THE JOURNAL OF FINANCE • VOL. LXV, NO. 4 • AUGUST 2010

Information Quality and Long-Run Risk: Asset Pricing Implications

HENGJIE AI*

ABSTRACT

I study the asset pricing implications of the quality of public information about persistent productivity shocks in a general equilibrium model with Kreps–Porteus preferences. Low information quality is associated with a high equity premium, a low volatility of consumption growth, and a low volatility of the risk-free interest rate. The relationship between information quality and the equity premium differs from that in endowment economies. My calibration improves substantially upon the Bansal–Yaron model in terms of the moments of the wealth–consumption ratio and the return on aggregate wealth.

IN THIS PAPER, I introduce learning into a production-based long-run risk model. Learning improves the asset pricing implications of the model along three dimensions: it increases the equilibrium equity premium, reduces the volatility of consumption growth, and reduces the volatility of the risk-free interest rate. I show that because of the production economy setting, the relationship between information quality and the equity premium differs from that found in pure exchange economies (see, e.g., Veronesi (2000)). Quantitatively, my model improves substantially upon the Bansal and Yaron (2004) model in terms of its predictions about the volatility of the return on aggregate wealth and the wealth–consumption ratio. Compared to Bansal and Yaron (2004), I use a lower relative risk aversion (RRA) parameter, two, and a higher intertemporal elasticity of substitution (IES) parameter, two. My model generates a lower risk premium on the market portfolio, 5.06% per year, but does roughly as well as the Bansal and Yaron (2004) model according to other measures commonly used to assess the performance of consumption-based asset pricing models, such as the key moments of the risk-free interest rate and the consumption growth rate.

More specifically, I consider a general equilibrium model with a linear production technology and Kreps–Porteus preferences (Kreps and Porteus (1978)).

*Hengjie Ai is at the Fuqua School of Business, Duke University. This paper is based on chapter three of my doctoral dissertation at the University of Minnesota. I thank the acting editor, Pietro Veronesi, and an anonymous referee for many constructive criticisms. I thank Ravi Bansal, Michele Boldrin, Michael Brandt, V.V. Chari, Qi Chen, James Choi, Robert Goldstein, Larry Jones, Narayana Kocherlakota, David Levine, Erzo Luttmer, Ellen McGrattan, Martin Schneider, Jan Werner, Ofer Zeitouni, and participants at numerous seminars and conferences for helpful comments. Financial support from the Graduate School Dissertation Fellowship of the University of Minnesota is gratefully acknowledged. All errors are mine.

Productivity contains both persistent and transitory components, which are assumed to be statistically independent. In equilibrium, the persistent component of productivity shocks translates into persistent expected consumption growth, generating long-run risk.

Without learning, long-run risk does not contribute to the risk premium on aggregate wealth in the above economy if the shocks to the persistent and the transitory components of productivity are independent. Consequently, the model implies a risk premium that is too low if it is calibrated to match the volatility of the consumption growth rate in the data. Moreover, without learning, the model implies that the volatility of the consumption growth rate will be higher than the volatility of the return on aggregate wealth. However, according to Lustig, Van Nieuwerburgh, and Verdelhan (2008), the volatility of the return on aggregate wealth is estimated to be 4.94% per quarter, which is much higher than the estimated volatility of consumption growth during the same period, 0.44% per quarter.

By introducing learning I substantially improve the asset pricing implications of the model along three dimensions. First, learning leads to a higher equity premium if the representative agent's relative risk aversion is greater than one. With a risk aversion of two and a volatility of consumption growth of 2.80%, my benchmark learning model generates a risk premium on aggregate wealth of 2.68% per year, which is slightly higher than its empirical estimate of 2.16% per year in Lustig, Van Nieuwerburgh, and Verdelhan (2008). Without learning, the same model would predict a 0.17% per year risk premium on aggregate wealth if the volatility of consumption is calibrated to match the data. The key to understanding this result is that learning creates a positive covariance between the realized return and the expected return on the production technology (which in equilibrium equals the return on aggregate wealth). Whenever a high realized return is observed, the agent optimally revises her posterior belief of the expected return upward. The positive covariance induces a negative hedging demand when the agent's risk aversion is greater than one, which in equilibrium must be met by a higher risk premium for holding the claim to the technology.

Second, learning reduces the volatility of consumption growth if the agent's intertemporal elasticity of substitution is greater than one. In Lustig, Van Nieuwerburgh, and Verdelhan (2008), the estimated standard deviation of the log return on aggregate wealth is 4.94% per quarter, and that of the log wealth–consumption ratio is 11.74% per quarter. The standard deviation of the log consumption growth rate is smaller by an order of magnitude, 0.44% per quarter. With an IES of two, my benchmark model with learning provides a coherent interpretation of the above empirical facts.

The above empirical evidence on the volatility of the return on aggregate wealth and the wealth–consumption ratio suggests a large positive covariance between the log return on aggregate wealth and the log wealth–consumption ratio. To see this, following Campbell and Shiller (1989), I consider the log-linear approximation of the log return on aggregate wealth:

$$r_{W,t+1} = \kappa + \rho(w_{t+1} - c_{t+1}) + c_{t+1} - w_t, \tag{1}$$

where $r_{W,t+1}$ is the log return on aggregate wealth, w_{t+1} is the log aggregate wealth, c_{t+1} is the log aggregate consumption, and κ and ρ are parameters of linearization as defined in Campbell and Shiller (1989). Equation (1) implies that the innovation of consumption growth can be decomposed into an innovation in the log return on aggregate wealth and an innovation in the log wealth–consumption ratio. That is, using $[E_{t+1} - E_t]$ to denote innovations in conditional expectations, I have

$$[E_{t+1} - E_t](\Delta c_{t+1}) = [E_{t+1} - E_t](r_{W,t+1}) - \rho[E_{t+1} - E_t](w_{t+1} - c_{t+1}). \qquad (2)$$

This suggests that the variance of log consumption growth can be represented as the sum of the variances of the return on wealth and the log wealth–consumption ratio less the covariance of the two:

$$\mathrm{var}_t[\Delta c_{t+1}] = \mathrm{var}_t[r_{W,t+1}] + \rho^2 \mathrm{var}_t[w_{t+1} - c_{t+1}] - 2\rho \mathrm{cov}_t[r_{W,t+1}, w_{t+1} - c_{t+1}]. \qquad (3)$$

From equation (3), we can see that a high volatility of the log return on aggregate wealth and a high volatility of the log wealth–consumption ratio can be reconciled with a low volatility of the consumption growth rate only if the variance terms on the right-hand side of equation (3) are offset by a large positive covariance between the log return on aggregate wealth and the log wealth–consumption ratio.

In my model, learning, together with an IES of two, generates an endogenous positive covariance, $\mathrm{cov}_t[r_{W,t+1}, w_{t+1} - c_{t+1}]$, and produces a quarterly volatility of consumption growth of 1.14%.[1] Without learning, the covariance term in equation (3) is zero, and consequently the model will produce a counterfactually high level of volatility of consumption growth.

Third, learning decreases the volatility of the risk-free interest rate. Fluctuations in the risk-free interest rate come from fluctuations in expected consumption growth. If the agent observes the persistent component of the return on technology, then news about the persistent component of technology growth translates fully into innovations in expected consumption growth. As information quality decreases, the information content in news becomes smaller, and the agent's belief about future consumption growth becomes less sensitive to news, thereby making the risk-free rate less volatile.

The effect of learning on the equity premium in my production economy differs from that in pure exchange economies. In my paper, the assumption of a linear production technology implies that the return on aggregate wealth is determined exogenously by the technology. This delivers a simple separation result: learning increases the equity premium if RRA is greater than one, and reduces the volatility of consumption growth if IES exceeds one. In contrast, in pure exchange economies, the return on aggregate wealth is endogenously determined in equilibrium. In general, the effect of learning on the equity

[1]This number, although somewhat higher than the estimate provided by Lustig, Van Nieuwerburgh, and Verdelhan (2008), is consistent with most of the empirical estimates when a longer sample is considered; see, for example, Mehra and Prescott (1985) and Bansal and Yaron (2004).

premium depends on both RRA and IES. In the case of constant relative risk aversion (CRRA) preferences, Veronesi (2000) establishes that in pure exchange economies, learning decreases the equity premium if RRA is higher than one.

To better understand the difference between production and pure exchange economies, I consider a pure exchange economy with a Kreps–Porteus utility function. My study confirms Veronesi's (2000) finding in the more general setting. Furthermore, the Kreps–Porteus utility function allows me to separate the role of IES from that of RRA. I show that Veronesi's (2000) findings are driven by the agent's attitude toward intertemporal substitution, but not risk aversion.

Many other papers study the role of learning in understanding asset returns in a general or partial equilibrium context. Detemple (1986), Dothan and Feldman (1986), and Gennotle (1986) are the first to study learning and asset pricing in a fully dynamic framework. More recently, Brennan and Xia (2001) emphasize the role of learning in understanding the volatility of the stock market and the equity premium. Brevik and D'Addona (2007) study the relationship between information quality and the equity premium in a pure exchange economy with recursive preferences. Gollier and Schlee (2006) provide general conditions under which information increases or decreases the equity premium in pure exchange economies with general expected utility functions. Hansen, Sargent, and Tallarini (1999) and Hansen and Sargent (2009) study filtering and asset pricing problems in economies with robustness concerns. Croce, Lettau, and Ludvigson (2007) examine the implications of investors' information for the cross-sectional properties of stock return and cash flow duration in long-run risk models. However, the above papers do not investigate the asset pricing implications of information quality in a long-run risk production economy, nor do they confront the model with empirical evidence on the statistical properties of the wealth–consumption ratio, which is the main focus of this paper.

The paper is organized as follows. Section I describes the economy. Section II characterizes the optimal consumption allocation and studies the qualitative asset pricing implications of information quality. Section III calibrates the model and demonstrates the quantitative importance of learning in accounting for the key features of the volatility of the return on wealth and the wealth–consumption ratio in the data. Section IV considers a pure exchange economy with learning and compares my result on the relationship between information quality and the equity premium with that of Veronesi (2000). Section V concludes.

I. The Economy

A. Preferences

Consider an infinite horizon economy populated by a continuum of agents with identical Kreps–Porteus preferences with constant RRA parameter γ and constant IES parameter ψ. Time is continuous, and the representative agent's preference is described by the stochastic differential utility (SDU) developed by Duffie and Epstein (1992a,b). The SDU is the continuous-time version of the recursive preference considered in Kreps and Porteus (1978), Epstein and Zin (1989), and Weil (1989).[2] Given a consumption process $\{C_s : s \geq 0\}$, for every $t \geq 0$, the date-t utility of the agent, denoted V_t, is defined recursively by

$$V_t = E_t\left[\int_t^{\infty} f(C_s, V_s)\, ds\right]. \tag{4}$$

In the above equation, $f(C, V)$ is the aggregator of the recursive preference.[3] I focus on the class of SDU that has constant RRA, γ, and constant IES, ψ, so that

$$f(C, V) = \frac{\beta}{1 - 1/\psi} \frac{C^{1-1/\psi} - ((1-\gamma)V)^{\frac{1-1/\psi}{1-\gamma}}}{((1-\gamma)V)^{\frac{1-1/\psi}{1-\gamma}-1}}. \tag{5}$$

I assume $\gamma \neq 1$ for simplicity; extension to the limiting case $\gamma = 1$ is straightforward. However, I do allow $\psi = 1$ with the understanding that, in this case,

$$f(C, V) = \beta(1-\gamma)V\left[\ln C - \frac{1}{1-\gamma}\ln[(1-\gamma)V]\right].$$

B. *The Technology*

The economy is endowed with a linear technology that allows the agent to produce consumption goods from capital. Let K_t denote the total capital stock at time t. Further, let C_t denote the rate of consumption flow, and I_t the rate of the investment flow. Finally, let $A_{t,\Delta}$ denote the total factor productivity of the technology during the infinitesimal time interval $[t, t + \Delta]$. The resource constraint during this time interval is

$$C_t\Delta + I_t\Delta = A_{t,\Delta}K_t. \tag{6}$$

Given I_t, the law of motion of capital during the time interval $[t, t + \Delta]$ is

$$K_{t+\Delta} = (1 - \delta_{t,\Delta})K_t + I_t\Delta, \tag{7}$$

[2]Representation of SDU in the infinite-horizon case is discussed in Duffie and Epstein (1992b). The existence and uniqueness of SDU of the Kreps and Porteus (1978) type are discussed in Duffie and Lions (1992) and Schroder and Skiadas (1999).

