrt{2\pi\sigma^2 t}}\,dx\,dt \\
&= c_0^\gamma D_0^\gamma f(e^{-m})^{1-\gamma} \int_{-\infty}^{-m-s_0} \int_{\bar{t}}^\infty e^{-\rho t} \frac{e^{-\frac{(x-\mu t)^2}{2\sigma^2 t}}}{\sqrt{2\pi\sigma^2 t}}\,dt\,dx \\
&\geq c_0^\gamma D_0^\gamma f(e^{-m})^{1-\gamma} \frac{e^{-(\rho+\mu\bar{t})}}{q} \int_{-\infty}^{-m-s_0} e^{\kappa x}\,dx \\
&= c_0^\gamma D_0^\gamma f(e^{-m})^{1-\gamma} \frac{e^{-(\rho+\mu\bar{t})}}{q} \frac{e^{-\kappa m}}{\kappa} \times e^{-\kappa s_0} \\
&= c_0^\gamma D_0^\gamma f(e^{-m})^{1-\gamma} \frac{e^{-(\rho+\mu\bar{t})}}{q} \frac{e^{-\kappa m}}{\kappa} \times D_0^{-\kappa} \\
&= c_0^\gamma f(e^{-m})^{1-\gamma} \frac{e^{-(\rho+\mu\bar{t})}}{q} \frac{e^{-\kappa m}}{\kappa} \\
&\geq c_0^\gamma c_2^{1-\gamma} \frac{e^{-(\rho+\mu\bar{t})}}{q} \frac{1}{\kappa} D_0^\alpha \times e^{m(\gamma-\kappa-1)} \\
&= c_4(D_0) e^{m(\alpha-1)}.
\end{aligned}$$

Now, since $\alpha > 1$ and m is arbitrary, $P(C_0)$ must therefore be infinite, and the equilibrium does not exist. Equivalently, we could have used the identity $\frac{1}{1-\gamma}\frac{P(C_0)}{C_0} = U$ to show that expected utility is negative infinity for this case. ■

Mathematica code

Price-dividend Ratios

We have verified numerically that the formulae for the prices given in Proposition 2 are indeed correct, both above and below the breakpoint. The following Mathematica code calculates the price-dividend ratios for different D, for a long, but finite, horizon economy ($T = 1000$), using direct numerical integration of (10), and produces results identical to those shown in Figure 1.

In[1]:= $\gamma = 5; \sigma = 4/100; \mu = 0.75/100; \rho = 1/100; \xi = \mu - \frac{\sigma^2}{2}; T = 1000; B = 1;$ PD={};

In[2]:= v=Range[1/4,8,1/4];

In[3]:= For[i=1,i<32,
e=Extract[v,i];
v=NIntegrate[$((B+e)/(B+e*$Exp[y]$))^{\gamma-1} *$ Exp[$-\rho*\tau-(y-\xi*\tau)^2/(2*\sigma^2*\tau)$]
/Sqrt[$2*\pi*\sigma^2*\tau$], $\{y,-\infty,\infty\},\{\tau,0,T\}$];
PD=Append[PD,{e,v}],
i=i+1];

In[4]:= ListPlot[PD,PlotJoined->True,PlotRange->All];

Long-term Risk-free Rate

We have verified numerically that the formulae for the long rate given in Proposition 6 are indeed correct, by directly evaluating Equation (18). The following Mathematica code calculates the yield for different maturities.

For example, with parameters $\rho = 1\%$, $\mu = 3.5\%$, $\sigma = 20\%$, $\gamma = 2.5$, the long rate is close to $r^l = \rho + \frac{\mu^2}{2\sigma^2} = 2.53\%$, in line with Equation (21). The list L provides pairs of time to maturity and yields, $\{t, r_t\}$. For example, the last element in L shows that for a time to maturity of 10,000 years the yield is 2.56% in this example.

By varying $B0$, $D0$, and γ in the code, it is easily verified that the long rate does not depend on these parameters. It can also be checked that for $\mu > \gamma\sigma^2$, Equation (22) provides the correct long rate.