[3]In general, a recursive preference is characterized by a pair of aggregators $(f, \mathcal{A})$. Duffie and Epstein (1992b) show that one can always normalize so that $\mathcal{A} = 0$. The aggregator f used here is the normalized aggregator.

where $\delta_{t,\Delta}$ is the depreciation rate of capital during the same time interval. Combining equations (6) and (7), the law of motion of capital can be written as

$$K_{t+\Delta} = (A_{t,\Delta} + 1 - \delta_{t,\Delta})K_t - C_t\Delta. \tag{8}$$

The above equation makes it clear that the productivity of the technology depends on $A_{t,\Delta} - \delta_{t,\Delta}$. I assume that $A_{t,\Delta} - \delta_{t,\Delta}$ consists of a persistent component, denoted by θ_t, and an i.i.d. component, modeled by increments of a Brownian motion $B_{K,t}$. That is,

$$A_{t,\Delta} - \delta_{t,\Delta} = \theta_t\Delta + \sigma_K(B_{K,t+\Delta} - B_{K,t}), \tag{9}$$

where σ_K is the volatility parameter of the return on the technology. Combining equations (8) and (9), and taking the limit as $\Delta \to 0$, the law of motion of capital can be written as

$$\forall t \geq 0, \quad dK_t = K_t[\theta_t dt + \sigma_K dB_{K,t}] - C_t dt. \tag{10}$$

The persistent productivity shock, θ_t, is modeled by an Ornstein–Uhlenbeck process. The law of motion of θ_t is described by the following stochastic differential equation (SDE):

$$\forall t \geq 0, \quad d\theta_t = a(\overline{\theta} - \theta_t)dt + \sigma_\theta dB_{\theta,t}, \tag{11}$$

where $a > 0$ determines the mean reversion rate of the persistent shock.

C. Information

I consider two alternative information structures in this paper. In an economy with complete information, $\forall t \geq 0,\ \theta_t$ is observable to the agent. In an economy with incomplete information, $\{\theta_t : t \geq 0\}$ is not observable. In the latter case, the agent observes two sources of noisy information about $\{\theta_t : t \geq 0\}$. First, the K_t process itself contains information about θ_t. In fact, equation (10) implies that

$$\frac{1}{K_t}(dK_t + C_t dt) = \theta_t dt + \sigma_K dB_{K,t}.$$

That is, knowing K_t and C_t is equivalent to observing the true θ_t plus a white noise $\sigma_K dB_{K,t}$. Second, the agent observes a noisy signal of θ_t, denoted $\{e_t : t \geq 0\}$, where

$$de_t = \theta_t dt + \sigma_e dB_{e,t}, \quad e_0 = 0. \tag{12}$$

Intuitively, e_t is determined by the true value of θ_t plus a white noise $\sigma_e dB_{e,t}$.

I assume that B_e is independent of $[B_K, B_\theta]$. The parameter $\sigma_e \geq 0$ is a measure of information quality. If $\sigma_e = 0$, then the $\{e_t : t \geq 0\}$ process carries perfect information about $\{\theta_t : t \geq 0\}$ and hence observing $\{e_t : t \geq 0\}$ is equivalent to observing $\{\theta_t : t \geq 0\}$ itself. If $\sigma_e > 0$, then $\{e_t : t \geq 0\}$ contains noisy information about $\{\theta_t : t \geq 0\}$. In the case $\sigma_e \to \infty$, $\{e_t : t \geq 0\}$ does not contain any information about $\{\theta_t : t \geq 0\}$ and the agent can update her belief about $\{\theta_t : t \geq 0\}$ only from the observed $\{K_t : t \geq 0\}$ process. I allow B_K and B_θ to be correlated and denote this correlation as

$$\rho = \text{corr}(B_K, B_\theta). \tag{13}$$

The agent's consumption policy can depend only on the information available and must satisfy the resource constraint. Formally, a consumption plan $\{C_t : t \geq 0\}$ is feasible with initial condition (θ_0, K_0) if:

1. For every $t \geq 0$, C_t depends only on the information available at time t.
2. Given $\{C_t : t \geq 0\}$, the solution to equations (10) and (11) satisfies $K_t \geq 0$ for all $t \geq 0$.

D. The Asset Market

I assume that n equities are traded on the market. Equities are indexed by $i \in \{1, 2, \dots, n\}$. I use boldface to indicate that the relevant variable is vector- or matrix-valued and superscripts to denote the index of an individual equity. Let $\boldsymbol{D}_t$ denote the vector of the rates of dividend payments of the n equities, and $\boldsymbol{P}_t$ denote the prices of the equities at time t, where

$$\boldsymbol{D}_t = \left[D_t^1,\ D_t^2, \dots,\ D_t^n\right]^T$$
$$\boldsymbol{P}_t = \left[P_t^1,\ P_t^2, \dots,\ P_t^n\right]^T.$$

For each i, $\{P_t^i : t \geq 0\}$ is assumed to be a diffusion process of the form

$$dP_t^i = P_t^i \left[\mu_{P,t}^i dt + \boldsymbol{\sigma}_{P,t}^i d\boldsymbol{B}_t\right]. \tag{14}$$

In the above equation, $\{\mathbf{B}_t : t \geq 0\}$ is a k-dimensional standard Brownian motion, $\mu_{P,t}^i$ is a scalar, and $\boldsymbol{\sigma}_{P,t}^i$ is a $1 \times k$ vector. It is convenient to define the cumulative return process, $\{R_t^i : t \geq 0\}$, for each equity i as in Duffie (2001):

$$R_0^i = 0, \quad \text{and} \quad dR_t^i = \left(\mu_{P,t}^i + \frac{D_t^i}{P_t^i}\right) dt + \boldsymbol{\sigma}_{P,t}^i d\boldsymbol{B}_t. \tag{15}$$

I use the vector notation

$$\boldsymbol{R}_t = \left[R_t^1, R_t^2, \dots, R_t^n\right]^T$$

and

$$d\boldsymbol{R}_t = \boldsymbol{\mu}_{R,t} dt + \boldsymbol{\sigma}_{R,t} d\boldsymbol{B}_t, \quad t \geq 0, \tag{16}$$

where $\boldsymbol{\mu}_{R,t}$ and $\boldsymbol{\sigma}_{R,t}$ are, respectively, the drift and diffusion coefficient of the cumulative return process of the form

$$\boldsymbol{\mu}_{R,t} = \left[\mu_{P,t}^1 + \frac{D_t^1}{P_t^1},\ \mu_{P,t}^2 + \frac{D_t^2}{P_t^2}, \dots,\ \mu_{P,t}^n + \frac{D_t^n}{P_t^n}\right]^T \tag{17}$$

and

$$\boldsymbol{\sigma}_{R,t} = \left[\boldsymbol{\sigma}^{1\ T}_{P,t},\ \boldsymbol{\sigma}^{2\ T}_{P,t}, \ldots, \boldsymbol{\sigma}^{n\ T}_{P,t}\right]^T. \tag{18}$$

A bond is traded and allows the agents in the economy to borrow from and lend to each other at a locally risk-free interest rate r_t. Therefore, $\mu^i_{R,t} - r_t$ is the instantaneous risk premium on equity i at time t.

The equity that pays aggregate consumption as its dividend is of particular importance. The value of this equity is equal to aggregate wealth, denoted W_t. I use $\{R_{W,t} : t \geq 0\}$ to denote the cumulative return process of aggregate wealth. I use $\mu_{W,t}$ and $\sigma_{W,t}$ to denote the drift and volatility of the cumulative return process of aggregate wealth.

In the competitive equilibrium of this economy, only the equity that pays aggregate consumption is in positive supply. All other assets are in net zero supply. The notion of competitive equilibrium is standard in this environment. In Section II, I first solve the planner's problem and appeal to the first and second fundamental welfare theorems to derive the implied asset prices.

II. Asset Pricing Implications

A. The Planner's Problem

The planner's problem is to maximize the utility of the representative agent, given in equation (4), subject to the feasibility constraints (10) and (11). If the representative agent observes θ_t, for every t, the optimal consumption plan, C_t, can be chosen as a function of the history $\{\theta_s, K_s : 0 \leq s \leq t\}$. Dynamic programming can be applied, and optimal consumption is a function of the state variables (θ_t, K_t).

In the incomplete information case, the planner's problem is to maximize the same objective function in (4) subject to the same constraints (10) and (11). However, feasibility now requires that the optimal consumption plan, C_t, depend only on the observables $\{K_s, e_s : 0 \leq s \leq t\}$, for every $t \geq 0$. In this case, I assume that the agent's date-0 prior belief about θ_0 is a Gaussian distribution with mean m_0 and variance Q_0.

The planner's problem in the incomplete information economy can be solved by a two-step procedure.[4] The first step is a learning problem, that is, deducing the conditional distribution of θ_t given observations. If the prior distribution of θ_0 is Gaussian, then the conditional distribution of θ_t given the observations of $\{K_s, e_s : 0 \leq s \leq t\}$ is Gaussian for all t. Therefore, the conditional distribution of θ_t can be characterized by the first two moments, $m_t = E_t[\theta_t]$ and $Q_t = \mathrm{var}_t[\theta_t]$. The moments m_t and Q_t can be obtained as the solution to the following Kalman filter (Liptser and Shiryaev (2001)):

$$dm_t = a(\overline{\theta} - m_t)\,dt + \left(\frac{1}{\sigma_K} Q_t + \rho\sigma_\theta\right) d\widetilde{B}_{K,t} + \frac{1}{\sigma_e} Q_t d\widetilde{B}_{e,t} \tag{19}$$

[4] A separation property applies here. The standard reference is Liptser and Shiryaev (2001).

$$dQ_t = \left\{\sigma_\theta^2 - 2aQ_t - \left[\left(\rho\sigma_\theta + \frac{1}{\sigma_K}Q_t\right)^2 + \left(\frac{1}{\sigma_e}Q_t\right)^2\right]\right\} dt. \tag{20}$$

In equation (19), $\widetilde{B}_{K,t}$ and $\widetilde{B}_{e,t}$ are defined recursively by

$$d\widetilde{B}_{K,t} = \frac{1}{\sigma_K}\left[\frac{1}{K_t}(dK_t + C_t dt) - m_t dt\right] \tag{21}$$

and

$$d\widetilde{B}_{e,t} = \frac{1}{\sigma_e}[de_t - m_t dt]. \tag{22}$$

Furthermore, Theorem 8.1 in Liptser and Shiryaev (2001) implies that $\{\widetilde{B}_{K,t}, \widetilde{B}_{e,t} : t \geq 0\}$ are independent Brownian motions.

Intuitively, equation (19) implies that the changes in m_t have three components: a predictable component, revisions that come from observations of K_t, and revisions that come from observations of the signal e_t. The locally predictable component, $a(\overline{\theta} - m_t)\,dt$, captures the mean reversion of m_t that comes from the mean reversion of θ_t. The definition of $d\widetilde{B}_{K,t}$ in equation (21) implies that $d\widetilde{B}_{K,t}$ is the revision of m_t that results from the deviation of the realized return on the technology, $\frac{1}{K_t}(dK_t + C_t dt)$, and the expected return on the technology, $m_t dt$. Similarly, $d\widetilde{B}_{e,t}$ is the revision of m_t that comes from news contained in e_t—that is, the difference between the realized changes in $e_t,\ de_t$, and the expected changes in $e_t,\ m_t dt$.

Equation (20) describes the law of motion of Q_t. As $t \longrightarrow \infty$, the posterior variance of $\theta_t,\ Q_t$, converges to the steady-state level Q, where Q is given by

$$Q = \frac{(1-\rho^2)\sigma_\theta^2}{\left(a + \rho\frac{\sigma_\theta}{\sigma_K}\right) + \sqrt{\left(a + \rho\frac{\sigma_\theta}{\sigma_K}\right)^2 + (1-\rho)\sigma_\theta^2\left(\sigma_K^{-2} + \sigma_e^{-2}\right)}}. \tag{23}$$

I further assume that the conditional variance starts at its steady-state level; therefore, $Q_t = Q$ for all t. Whenever necessary, I use the notation $Q(\sigma_e)$ to emphasize the dependence of Q on the information quality parameter σ_e. When $\sigma_e = 0,\ e_t$ perfectly reveals the true value of θ_t and thus $Q = 0$. In general, the steady-state posterior variance, Q, increases as information about θ_t contained in e_t becomes noisier. Mathematically,

$$Q(0) = 0; \quad \frac{\partial}{\partial \sigma_e} Q(\sigma_e) > 0. \tag{24}$$

Using equation (21), I can write the law of motion of K_t in terms of the innovation process $d\widetilde{B}_{K,t}$:

$$dK_t = K_t[m_t dt + \sigma_K d\widetilde{B}_{K,t}] - C_t dt. \tag{25}$$

Equations (19) and (25) together imply that the conditional covariance between the return on the technology and the expected return on the technology, m_t, is

$$\mathrm{cov}_t\left(\frac{dK_t}{K_t}, dm_t\right) = \rho\sigma_K\sigma_\theta + Q. \tag{26}$$

If θ_t is observable, the expected return on the technology is θ_t. Equations (10) and (11) imply that the conditional covariance between the return on the technology and the expected return on the technology is instead given by

$$\mathrm{cov}_t\left(\frac{dK_t}{K_t}, d\theta_t\right) = \rho\sigma_K\sigma_\theta. \tag{27}$$

Comparing equations (26) and (27) makes it clear that learning creates an additional positive covariance between the innovations in the realized return on the technology and the innovations in the expected return, m_t. Learning implies that whenever the realized return on the technology is high, the agent will revise her posterior belief about the expected return upwards, creating a positive covariance between the two.

It is convenient to denote

$$\sigma_m = \sqrt{\left(\frac{1}{\sigma_K}Q + \rho\sigma_\theta\right)^2 + \left(\frac{1}{\sigma_e}Q\right)^2} \tag{28}$$

and define

$$\tilde{B}_{m,t} = \frac{1}{\sigma_m}\left[\left(\frac{1}{\sigma_K}Q + \rho\sigma_\theta\right)\widetilde{B}_{K,t} + \frac{1}{\sigma_e}Q\widetilde{B}_{e,t}\right]. \tag{29}$$

Therefore, $\tilde{B}_m$ is a standard Brownian motion, and the law of motion of $\{m_t : t \geq 0\}$ can be written as

$$dm_t = a(\overline{\theta} - m_t)dt + \sigma_m d\widetilde{B}_{m,t}. \tag{30}$$

The second step in solving the planner's problem in the incomplete information case is a dynamic programming problem. By taking the posterior distribution of θ as a state variable, the second-step problem can be made recursive and solved by standard dynamic programming techniques. Since $Q_t = Q$ for all $t \geq 0$, the conditional mean m_t is sufficient to keep track of the conditional distribution of θ_t and thus can serve as the state variable in the second-stage dynamic programming problem. Using equations (19) and (25), the second-stage optimization problem can now be written as

$$\begin{aligned} V(K, m; \sigma_e) &= \max_{c_t \geq 0} E_0\left[\int_0^\infty f(C_t, V_t)\, dt\right] \\ \text{subject to:}\quad dK_t &= K_t[m_t dt + \sigma_K dB_{K,t}] - C_t dt, \quad K_0 = K,\ K_t \geq 0, \\ dm_t &= a(\overline{\theta} - m_t)dt + \sigma_m d\widetilde{B}_{m,t}, \quad m_0 = m. \end{aligned} \tag{31}$$

The value function of the planner's problem is denoted as $V(K, m; \sigma_e)$. I use this notation to emphasize the dependence of the value function on the information quality parameter σ_e.

The optimization problem in (31) is identical to the planner's problem of a complete information economy with m_t as the observable state variable.[5] The key difference between the optimization problem in (31) and the planner's problem in the same economy with complete information is that the covariances between the expected return and the realized return on the technologies are different, as summarized in equations (26) and (27). Therefore, learning affects the equilibrium prices and quantities of the model by endogenously generating the conditional distribution of the state variables.

In the rest of this section, I focus on the planner's problem in (31). The solution to the planner's problem in the case of the complete information economy can be obtained as the special case of (31) with $\sigma_e = 0$. The following proposition characterizes the solution to the planner's problem.

PROPOSITION 1: *Consider the social planner's problem, (31).*

- *The value function is of the form*

$$V(K, m; \sigma_e) = H(m; \sigma_e)\frac{K^{1-\gamma}}{1-\gamma}, \tag{32}$$

where $H(m; \sigma_e)$ satisfies the ordinary differential equation (ODE) in equation (IA.3) in the Internet Appendix.[6] *Moreover,*

$$\frac{H'(m; \sigma_e)}{H(m; \sigma_e)} < 0\ (> 0) \quad \textit{if} \quad \gamma > 1\ (\gamma < 1). \tag{33}$$

- *The optimal consumption policy function is given by*

$$C(K, m; \sigma_e) = x(m; \sigma_e)^{-1} K, \tag{34}$$

where $x(m; \sigma_e)$ is the wealth–consumption ratio, which satisfies

$$x(m; \sigma_e) = \beta^{-\psi} H(m; \sigma_e)^{-\frac{1-\psi}{1-\gamma}}. \tag{35}$$

Furthermore, $x(m; \sigma_e)$ is strictly increasing (decreasing) in m if $\psi > 1$ ($\psi < 1$).

Proof: See the Internet Appendix.

The first part of the proposition implies that the value function is homogeneous of degree $1-\gamma$ in K. This is due to the fact that the utility function is homogeneous of degree $1-\gamma$ in consumption, and the constraint set is linearly homogeneous in K. It follows that the value function is multiplicatively separable in m and K. Monotonicity of the value function with respect to m implies that $H(m; \sigma_e)$ is increasing or decreasing in m depending on the sign of γ.

The ratio of total capital stock to the rate of consumption flow, $x(m; \sigma_e)$, defined in equation (35), is also the wealth–consumption ratio of the representative

[5] A formal statement of this result can be found in Proposition 6 of Veronesi (1999).

[6] The Internet Appendix is available online at http://www.afajof.org/supplements.asp.

agent. This is because aggregate wealth measured in terms of current-period consumption goods, W_t, is equal to K_t for all $t \geq 0$. Intuitively, $W_t = K_t$ is true for two reasons. First, since capital is the only factor of production in this economy, the value of the total capital stock is equal to aggregate wealth. Second, equation (10) implies that one unit of C_t can always be transformed freely into one unit of K_t, and consequently the relative price of capital measured in current-period consumption is always equal to one.

The second part of Proposition 1 implies that the wealth–consumption ratio, $x(m; \sigma_e)$, is an increasing (decreasing) function of the state if $\psi > 1 (\psi < 1)$. This can be explained by the interaction between the income effect and the substitution effect as follows. A higher value of m_t implies a higher expected return on the technology. On the one hand, this means that the agent is wealthier and, other things being equal, will consume a greater proportion of her total wealth. This is the income effect, which tends to decrease the wealth–consumption ratio when m_t is higher. On the other hand, a higher expected return on the technology also encourages the agent to save more in order to raise future consumption. This is the intertemporal substitution effect, which tends to increase the wealth–consumption ratio for higher values of m_t. If $\psi > 1$, IES is large and therefore the substitution effect dominates the income effect. Consequently, the wealth–consumption ratio is increasing in m_t. By the same rationale, if $\psi < 1$, the income effect dominates and the wealth–consumption ratio decreases with m_t.

The above proposition also implies that equilibrium consumption growth will contain a slow-moving mean-reverting component. Therefore, the economy considered in this paper is essentially a continuous-time version of the Bansal and Yaron (2004) economy with production. To see this point, note that equations (34) and (35) imply that

$$C_t = \beta^{\psi} H(m_t; \sigma_e)^{\frac{1-\psi}{1-\gamma}} K_t. \tag{36}$$

By Ito's lemma, consumption is a diffusion process and can be written as

$$dC_t = C_t \left[\mu_C(m_t) dt + \sigma_K d\widetilde{B}_{K,t} + \frac{1-\psi}{1-\gamma} \frac{H'(m_t; \sigma_e)}{H(m_t; \sigma_e)} \sigma_m d\widetilde{B}_{m,t} \right], \tag{37}$$

where $\mu_C(m)$ is an increasing function of m. Since $\{m_t : t \geq 0\}$ is a mean-reverting process, the expected consumption growth is also mean-reverting. As demonstrated by Bansal and Yaron (2004), with Kreps–Porteus preferences, fluctuations in m will require a significant risk premium in equilibrium. In Section II.B, I study how information quality affects the risk premium associated with fluctuations in m.

B. The Equity Premium

In this subsection, I establish that learning increases the risk premium on aggregate wealth when $\gamma > 1$. To solve for the risk premium on aggregate wealth, I first consider a general specification of the portfolio choice problem of the representative consumer. Suppose the drift and diffusion coefficients of the cumulative return processes of the n equities specified in equation (16) are functions of the state variable m, that is,

$$\boldsymbol{\mu}_{R,t} = \boldsymbol{\mu}_R(m_t), \quad \boldsymbol{\sigma}_{R,t} = \boldsymbol{\sigma}_R(m_t), \quad t \geq 0,$$

where $\boldsymbol{\mu}_R(m)$ is an $\mathcal{R}^n$-valued measurable function of m, and $\boldsymbol{\sigma}_R(m)$ is an $\mathcal{R}^{n\times k}$-valued one. In equilibrium, the drift and diffusion parameters of the cumulative return on aggregate wealth are also functions of m. I denote

$$\mu_{W,t} = \mu_W(m_t); \quad \boldsymbol{\sigma}_{W,t} = \boldsymbol{\sigma}_W(m_t),$$

where $\boldsymbol{\mu}_W(m)$ is a real-valued measurable function of m, and $\boldsymbol{\sigma}_W(m)$ is an $\mathcal{R}^k$-valued one. I also use $\boldsymbol{\sigma}_{R,m}(m_t)$ to denote the $n \times 1$ vector of covariances of the returns with the state variable m, which is also assumed to be a function of the state variable m:

$$\boldsymbol{\sigma}_{R,m}(m_t) = \mathrm{cov}_t[d\boldsymbol{R}_t, dm_t].$$

Finally, I assume that the risk-free interest rate r_t is a function of the state variable, m_t, denoted $r(m_t)$.