In[1]:= $B0 = 2; D0 = 1; \sigma = 0.2; \mu = 0.035; \gamma = 2.5; \rho = 0.01;$
Off[Integrate::gener];

In[2]:= $L = \{\}; T = \{1, 10, 100, 1000, 10000, -1\};$

In[3]:= For[t = First[T], $t > 0$,
P= N[Integrate[$(B0+D0)^\gamma *$ Exp[$-\rho t$]$*$

```
1/Sqrt[2π*σ^2 t]*Exp[−(y−μt)^2/(2*σ^2 t)]/(B0+D0*Exp[y])^γ,
{y, −∞, ∞}]];
r = −Log[P]/t;
L = Append[L, {t,r}]; T = Delete[T, 1]; t= First
[T];]

In[4]:= L (* L is a list with elements {t,r_t}, from
    numerical calculations*)

Out[4]= {{1,0.0362381},{10,0.0350963},{100,0.0307781},
    {1000,0.026798},{10000,0.0255731}}

In[5]:= r_l = If[μ < γσ^2, ρ+μ^2/(2σ^2), ρ+γμ−γ^2σ^2/2](*Theoretical
    value of long rate *)

Out[5]= 0.0253125
```


References

Abel, A. B. 1990. Asset Prices under Habit Formation and Catching Up with the Joneses. *American Economic Review* 80:38–42.

Abramowitz, M., and I. A. Stegun. 1964. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York: Dover.

Bansal, R., and A. Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59:1481–509.

Barro, R. 2005. Rare Events and the Equity Premium. Working Paper, Harvard University, Department of Economics.

Berk, J., and J. Walden. 2009. Limited Capital Market Participation and Human Capital Risk. Working Paper, Stanford Graduate School of Business.

Bhamra, H., L. Kuehn, and I. Strebulaev. 2010. The Levered Equity Risk Premium and Credit Spreads: A Unified Framework. *Review of Financial Studies* 23:645–703.

Blanchard, O. J. 1979. Speculative Bubbles, Crashes, and Rational Expectations. *Economics Letters* 3:387–89.

Blanchard, O. J., and M. W. Watson. 1982. Bubbles, Rational Expectations, and Financial Markets. In P. Wachtel (ed.), *Crises in the Economic and Financial Structure*. Lexington, MA: D. C. Heath and Company, 295–315

Campbell, J. Y. 1986. Bond and Stock Returns in a Simple Exchange Model. *Quarterly Journal of Economics* 101:785–804.

———. 2003. Consumption-based Asset Pricing. In G. M. Constantinides, M. Harris, and R. M. Stulz (eds.), *Handbook of the Economics of Finance*. Amsterdam: North Holland, 805-87

Campbell, J. Y., and J. H. Cochrane. 1999. By Force of Habit: A Consumption-based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* 107:205–51.

Cochrane, J. H. 1992. Explaining the Variance of Price-dividend Ratios. *Review of Financial Studies* 5:243–80.

———. 2001. *Asset Pricing*. Princeton, NJ: Princeton University Press.

Cochrane, J. H., F. A. Longstaff, and P. Santa-Clara. 2008. Two Trees. *Review of Financial Studies* 21:347–85.

Constantinides, G. M. 1990. Habit Formation: A Resolution of the Equity Premium Puzzle. *Journal of Political Economy* 98:519–43.

Diba, B. T., and H. I. Grossman. 1988. Explosive Rational Bubbles in Stock Prices? *American Economic Review* 78:520–30.

Dixit, A. K., and R. S. Pindyck. 1994. *Investment under Uncertainty*. Princeton, NJ: Princeton University Press.

Duffie, D., and L. G. Epstein. 1992. Stochastic Differential Utility. *Econometrica* 60:353–94.

Dybvig, P. H., J. E. Ingersoll Jr., and S. A. Ross. 1996. Long Forward and Zero-coupon Rates Can Never Fall. *Journal of Business* 69:1–25.

Engsted, T. 2006. Explosive Bubbles in the Cointegrated VAR Model. *Finance Research Letters* 3:154–62.

Ferson, W. E., and G. M. Constantinides. 1991. Habit Persistence and Durability in Aggregate Consumption: Empirical Tests. *Journal of Financial Economics* 29:199–240.

Froot, K. A., and M. Obstfeld. 1991. Intrinsic Bubbles: The Case of Stock Prices. *American Economic Review* 81:1189–214.

Gabaix, X., and D. Laibson. 2001. The 6D Bias and the Equity Premium Puzzle. In B. Bernanke and K. Rogoff (eds.), *NBER Macroeconomics Annual*. Cambridge, MA: MIT Press.