Denote ϕ_i as the fraction of the agent's total wealth invested in equity i, and denote

$$\boldsymbol{\phi} = [\phi_1, \phi_2, \ldots, \phi_n]^T.$$

The law of motion of the representative agent's wealth is written as

$$dW_t = W_t\{\boldsymbol{\phi}^T[\boldsymbol{\mu}_R(m_t) - r(m_t)] + r(m_t)\}dt + W_t\boldsymbol{\phi}^T\boldsymbol{\sigma}_R(m_t)d\boldsymbol{B}_t - C_t dt. \quad (38)$$

In the competitive equilibrium, the agent's objective is to maximize the utility function in equation (4) subject to the constraints (38) and (30). The value function of this maximization problem is a function of W_t and m_t, denoted $U(W, m)$. The following lemma characterizes the solution to the optimal portfolio choice problem. The proof of the lemma can be found in Duffie and Epstein (1992a).

LEMMA 1: *The optimal consumption policy and the optimal portfolio choice policy are functions of the state variables* (W, m). *Let* $C^*(W, m)$ *and* $\phi^*(m)$ *denote the consumption and portfolio policy functions. Then, for all* (W, m),

$$U_W(W, m) = f_C(C^*(W, m), U(W, m)) \quad (39)$$

and

$$[\boldsymbol{\mu}_R(m) - r(m)] + \frac{U_{W,m}(W, m)}{U_W(W, m)}\boldsymbol{\sigma}_{R,m}(m) + \frac{WU_{WW}(W, m)}{U_W(W, m)}\boldsymbol{\sigma}_R(m)\boldsymbol{\sigma}_R(m)^T\boldsymbol{\phi}^*(m) = 0. \quad (40)$$

Note that equation (38) implies that the diffusion parameter of the cumulative return on total wealth satisfies

$$\boldsymbol{\sigma}_W(m) = \boldsymbol{\phi}^*(m)^T\boldsymbol{\sigma}_R(m). \quad (41)$$

Using equations (40) and (41), I obtain Merton's intertemporal capital asset pricing formula (Merton (1973)):

$$\boldsymbol{\mu}_R(m) - r(m) = -\frac{WU_{WW}(W,m)}{U_W(W,m)}\operatorname{cov}\left[d\boldsymbol{R}, \frac{dW}{W}\middle| m\right] - \frac{U_{W,m}(W,m)}{U_W(W,m)}\operatorname{cov}[d\boldsymbol{R}, dm\,|\,m]. \tag{42}$$

Equation (42) reveals that the risk premium on any asset can be decomposed into two components: compensation for its covariance with the return on aggregate wealth, and compensation for its covariance with the innovations of the expected return on aggregate wealth. In fact, as I show in the Internet Appendix, the state price density of this economy, denoted $\{\pi_t : t \geq 0\}$, satisfies

$$\frac{d\pi_t}{\pi_t} - E_t\left[\frac{d\pi_t}{\pi_t}\right] = \frac{W_t U_{WW}(W_t, m_t)}{U_W(W_t, m_t)}\boldsymbol{\sigma}_{R,W}(m_t)d\boldsymbol{B}_t + \frac{U_{W,m}(W_t, m_t)}{U_W(W_t, m_t)}\sigma_m d\widetilde{B}_{m,t}. \tag{43}$$

Therefore, the above interpretation is indeed appropriate. I define the term

$$-\frac{WU_{WW}(W,m)}{U_W(W,m)}\operatorname{cov}\left[d\boldsymbol{R}, \frac{dW}{W}\middle| m\right] \tag{44}$$

as the myopic demand component of the risk premium. This term accounts for the total risk premium if the total demand of the equity coincides with the myopic demand, that is, if the intertemporal hedging demand (Merton (1971)) of the equity is zero. The second term in (42),

$$-\frac{U_{W,m}(W_t, m_t)}{U_W(W_t, m_t)}\operatorname{cov}_t[dR_t, dm_t], \tag{45}$$

is defined as the hedging demand component of the risk premium, since this term reflects the adjustment of the risk premium on the equity due to the agent's intertemporal hedging demand. This decomposition turns out to be very useful when I calibrate the model to the empirical evidence on the return properties of aggregate wealth, and when I compare my result on the relationship between the equity premium and information quality with that of Veronesi (2000) in pure exchange economies.

In this economy, $W_t = K_t$ for all t; therefore, the cumulative return process of aggregate wealth is exogenously determined by the technology:

$$dR_{W,t} = \frac{1}{K_t}[dK_t + C_t dt] = m_t dt + \sigma_K d\widetilde{B}_{K,t}. \tag{46}$$

The fact that $W_t = K_t$ also implies that the value function in the planner's problem coincides with the value function in Lemma 1. Using Proposition 1, equation (42) can be written as

$$\boldsymbol{\mu}_R(m) - r(m) = \gamma \operatorname{cov}\left[d\boldsymbol{R}, \frac{dW}{W}\middle| m\right] - \frac{H'(m)}{H(m)}\operatorname{cov}[d\boldsymbol{R}, dm\,|\,m]. \tag{47}$$

Equations (46) and (47) can be used to derive the formula for the risk premium on aggregate wealth, which is summarized in the following proposition.

PROPOSITION 2: *The risk premium on aggregate wealth is given by*

$$\mu_{W,t} - r_t = \gamma\sigma_K^2 - \frac{H'(m_t;\sigma_e)}{H(m_t;\sigma_e)}(\rho\sigma_K\sigma_\theta + Q). \tag{48}$$

Together with the inequality in (33), the above proposition implies that learning increases the risk premium on aggregate wealth when $\gamma > 1$. In this economy, learning does not affect the myopic demand component of the equity premium, and it increases the hedging demand component of the equity premium when $\gamma > 1$.

To understand the above proposition, first consider the special case $\rho = 0$, that is, the case in which the innovations in the return on the technology and the innovations in the state variable θ are independent. Here, the long-run risk in consumption growth does not contribute to the risk premium on aggregate wealth if θ is perfectly observable. Learning restores the long-run risk premium and increases the overall risk premium on aggregate wealth if $\gamma > 1$.

If θ is observable, that is, $\sigma_e = 0$, then equation (24) implies that the risk premium on aggregate wealth is

$$\mu_{W,t} - r_t = \gamma\sigma_K^2. \tag{49}$$

Using Proposition 1, it can be shown that

$$\mu_{W,t} - r_t = \gamma \operatorname{cov}_t\left[\frac{dC_t}{C_t}, \frac{dW_t}{W_t}\right].$$

That is, the Lucas (1978)–Breeden (1979) formula applies in this case. Despite the recursive preferences and the persistent productivity shocks, the risk premium on aggregate wealth is the same as that in an economy with CRRA preferences and i.i.d. productivity growth. This is not surprising, given the assumption $\rho = 0$. Intuitively, the long-run risk comes from $dB_{\theta,t}$, and fluctuations in the return on the technology come from $dB_{K,t}$. Since $dB_{K,t}$ and $dB_{\theta,t}$ are uncorrelated, the return on the technology, and hence aggregate wealth, are not exposed to the long-run risk in consumption.

If, instead, $\rho = 0$ but θ is not observable, the long-run risk is priced and the total risk premium is given by

$$\mu_{W,t} - r_t = \gamma\sigma_K^2 - \frac{H'(m_t;\sigma_e)}{H(m_t;\sigma_e)}Q.$$

Here, learning creates an additional term $-\frac{H'(m_t;\sigma_e)}{H(m_t;\sigma_e)}Q$ in the equity premium. From Proposition 1, the sign of $\frac{H'(m;\sigma_e)}{H(m;\sigma_e)}$ depends on γ; therefore, learning increases the equity premium if $\gamma > 1$, and decreases the equity premium if $\gamma < 1$. Intuitively, although $dB_{K,t}$ is uncorrelated with the fluctuation in the state variable θ_t, it is correlated with the posterior mean of θ_t, m_t, because of learning. Since m_t drives the persistent component of consumption growth, as shown in equation (37), the risk in $dB_{K,t}$ will be priced.

In the general case $\rho \neq 0$, learning enhances the hedging demand component of the risk premium on aggregate wealth by increasing the premium for long-run risk in holding aggregate wealth. To see this, by equation (48), the long-run average equity premium in the economy is given by

$$E^*[\mu_{R,W,t} - r_t] = \gamma\sigma_K^2 - (\rho\sigma_K\sigma_\theta + Q)E^*\left[\frac{H'(m;\sigma_e)}{H(m;\sigma_e)}\right], \tag{50}$$

where the notation $E^*[\cdot]$ denotes that the expectation in (50) is taken with respect to the steady-state distribution of m. Numerical results indicate that the term $E[\frac{H'(m;\sigma_e)}{H(m;\sigma_e)}]$ is hardly affected by σ_e. In fact, using the log-linear approximation method proposed in Campbell et al. (2004), one can show that

$$\frac{H'(m_t;\sigma_e)}{H(m_t;\sigma_e)} \approx \frac{1-\gamma}{a+\kappa_1}, \tag{51}$$

where $\kappa_1 > 0$ is a constant as given in the Internet Appendix. Therefore,

$$\frac{\partial}{\partial\sigma_e}E^*[\mu_{R,W,t} - r_t] \approx \frac{\gamma-1}{a+\kappa_1}\frac{\partial}{\partial\sigma_e}Q\,(\sigma_e). \tag{52}$$

This implies that if $\gamma > 1(\gamma < 1)$, the equity premium increases (decreases) as the noise contained in the $\{e_t\}_{t\geq 0}$ process increases.

The key implication of Proposition 2 is that the direction of the effect of information quality on the equity premium depends on the RRA parameter. This is because learning affects only the hedging demand component of the equity premium, the sign of which is determined by the risk-aversion parameter.

The hedging demand component of the risk premium on aggregate wealth is

$$-\frac{U_{W,m}(W_t,m_t)}{U_W(W_t,m_t)}\mathrm{cov}_t\left[\frac{dW_t}{W_t},dm_t\right] = -\frac{H'(m_t)}{H(m_t)}\mathrm{cov}_t\left[\frac{dW_t}{W_t},dm_t\right]. \tag{53}$$

Proposition 1 implies that the coefficient $-\frac{H'(m_t)}{H(m_t)}$ is positive (negative) if $\gamma > 1(\gamma < 1)$. The interpretation is that if $\gamma > 1$, the agent dislikes assets with an expected return that is positively correlated with the return on her total wealth. Consequently, the hedging demand for the asset is negative when the covariance term in (53) is positive. Similarly, if $\gamma < 1$ and the expected return on the asset and the return on the aggregate wealth are positively correlated, the hedging demand will be positive.