Geweke, J. 2000. Computational Experiments and Reality. Working Paper, University of Iowa.

———. 2001. A Note on Some Limitations of CRRA Utility. *Economic Letters* 71:341–45.

Gilles, C., and S. F. LeRoy. 1997. Bubbles as Payoffs at Infinity. *Economic Theory* 9:261–81.

Gradshteyn, L. S., and I. M. Ryzhik. 2000. *Tables of Integrals, Series, and Products*. New York: Academic Press.

Hansen, L. P., and R. Jagannathan. 1991. Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy* 99:225–62.

Heaton, J., and D. J. Lucas. 1996. Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing. *Journal of Political Economy* 104:443–87.

Ingersoll, J. E. Jr. 1987. *Theory of Financial Decision Making*. Totowa, NJ: Rowman and Littlefield.

Kim, T. S., and E. Omberg. 2002. Dynamic Nonmyopic Portfolio Optimization. *Review of Financial Studies* 9:141–61.

Kogan, L., S. Ross, J. Wang, and M. Westerfield. 2006. The Price Impact and Survival of Irrational Traders. *Journal of Finance* 61:195–229.

Labadie, P. 1989. Stochastic Inflation and the Equity Premium. *Journal of Monetary Economics* 24:277–98.

LeRoy, S. F. 2006. Excess Volatility. In S. N. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics*, 2nd ed. London: Palgrave Macmillan.

LeRoy, S., and R. Porter. 1981. The Present Value Relation: Tests Based on Variance Bounds. *Econometrica* 49:555–77.

Lucas, R. 1978. Asset Prices in an Exchange Economy. *Econometrica* 46:1429–45.

Malloy, C. J., T. J. Moskowitz, and A. Vissing-Jørgensen. 2009. Long-run Stockholder Consumption Risk and Asset Returns. *Journal of Finance* 64:2427–79.

Martin, I. 2009. The Lucas Orchard. Working Paper, Stanford Graduate School of Business.

Mehra, R., and E. C. Prescott. 1985. The Equity Premium: A Puzzle. *Journal of Monetary Economics* 15: 145–61.

Nielsen, L. 1984. Unbounded Expected Utility and Continuity. *Mathematical Social Sciences* 8:201–16.

Nielsen, L. 1987. Corrigenda. *Mathematical Social Sciences* 14:193–94.

Parker, J. 2001. The Consumption Risk of the Stock Market. *Brookings Papers on Economic Activity* 2:279–348.

Rietz, T. A. 1988. The Equity Risk Premium: A Solution. *Journal of Monetary Economics* 22:117–31.

Roche, H. 2001. Asset Pricing in a Lucas Economy with Recursive Utility and Heterogeneous Agents. Working Paper, Departmento de Administracion, Instituto Tecnologico Autonomo de Mexico.

Rubinstein, M. 1976. The Strong Case for the Generalized Logarithmic Utility Model as a Premier Model of Financial Markets. *Journal of Finance* 31:551–71.

———. 1983. Displaced Diffusion Option Pricing. *Journal of Finance* 38:213–17.

Savov, A. 2011. Asset Pricing with Garbage. *Journal of Finance*, Forthcoming.

Shiller, R. J. 1981. Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends? *American Economic Review* 71:421–36.

Triplett, J. E. 1997. Measuring Consumption: The Post-1973 Slowdown and the Research Issues. *Federal Reserve Bank of St. Louis Review* 79:9–42.

Tsionas, E. 2005. Likelihood Evidence on the Asset Returns Puzzle. *Review of Economic Studies* 72:917–46.

Wachter, J. 2002. Portfolio and Consumption Decisions Under Mean-reverting Returns: An Exact Solution for Complete Markets. *Journal of Financial and Quantitative Analysis* 37:63–91.

Weil, P. 1989. The Equity Premium Puzzle and the Risk-free Rate Puzzle. *Journal of Monetary Economics* 24:401–21.

Weitzman, M. L. 1998. Why the Far-distant Future Should Be Discounted at Its Lowest Possible Rate. *Journal of Environmental Economics and Management* 36:201–8.

———. 2001. Gamma Discounting. *American Economic Review* 91:260–71.

———. 2007. Subjective Expectations and Asset-return Puzzles. *American Economic Review* 97:1102–30.

West, K. D. 1987. A Specification Test for Speculative Bubbles. *Quarterly Journal of Economics* 102:553–80.