As shown in equations (26) and (27), if θ is not observable, learning creates an additional positive covariance between innovations in the return on aggregate wealth and innovations in the expected return on aggregate wealth. The positive covariance induces a negative (positive) hedging demand if $\gamma > 1(\gamma < 1)$. In equilibrium, the expected return on equity has to adjust to equate supply and demand; therefore, if $\gamma > 1$, the negative hedging demand created by learning translates into a higher equity premium in equilibrium. More generally, the equity premium increases as the information contained in $\{e_t : t \geq 0\}$ becomes noisy, as shown in equation (52).

The setup of the production economy allows me to study not only the effect of information quality on the equity premium, but also its implications for equilibrium allocations, and, in particular, for the volatility of consumption growth rates. In the next subsection, I analyze the effect of learning on the volatility of the consumption growth rate and on that of the risk-free interest rate.

C. Volatility of Consumption Growth and the Risk-Free Interest Rate

The high equity premium, the low volatility of consumption growth, and the low volatility of the risk-free rate observed in the data are among the most important challenges to consumption-based asset pricing models. In Section II.B, I provide conditions under which learning either increases or decreases the equity premium. This subsection is devoted to analysis of the effect of learning on the volatility of both consumption growth and the risk-free interest rate.

First, learning decreases (increases) the conditional volatility of consumption growth if $\psi > 1 (\psi < 1)$. To see this, using equation (37),

$$\mathrm{var}_t(d \ln C_t) = \left(\frac{1-\psi}{1-\gamma}\frac{H'(m_t;\sigma_e)}{H(m_t;\sigma_e)}\right)^2 \sigma_m^2 + \sigma_K^2 + 2(\rho\sigma_K\sigma_\theta + Q)\frac{1-\psi}{1-\gamma}\frac{H'(m_t;\sigma_e)}{H(m_t;\sigma_e)}. \tag{54}$$

Applying the definition of σ_m in (28) and equation (23), I obtain

$$\sigma_m^2 = \sigma_\theta^2 - 2aQ.$$

By the log-linear approximation in (51), equation (54) can be written as

$$\mathrm{var}_t(d \ln C_t) \approx \left(\frac{1-\psi}{a+\kappa_1}\right)^2 \left(\sigma_\theta^2 - 2aQ\right) + \sigma_K^2 + 2(\rho\sigma_K\sigma_\theta + Q)\frac{1-\psi}{a+\kappa_1}.$$

Therefore, together with (24), the above equation implies

$$\frac{\partial}{\partial\sigma_e}\mathrm{var}_t(d \ln C_t) \approx 2(1-\psi)\frac{a\psi+\kappa_1}{(a+\kappa_1)^2}\frac{\partial}{\partial\sigma_e}Q(\sigma_e) < 0\ (>0) \quad \text{if} \quad \psi > 1 (\psi < 1). \tag{55}$$

That is, the conditional volatility of consumption growth is decreasing (increasing) in σ_e if $\psi > 1 (\psi < 1)$.

To understand equation (55) intuitively, note that the identity

$$\ln C_t = \ln W_t - \ln \frac{W_t}{C_t}$$

implies the variance decomposition

$$\mathrm{var}_t[d\ln C_t] = \mathrm{var}_t[d\ln K_t] + \mathrm{var}_t[d\ln x(m_t;\sigma_e)] - 2\mathrm{cov}_t[d\ln K_t, d\ln x(m_t;\sigma_e)]. \tag{56}$$

Equation (56) is essentially the continuous-time version of equation (3). It implies that if the return on aggregate wealth ($d\ln K_t$) and the innovation in the wealth–consumption ratio ($d\ln x(m_t;\sigma_e)$) are uncorrelated, then the conditional variance of consumption growth is the sum of the conditional variance of the return on aggregate wealth and that of the wealth–consumption ratio. When $\psi > 1$, learning creates a positive correlation between the return on aggregate wealth and the wealth–consumption ratio, thereby reducing the volatility of consumption growth. An IES higher than one implies that the agent has a strong incentive to substitute future consumption for today's consumption. Consequently, the agent optimally chooses a high wealth–consumption ratio whenever the expected future return on the technology is high. This generates an endogenous positive covariance between the return on aggregate wealth and the innovation in the wealth–consumption ratio, since high expected returns are associated with high realized returns because of learning. The positive covariance offsets the two variance terms in equation (56) and reduces the overall conditional volatility of the consumption growth rate.

In addition, learning decreases the steady-state volatility of the risk-free interest rate. To understand this result, note that since the expected return on aggregate wealth is m_t, equation (48) implies that the risk-free interest rate can be written as

$$r_t = m_t - \gamma\sigma_K^2 + \frac{H'(m_t;\sigma_e)}{H(m_t;\sigma_e)}(\rho\sigma_K\sigma_\theta + Q) \approx m_t - \gamma\sigma_K^2 + \frac{1-\gamma}{a+\kappa_1}(\rho\sigma_K\sigma_\theta + Q),$$

where the approximate equality uses the log-linearization approximation in (51). Therefore, the unconditional variance of the risk-free interest rate is approximately equal to the unconditional variance of the posterior mean m_t:

$$\mathrm{var}^*(r_t) \approx \mathrm{var}^*(m_t) = \frac{1}{2a}\sigma_m^2 = \frac{1}{2a}\left[\sigma_\theta^2 - 2aQ\right],$$

where the notation $\mathrm{var}^*[\cdot]$ denotes the unconditional variance evaluated at the steady-state distribution of m_t. Consequently,

$$\frac{\partial}{\partial\sigma_e}\mathrm{var}^*(r_t) = -\frac{\partial}{\partial\sigma_e}Q(\sigma_e) < 0.$$

Intuitively, fluctuations in the interest rate in this economy come from fluctuations in expected consumption growth, which in turn come from fluctuations in the posterior belief m_t. If information is imprecise, then the revision of the posterior belief m_t is small whenever new information arrives. In the extreme case where there is absolutely no new information, m_t would be a constant: the agent does not change her belief at all. In the model, the $\{K_t\}_{t\geq 0}$ process always carries some information about $\{\theta_t\}_{t\geq 0}$; therefore, the steady-state volatility of the interest rate achieves its minimum when $\sigma_e = \infty$. When information quality improves, the posterior mean becomes more sensitive to the arrival of new information. In fact, when θ_t is known, m_t moves one for one with movements in θ_t and $\text{var}^*(r_t) \approx \frac{1}{2a}\sigma_\theta^2$. This is the case in which the steady-state volatility of m_t, and consequently the steady-state volatility of the risk-free interest rate, is maximized.

To summarize, my results imply that in an economy with linear production technology and SDU, the direction in which learning affects the equity premium is completely determined by the RRA parameter, while the sign of the effect of learning on the volatility of consumption growth depends only on the IES parameter. I also show that learning reduces the volatility of the risk-free interest rate regardless of the preference parameters. My results are similar in spirit to the findings of Tallarini (2000) in the sense that in production economy models with Kreps–Porteus preferences, the quantity implications are primarily determined by IES and the risk premiums mostly depend on RRA.

In Section III, I calibrate the model to the U.S. economy and demonstrate the quantitative importance of learning in understanding some of the recent empirical evidence on the statistical properties of the wealth–consumption ratio provided by Lustig, Van Nieuwerburgh, and Verdelhan (2008).

III. Calibration

The purpose of this section is to assess the model's ability to account for some of the key asset pricing statistics in the data. Although the long-run risk model developed by Bansal and Yaron (2004) successfully explains many salient features of the asset market data, the volatility of the return on aggregate wealth and the volatility of the wealth–consumption ratio implied by the model are too low relative to the empirical evidence presented in Lustig, Van Nieuwerburgh, and Verdelhan (2008). My benchmark model substantially improves the performance of the Bansal and Yaron (2004) model relative to the empirical findings of Lustig, Van Nieuwerburgh, and Verdelhan (2008). Other asset pricing statistics of the model, such as the unconditional moments of the consumption growth rate, the risk premium on the market portfolio, and the risk-free interest rate, are largely consistent with their counterparts in the data and in the Bansal and Yaron (2004) model.

A. Parameter Choices

The choices of the preference parameters are as follows: $\beta = 0.014,\ \gamma = 2$, and $\psi = 2$. I calibrate the discount rate $\beta = 0.014$. This is equivalent to an annual discount factor of 0.986 per year in discrete time models. I choose the risk-aversion parameter $\gamma = 2$. This choice of risk-aversion parameter is consistent with that used in most of the macroeconomics literature, for example, Rouwenhorst (1995). In the benchmark model, I use an IES parameter $\psi = 2$.

Robust analysis with respect to the IES parameter is reported in the next subsection.

I choose the technology parameters as follows: $\sigma_K = 0.099$, $a = 0.027$, $\bar{\theta} = 0.035$, and $\sigma_\theta = 0.005$. I choose σ_K to match the volatility of the return on aggregate wealth and $\bar{\theta}$ to match the level of the mean risk-free interest rate. The estimated volatility of the return on aggregate wealth by Lustig, Van Nieuwerburgh, and Verdelhan (2008) is 4.94% at the quarterly level. With $\sigma_K = 0.099$, equation (46) implies that the instantaneous volatility of the return on aggregate wealth is

$$\lim_{\Delta \to 0} Std_t \left[\frac{R_{t+\Delta}}{R_t} - 1 \right] = \sigma_K = 9.9\%.$$

The annualized volatility of the return on aggregate wealth is also affected by the fluctuations in θ and is slightly higher than 9.9%. This brings the quarterly volatility of the return on aggregate wealth in my model to 4.97%, which closely matches the estimate in Lustig, Van Nieuwerburgh, and Verdelhan (2008). Changes in $\bar{\theta}$ primarily affect the risk-free interest rate without changing the risk premium. I choose $\bar{\theta} = 0.035$ so that the average of the annualized risk-free interest rate in the model is 0.86%, which matches the point estimate of the same moment in Bansal and Yaron (2004).

Because of the statistical difficulty of estimating the long-run risk parameters, there is little independent empirical evidence available to discipline the choice of a and σ_θ. Here I choose $a = 0.027$ and $\sigma_\theta = 0.005$ and verify that the key statistics of consumption dynamics generated by the model are consistent with those in the data. Equation (3) shows that, without learning, the volatility of consumption growth is at least as high as the volatility of the return on aggregate wealth. Consequently, it is impossible for the model to match the second moment of the consumption growth rate in the data. I set $\sigma_e = \infty$ in the benchmark model to maximize the impact of learning. In Section III.B, I show that the key moments of the equilibrium consumption growth rate generated in the benchmark model with learning are largely consistent with the data. I also change σ_e and evaluate the effect of learning on the dynamics of the consumption growth rate and the wealth–consumption ratio.

B. Dynamics of Aggregate Consumption and Aggregate Wealth

In this subsection, I discuss the model's implications on the dynamics of the consumption growth rate and the return on aggregate wealth. I show that the moments of consumption growth rate generated by the model are largely consistent with those in the data, and the model improves substantially upon the Bansal and Yaron (2004) model in terms of the statistical properties of the wealth–consumption ratio and the return on aggregate wealth.

The model is solved numerically using the Markov chain approximation method developed in Kushner and Dupuis (2001). Figure 1 plots the risk premium on aggregate wealth as a function of the state variable m. The solid

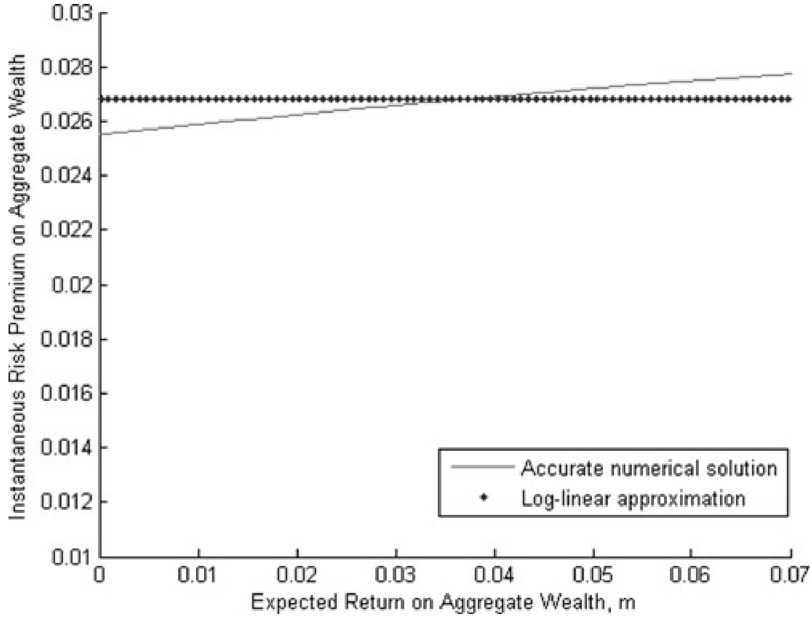

Figure 1. Expected return on aggregate wealth. The dotted line is the risk premium on aggregate wealth obtained by log-linear approximation. The solid line is the accurate numerical solution for the risk premium on aggregate wealth computed using the Markov chain approximation method.

line is the instantaneous risk premium on aggregate wealth as given in (48). The dotted line represents the risk premium on aggregate wealth obtained by using the log-linear approximation in equation (51). Log-linear approximation assumes that $\ln H(m)$ is linear in m, and therefore $\frac{H'(m)}{H(m)}$ is a constant. Figure 1 shows that the risk premium on aggregate wealth is slightly increasing in m. This is because $\frac{H'(m)}{H(m)}$ is slightly decreasing in m. The log-linearization method gives a fairly good approximation of the level of the risk premium, but does not capture the dependence of the risk premium on m. Figure 2 plots the expected return on aggregate wealth and the risk-free interest rate as a function of m. As m rises, the expected return on aggregate wealth increases and most of the increase is due to changes in the risk-free interest rate. Empirical evidence suggests that risk premia are countercyclical (see, e.g., Cochrane (2005)). A richer model that incorporates countercyclical volatility of the return on aggregate wealth could generate the pattern of time-varying risk premia consistent with the empirical evidence. However, I abstract from time-varying volatility in order to focus on the effect of learning on the unconditional moments of returns and consumption growth rates.

Table I reports the moments of the wealth–consumption ratio generated by the model. Estimates of these moments from the data, and the corresponding moments produced by the Bansal and Yaron (2004) model, are taken from

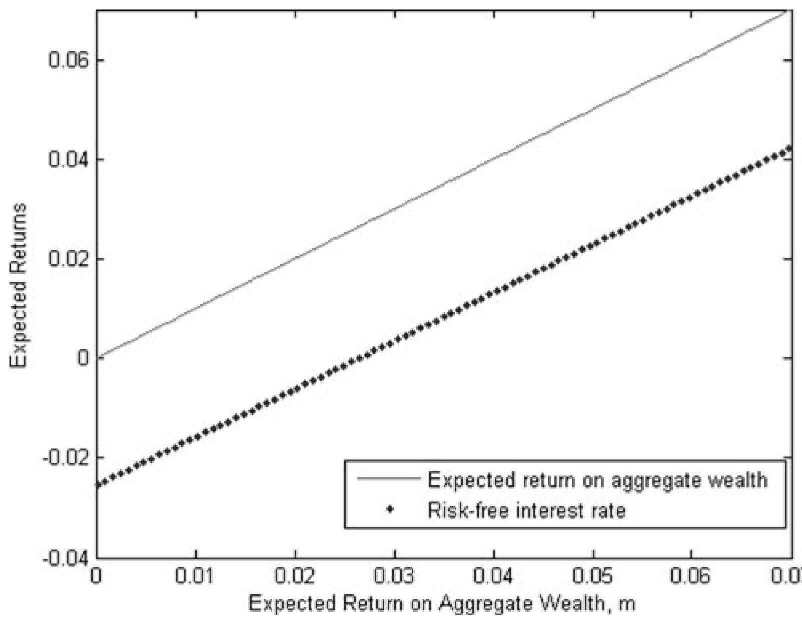

Figure 2. Expected returns as a function of *m*. The dotted line is the expected return on aggregate wealth, m. The dashed line is the risk-free interest rate as a function of m. Most of the variation in the expected return on aggregate wealth comes from variation in the risk-free interest rate. The graph is based on the accurate numerical solution using the Markov chain approximation method.

Lustig, Van Nieuwerburgh, and Verdelhan (2008). Since all moments in Lustig, Van Nieuwerburgh, and Verdelhan (2008) are calculated at the quarterly frequency, I simulate my continuous-time model and aggregate quantities to a quarterly frequency. This avoids the time aggregation issue and makes all moments directly comparable to each other.

My model reproduces the large volatility of the wealth–consumption ratio and the return on aggregate wealth in the data. Note that the quarterly standard deviation of the return on aggregate wealth produced by the Bansal and Yaron (2004) model, 1.64%, is much smaller than its empirical counterpart of 4.94%. My model closely matches the volatility of the return on aggregate wealth because the parameter σ_K is calibrated to target this moment. The model also produces a highly volatile wealth–consumption ratio process, as in the data. The quarterly standard deviation of the wealth–consumption ratio produced by the benchmark model, 11.01%, although still lower than the point estimate in Lustig, Van Nieuwerburgh, and Verdelhan (2008), 17.24%, is substantially higher than the same moment in the Bansal and Yaron (2004) model, 2.35%. The standard deviation of the first difference in the wealth–consumption ratio produced by the model, 4.34%, closely matches the same moment in the

Table I

Moments of the Wealth-to-Consumption Ratio

This table displays the unconditional moments of the log wealth–consumption ratio, $w - c$; its first difference, Δwc; the log return on the aggregate wealth, r_W; the consumption growth rate g_C; and the log risk-free interest rate, r_f. The expressions $E[\cdot]$ and $Std[\cdot]$ denote the means and standard deviations of the quantities in square brackets. The expressions $AC(1)[w - c]$ and $AC(4)[w - c]$ are, respectively, the first- and fourth-order autocorrelations of the log wealth–consumption ratio. The expression corr$[\Delta c, \Delta wc]$ stands for the correlation between consumption growth rates and the log wealth–consumption ratio. The second column (Data) contains moments estimated in Lustig, Van Nieuwerburgh, and Verdelhan (2008). The third column (BY) contains moments generated from simulations of the Bansal and Yaron (2004) model. Moments in the second and the third columns are reported in Lustig, Van Nieuwerburgh, and Verdelhan (2008). The last column contains moments generated from the benchmark model in this paper. The statistics in the last column are simulated from the continuous-time model and aggregated to a quarterly level. Moments in the third and the last columns are averages and standard errors (in parentheses) of 5,000 simulations of 220 quarters of data. Moments in the second column are the point estimates and standard errors (in parentheses) obtained by the bootstrap procedure described in Lustig, Van Nieuwerburgh, and Verdelhan (2008).

	Data	BY	Model
$E\,[w - c](s.e.)$	5.86 (0.49)	5.85 (0.01)	5.81 (0.28)
$Std\,[w - c](s.e.)$	17.24% (4.30)	2.35% (0.43)	11.01% (4.06)
$AC(1)\,[w - c](s.e.)$	0.96 (0.03)	0.91 (0.03)	0.88 (0.07)
$AC(4)\,[w - c](s.e.)$	0.85 (0.08)	0.70 (0.10)	0.65 (0.19)
$Std\,(\Delta wc)(s.e.)$	4.86% (1.16)	0.90% (0.05)	4.34% (0.38)
corr $[\Delta c, \Delta wc](s.e.)$	0.11 (0.06)	−0.06 (0.06)	0.43 (0.09)
$Std\,[g_C](s.e.)$	0.44% (0.03)	1.43% (0.08)	1.14% (0.18)
$Std\,[r_W](s.e.)$	4.94% (1.16)	1.64% (0.09)	4.97% (0.43)
$E\,[r_W - r_f](s.e.)$	0.54% (0.16)	0.40% (0.01)	0.66% (0.60)

data, 4.86%. The Bansal and Yaron (2004) model, on the other hand, produces variation in the wealth–consumption ratio of 0.90% per quarter.

Other statistics of the wealth–consumption ratio generated by the model are largely consistent with the data. The autocorrelations of the wealth–consumption ratio produced by the model are similar to those in Bansal and Yaron (2004), and are well within two standard deviations of their empirical estimates. In the data, the consumption growth rate and the first difference in the wealth–consumption ratio are moderately positively correlated (0.11 at the quarterly level). My model produces a larger correlation, 0.43 at the quarterly level, while the Bansal and Yaron (2004) model produces a small negative correlation. The quarterly risk premium on the aggregate wealth generated by my model, 0.66%, is slightly higher than the point estimate in Lustig, Van Nieuwerburgh, and Verdelhan (2008), 0.54%. The Bansal and Yaron (2004) model generates a slightly lower risk premium on aggregate wealth, 0.40%.

Table II compares the moments of the consumption growth rate in the data with those generated by the Bansal and Yaron (2004) model and those

Table II

Moments of Annualized Consumption Growth

This table compares the unconditional moments of annualized consumption growth in the data, in Bansal and Yaron's (2004) model, and in the benchmark incomplete information economy. The expressions $E[g_C]$ and $Std[g_C]$ denote the mean and standard deviation of annualized consumption growth rates. The expression $AC(j)$ is the jth autocorrelation. The second column (Data) contains the point estimates and the standard errors (in parentheses) of the estimates of these moments computed from U.S. Bureau of Economic Analysis data for the period 1929 to 1998 and are taken from Bansal and Yaron (2004) directly. The third column (BY) contains the consumption growth statistics reported in Bansal and Yaron (2004). They are generated from 1,000 simulations of 840 months of data in the Bansal and Yaron (2004) model and time-aggregated to an annual frequency. The last column (Model) contains the averages and standard errors (in parentheses) of the corresponding moments across 5,000 simulations of the benchmark model in this paper for 840 months. All simulations are done in the continuous-time model and aggregated to an annual frequency.

	Data	BY	Model
$E[g_C]$	1.8	1.8	1.79 (2.21)
$Std[g_C]$	2.93 (0.69)	2.72	2.80 (0.52)
$AC(1)$	0.49 (0.14)	0.48	0.48 (0.15)
$AC(2)$	0.15 (0.22)	0.29	0.28 (0.19)
$AC(5)$	−0.08 (0.10)	0.13	0.12 (0.06)
$AC(10)$	0.05 (0.09)	0.01	0.01 (0.07)

generated by my benchmark model. The estimates of the moments of the consumption growth rate are taken from Bansal and Yaron (2004). Since the estimation and calibration in Bansal and Yaron (2004) are done at an annual frequency, I simulate a continuous-time model and time-aggregate quantities to an annual frequency. Table II shows that the consumption growth statistics in my model are broadly consistent with their empirical counterparts, and with those produced in the Bansal and Yaron (2004) model. The mean and the standard deviation of the consumption growth rates in the model closely match their empirical counterparts. The autocorrelation functions of consumption growth rates are very similar to those in the Bansal and Yaron (2004) model and are largely consistent with the pattern in the data.

Learning and a high IES are together responsible for the large volatilities of the return on aggregate wealth and the wealth–consumption ratio generated by my model. To understand the intuition for this, note that equation (3) implies that a large covariance term, $\mathrm{cov}_t(r_{W,t+1}, w_{t+1} - c_{t+1})$, is necessary for generating the right pattern of volatility of the consumption growth rate, the wealth–consumption ratio, and the return on aggregate wealth. Following Campbell (1999), I write the innovation of the log wealth–consumption ratio as

$$(w_{t+1} - c_{t+1}) - E_t[w_{t+1} - c_{t+1}] = (\psi - 1)(E_{t+1} - E_t)\left[\sum_{j=0}^{\infty} \rho^i r_{W,t+j+1}\right].$$

Therefore, the covariance term in equation (3) can be written as

$$\begin{aligned}&\mathrm{cov}_t(r_{W,t+1}, w_{t+1} - c_{t+1})\\&\quad = (\psi - 1)\mathrm{cov}_t\left(r_{W,t+1} - E_t[r_{W,t+1}],\ (E_{t+1} - E_t)\left[\sum_{j=0}^{\infty}\rho^i r_{W,t+j+1}\right]\right).\end{aligned} \tag{57}$$

Equation (57) is exactly the discrete-time analogue of the variance decomposition in equation (56). The above equation shows that learning creates a positive covariance between the innovation in return, $r_{W,t+1} - E_t[r_{W,t+1}]$, and the innovation in expected future returns, $(E_{t+1} - E_t)[\sum_{j=0}^{\infty}\rho^i r_{W,t+j+1}]$. With $\psi > 1$, this implies a positive covariance between the return on aggregate wealth and the wealth–consumption ratio and helps produce a small volatility of consumption growth rates.

Equation (57) implies that the choice of the IES parameter is important for generating a small volatility of the consumption growth rates. Empirical evidence on the magnitude of the IES parameter is mixed. While Hansen and Singleton (1982), Attanasio and Weber (1989), and Vissing-Jorgensen (2002) estimate the IES parameter to be larger than one, other studies, for example, Hall (1988), Campbell (1999), and Browning, Hansen, and Heckman (1999), argue that the IES parameter is well below one. Bansal and Yaron (2004) calibrate the IES parameter to be 1.5 and argue that ignoring the time-varying volatility in consumption growth leads to serious downward bias in the estimates of the IES parameter. Bansal, Kiku, and Yaron (2007) estimate the IES parameter to be 2.43 with a standard deviation of 1.3. Key implications of the model for different choices of the IES parameter are reported in Table III. Changes in the IES parameter have a negligible effect on the statistical properties of asset returns. As shown in Table III, the effect of IES on the volatility of the log wealth–consumption ratio is quite significant. The volatility of the log wealth–consumption ratio is minimized when $\psi = 1$. In this case, the ratio of wealth to instantaneous consumption is constant by Proposition 1, and the volatility of the annualized log wealth–consumption ratio comes completely from the volatility of time-aggregated consumption. For $\psi < 1$, the wealth–consumption ratio is decreasing in m_t. Consequently, the volatility of the return on aggregate wealth and the volatility of the wealth–consumption ratio reinforce each other, making consumption growth more volatile than the return on aggregate wealth. For $\psi > 1$, as discussed in Section II.C, changes in the wealth–consumption ratio offset the volatility in the return on aggregate wealth because of learning. Therefore, as ψ increases, the volatility of the log wealth–consumption ratio rises, while the volatility of consumption growth decreases.

Table IV compares the model's implications for different choices of the information quality parameter σ_e. As I decrease σ_e from ∞ to zero, the mean consumption growth rate increases slightly. The volatility of consumption growth increases sharply from 2.80% to 13.1% per year. In the case with no learning, the volatility of consumption growth is higher than the volatility of the

Table III

The Role of Intertemporal Elasticity of Substitution

This table documents the effects of the choice of the IES parameter, ψ, on the key implications of the model. The expressions $Std\,[r_W]$, $Std[w-c]$, and $Std\,[\Delta wc]$ denote the standard deviations of the log return on aggregate wealth, the log wealth–consumption ratio, and the first difference in the log wealth–consumption ratio. The terms $E\,[g_C]$ and $Std\,[g_C]$ denote the mean and standard deviation of annualized consumption growth rates. $AC(1)$ is the first-order autocorrelation of consumption growth rates. $E\,[r_f]$ and $Std\,[r_f]$ denote, respectively, the annualized mean and standard deviation of the risk-free interest rate generated by the model. $E\,[r_W - r_f]$ stands for the annualized risk premium on aggregate wealth. All quantities are computed at an annual level. All reported statistics are the averages and standard deviations (in parentheses) of the corresponding moments across 5,000 simulations of the benchmark model for 70 years. All simulations are done in the continuous-time model and aggregated to an annual frequency.

ψ	0.5	1.0	1.5	2.0	2.5
			Wealth-to-Consumption Ratio		
$Std\,[r_W]$	10.3 (0.92)	10.3 (0.93)	10.2 (0.93)	10.3 (0.92)	10.3 (0.93)
$Std\,[w-c]$	12.2 (3.2)	5.74 (0.49)	10.1 (2.7)	18.5 (6.0)	27.5 (9.5)
			Consumption Growth Rate		
$E\,[g_C]$	1.63 (2.03)	1.69 (2.04)	1.69 (2.05)	1.79 (2.21)	1.85 (2.24)
$Std\,[g_C]$	11.1 (1.0)	8.13 (0.75)	5.36 (0.53)	2.80 (0.52)	2.04 (0.62)
$AC(1)$	0.23 (0.11)	0.25 (0.11)	0.29 (0.12)	0.48 (0.15)	0.77 (0.14)
			Returns		
$E\,[r_f]$	0.84 (1.02)	0.86 (1.00)	0.84 (0.96)	0.86 (0.98)	0.85 (0.94)
$Std\,[r_f]$	0.80 (0.28)	0.79 (0.27)	0.78 (0.27)	0.77 (0.26)	0.75 (0.26)
$E\,[r_W - r_f]$	2.69 (1.22)	2.68 (1.24)	2.69 (1.25)	2.68 (1.26)	2.69 (1.27)

Table IV

The Effect of Information Quality

This table documents the effects of information quality on the moments of consumption growth, the risk premium on aggregate wealth, the risk-free interest rate, and the wealth–consumption ratio for various values of the information quality parameter, σ_e. The terms $E[g_C]$ and $Std[g_C]$ denote the mean and standard deviation of annualized consumption growth rates. The expressions $E[r_W - r_f]$ and $E[r_f]$ are, respectively, the annualized risk premium on the aggregate wealth and the annualized risk-free interest rate. The expressions $Std[r_W]$, $Std[r_f]$, and $Std[w-c]$ denote the annualized standard deviations of the return on aggregate wealth, the risk-free interest rate, and the log wealth–consumption ratio. All reported statistics are the averages and standard errors (in parentheses) of the corresponding moments across 5,000 simulations of the benchmark model for 70 years. All simulations are done in the continuous-time model and aggregated to an annual frequency.

σ_e	$E[g_C]$	$Std[g_C]$	$E[r_W - r_f]$	$Std[r_W]$	$E[r_f]$	$Std[r_f]$	$Std[w-c]$
0.00	1.92 (2.60)	13.1 (1.62)	1.96 (1.23)	10.3 (1.0)	1.57 (1.66)	1.33 (0.48)	32.6 (11.6)
0.01	1.91 (2.55)	11.8 (1.51)	2.07 (1.23)	10.3 (1.0)	1.47 (1.58)	1.25 (0.45)	30.1 (10.7)
0.02	1.88 (2.55)	10.8 (1.47)	2.18 (1.24)	10.2 (1.0)	1.35 (1.51)	1.17 (0.42)	28.3 (10.1)
0.05	1.81 (2.37)	8.41 (1.29)	2.39 (1.24)	10.3 (1.0)	1.13 (1.30)	1.02 (0.37)	24.6 (8.7)
0.10	1.80 (2.27)	5.91 (1.13)	2.55 (1.25)	10.3 (1.0)	0.97 (1.12)	0.88 (0.32)	21.3 (7.4)
∞	1.79 (2.21)	2.80 (0.52)	2.68 (1.26)	10.3 (1.0)	0.86 (0.98)	0.77 (0.26)	18.5 (6.0)

return on aggregate wealth because changes in the wealth–consumption ratio are independent of the innovations in the return on aggregate wealth and add to the volatility of consumption growth. As σ_e changes from ∞ to zero, the risk premium on aggregate wealth decreases from 2.68% to 1.96%, despite the huge volatility in consumption growth rates associated with small values of σ_e. Furthermore, the mean risk-free interest rate increases from 0.86% to 1.57%, which echoes the change in the equity premium. The standard deviation of the risk-free interest rate increases from 0.77% to 1.33%. These all confirm the theoretical results obtained in Sections II.B and II.C.

The risk premium on aggregate wealth produced by the model is fairly low, 2.68% per year. To address the equity premium puzzle (Mehra and Prescott (1985)), in the next subsection I examine whether the benchmark model is able to reproduce the key features of the market risk premium when the dividend process of the market equity is calibrated to the data.

C. The Equity Premium Puzzle

Following Campbell (1996) and Bansal and Yaron (2004), I model the aggregate dividend process and the aggregate consumption process separately. The market return is assumed to be the return on the claim to the dividend process $\{D_t : t \geq 0\}$. The log dividend is assumed to satisfy

$$d \ln D_t = \phi d \ln C_t - A dt + \sigma_D dB_{D,t},$$

where the Brownian motion $B_{D,t}$ is independent of all other shocks in the economy. As in Bansal and Yaron (2004), the parameter ϕ captures the idea of leverage. The parameters ϕ, A, and σ_D are chosen as follows: $A = 0.034$, $\phi = 2.05$, $\sigma_D = 0.11$. The parameter A determines the mean growth rate of dividends and has virtually no effect on the risk premium; $A = 0.034$ is chosen to match the mean growth rate of dividends in the data. The parameters ϕ and σ_D are chosen to jointly match the standard deviation of the dividend growth rate, and the correlation between the dividend growth rate and the consumption growth rate in the data. The moments of dividend growth and the market return are reported in Table V. The autocorrelation of the dividend growth rate produced by this parameter choice is slightly higher than its counterpart in the data, but is similar to the same moment produced by the Bansal and Yaron (2004) model and well within one standard deviation of its point estimate.

I use a low risk-aversion parameter, $\gamma = 2$. In fact, the high volatility of the return on aggregate wealth and the moderate risk premium on the return on aggregate wealth suggests that risk aversion must be low. To see this, note that in any economy where the expected return on wealth is a Markov diffusion process, Lemma 1 and equation (42) imply that the risk premium on aggregate wealth be given by

$$\mu_{W,t} - r_t = \gamma \operatorname{var}_t\left[\frac{dW_t}{W_t}\right] - \frac{U_{W,m}(W_t, m_t)}{U_W(W_t, m_t)} \operatorname{cov}_t\left[\frac{dW_t}{W_t}, dm_t\right]. \tag{58}$$

Table V

The Risk Premium of the Market Return

This table reports the key asset pricing statistics in the data, in the Bansal and Yaron (2004) model, and in the benchmark incomplete information economy. All moments are annualized. The expressions $Std[g_D]$ and $AC(1)$ denote the standard deviation and the first-order autocorrelation of log dividend growth rate. $\mathrm{corr}(g, g_D)$ is the correlation of log dividend growth and log consumption growth. $E[r_f]$ and $E[r_M - r_f]$ are, respectively, the annualized mean risk-free rate and the annualized risk premium on the market return. $E[\exp(p-d)]$ is the average price-to-dividend ratio of the market portfolio. $Std(r_f)$, $Std(r_M)$, and $Std[p-d]$ denote the standard errors of the risk-free interest rate, the market return, and the log price-to-dividend ratio, respectively. The second column (Data) contains the point estimates and standard errors (in parentheses) of these moments provided in Bansal and Yaron (2004). The third column (BY) contains the moments generated by the Bansal and Yaron (2004) model reported in Bansal and Yaron (2004). The last column (Model) contains the moments generated by the benchmark model in this paper. Moments in the last column are averages and standard deviations (in parentheses) of 5,000 simulations of 840 months of data. The simulation is done in continuous time and aggregated to an annual frequency.

	Data	BY	Model
$Std[g_D]$	11.49% (1.98)	10.96%	10.65% (11.09)
$AC(1)$	0.21 (0.13)	0.33	0.30 (1.3)
$\mathrm{corr}(g, g_D)$	0.55 (0.34)	0.31	0.53 (0.11)
$E[r_f]$	0.86% (0.42)	0.93	0.86% (0.98)
$Std(r_f)$	0.97% (0.28)	0.57	0.77% (0.26)
$E[r_M - r_f]$	6.33% (2.15)	6.84	5.06% (3.14)
$Std[r_M]$	19.42% (3.07)	18.65	23.20% (3.27)
$E[\exp(p-d)]$	26.56 (2.53)	19.98	27.22 (1.50)
$Std[p-d]$	0.29 (0.04)	0.21	0.34 (0.12)

Lustig, Van Nieuwerburgh, and Verdelhan (2008) estimate the quarterly standard deviation of return on aggregate wealth to be 4.94%. Assuming homoskedasticity, this alone will generate a 0.49% risk premium on aggregate wealth with $\gamma = 2$ through the myopic demand component of the equity premium in (58). In the data, the risk premium on aggregate wealth is 0.54% per quarter (Lustig, Van Nieuwerburgh, and Verdelhan (2008)). As long as the covariance between the return and the expected return is positive, higher risk aversion will generate too much risk premium in aggregate wealth. In fact, in my model, the second term is fairly large because of learning, and consequently the model overstates the risk premium on aggregate wealth: 0.66% on a quarterly level. In Bansal and Yaron (2004), the volatility of the return on aggregate wealth is too low relative to the estimate in Lustig, Van Nieuwerburgh, and Verdelhan (2008), and thus they have to use a large risk-aversion parameter to produce a large market price of risk.

The model generates a risk premium on the market return of 5.06% per year with a risk-aversion parameter of two. Both the transitory and the persistent components of the return on technology are responsible for the large risk premium. Using log-linearization, by equation (42), the risk premium on an asset with cumulative return R_t can be written as

$$[\mu_R(m_t) - r(m_t)] = \gamma \mathrm{cov}_t \left[dR_t, \frac{dW_t}{W_t} \right] + \frac{\gamma - 1}{a + \kappa_1} \mathrm{cov}_t[dR_t, dm_t]. \tag{59}$$

The above equation makes it clear that a high volatility of return on aggregate wealth, σ_K, enhances the myopic demand component of the risk premium, and a low mean reversion parameter, a, strengthens the hedging demand component of the equity premium. As explained in Section III.B, the volatility of the return on aggregate wealth in my model is much larger than that in Bansal and Yaron (2004). The expected return on aggregate wealth is also more persistent in my model. The monthly autocorrelation of m_t is $e^{-\frac{a}{12}} = 0.998$, higher than that used in Bansal and Yaron (2004), 0.979. The hedging demand component of the equity premium is further enhanced by learning. Consequently, the model generates a significant risk premium on the market return despite the low risk aversion.

My model abstracts from time-varying volatility in consumption growth rates. The volatility of return on aggregate wealth is constant by assumption. Numerical results show that the variation in the instantaneous volatility of consumption growth rate is negligible. Bansal and Yaron (2004) demonstrate that time-varying volatility is important in accounting for many stylized empirical facts. First, the time-varying volatility risk in consumption growth contributes to a sizable portion of the equity premium in their calibration. My model understates the risk premium on the market return: 5.06% per year compared to 6.84% per year in Bansal and Yaron's (2004) calibration. Incorporating time-varying volatility in consumption growth will increase the risk premium produced by the model. Second, in Bansal and Yaron (2004), time-varying volatility helps generate the predictability of excess returns by dividend yield. As explained earlier, in my model, since the wealth–consumption ratio, the risk-free interest rate, and the risk premium are all increasing functions of the state variable m, a high dividend yield forecasts a low risk premium and low future returns in the model. These implications of the model are clearly inconsistent with empirical evidence. Since the main focus of the paper is on the effect of learning on the unconditional moments of the wealth–consumption ratio and the return on aggregate wealth, I abstract from time-varying volatility and do not attempt to address empirical evidence on the conditional moments of stock returns. A richer setup could potentially further improve the performance of the model in terms of the level of the risk premium on the market and the predictability of stock returns.

IV. Comparison with Pure Exchange Economies

The relationship between information quality and the equity premium has been studied by Veronesi (2000), Brennan and Xia (2001), and Brevik and D'Addona (2007), among others, in pure exchange economies. In pure exchange economies with CRRA utility, Veronesi (2000) establishes the result that learning leads to a lower (higher) equity premium if RRA is higher (lower) than one, which is the opposite of my result in equation (52). In this section, I consider a pure exchange economy and compare the relationship between information quality and the equity premium in my production economy with that in pure exchange economies.

In the linear production economy studied in this paper, as discussed in Section II.B, learning affects the hedging demand component of the equity premium, and the direction of this effect depends on RRA. Learning does not affect the myopic demand component of the equity premium. This is because the return on aggregate wealth is exogenously determined by the linear production technology, and the sign of the myopic demand component of the equity premium depends only on risk aversion and the volatility of the return on aggregate wealth, as shown in equation (42).

In the rest of the section, I demonstrate that in pure exchange economies, learning affects not only the hedging demand component, but also the myopic demand component of the equity premium. The sign of the effect of learning on the hedging demand component of the equity premium depends only on RRA, whereas the direction in which learning affects the myopic demand component of the equity premium is determined by the IES parameter. I also show that in the case of CRRA, the effect on the myopic demand component of the equity premium always dominates. Consequently, my result confirms Veronesi's (2000) finding and further reveals that Veronesi's (2000) result is driven by the agent's IES, not RRA.

Consider a pure exchange economy in which the representative agent has the Kreps–Porteus SDU as specified in (4). The endowment process of the economy is assumed to be

$$dY_t = Y_t[\theta_t dt + \sigma_Y dB_{Y,t}], \tag{60}$$

where $\{\theta_t\}_{t\geq 0}$ follows the same process as described in (11). Again, assume

$$\mathrm{corr}(B_Y, B_\theta) = \rho.$$

If the agent does not observe $\{\theta_t\}_{t\geq 0}$, she must update her belief based on the observed consumption process and an additional source of information denoted by $\{e_t\}_{t\geq 0}$, where

$$de_t = \theta_t dt + \sigma_e dB_{e,t}.$$

I again assume that B_e is independent of $[B_Y, B_\theta]$. The posterior mean of θ obeys the following stochastic differential equation (again, assuming that the posterior variance, Q_E, starts at its steady-state level given in equation (IA.16) in the Internet Appendix):

$$dm_t = a(\overline{\theta} - m_t)dt + \left(\frac{1}{\sigma_K} Q_E + \rho\sigma_\theta\right) d\widetilde{B}_{Yt} + \frac{1}{\sigma_e} Q_E d\widetilde{B}_{e,t}, \tag{61}$$

where $\widetilde{B}_{Yt}$ and $\widetilde{B}_{e,t}$ are innovation processes defined by

$$d\widetilde{B}_{Y,t} = \frac{1}{\sigma_Y}\left[\frac{dY_t}{Y_t} - m_t dt\right], \quad d\widetilde{B}_{e,t} = \frac{1}{\sigma_e}[de_t - m_t dt]$$

and Q_E is the steady-state posterior variance of θ in the exchange economy. The expression of Q_E is given in the Internet Appendix.

Let W_t denote the wealth of the representative agent. Homogeneity implies that the value function of the agent's optimal portfolio choice problem can be written as

$$U(W,m) = \frac{1}{1-\gamma} G(m)^{\frac{1}{\psi}} W^{1-\gamma}, \tag{62}$$

where the function $G(m)$ satisfies an ODE given in the Internet Appendix. Using equation (44), the myopic demand component of the risk premium on aggregate wealth, denoted MD_t, is written as

$$MD_t = \gamma \mathrm{var}_t \left[\frac{dW_t}{W_t} \right]. \tag{63}$$

Similarly, equation (44) implies that the hedging demand component of the risk premium on aggregate wealth is given by

$$HD_t = -\frac{1}{\psi} \frac{G'(m_t)}{G(m_t)} \mathrm{cov}_t \left[\frac{dW_t}{W_t}, dm_t \right]. \tag{64}$$

In the Internet Appendix, I derive the following log-linearization approximation of equations (63) and (64):

$$MD_t \approx \gamma \left\{ \sigma_C^2 + \frac{\left(1-\frac{1}{\psi}\right)^2}{(a+\varpi_1)^2} \left(\sigma_\theta^2 - 2aQ_E\right) + 2\frac{1-\frac{1}{\psi}}{a+\varpi_1} (\rho \sigma_C \sigma_\theta + Q_E) \right\}, \tag{65}$$

$$HD_t \approx -\frac{1}{\psi} \left\{ \frac{\left(1-\frac{1}{\psi}\right)(1-\gamma)}{(a+\varpi_1)^2} \left(\sigma_\theta^2 - 2aQ_E\right) + \frac{1-\gamma}{a+\varpi_1} (\rho \sigma_C \sigma_\theta + Q_E) \right\}. \tag{66}$$

Equations (65) and (66) reveal the effect of information quality on the myopic demand component and the hedging demand component of the equity premium. First, if $\psi > 1 (\psi < 1)$, learning increases (decreases) the myopic demand component of the equity premium by increasing (decreasing) the volatility of the return on wealth. To see this, using log-linear approximation around $\psi = 1$,

$$\frac{\partial}{\partial \sigma_e} \left[MD_t \right] \approx 2\gamma \left(1 - \frac{1}{\psi} \right) \frac{(\varpi_1 + a/\psi)}{(a+\varpi_1)^2} \frac{\partial Q_E}{\partial \sigma_e}, \tag{67}$$

where $\varpi_1 > 0$ is a constant defined in the Internet Appendix.[7] By (24), it is clear that the myopic demand component of the equity premium is increasing in σ_e if $\psi > 1$.

[7] Details of the log-linear approximation can be found in the Internet Appendix.

When $\psi > 1$, learning increases the myopic demand component of the risk premium on aggregate wealth by raising the volatility of the return on aggregate wealth. To see the intuition for this, consider the following variance decomposition:

$$\mathrm{var}_t\left(d\ln W_t\right) = \mathrm{var}_t\left(d\ln C_t\right) + \mathrm{var}_t\left(\ln x_t\right) + 2\mathrm{cov}_t\left(d\ln C_t,\ d\ln x_t\right), \quad (68)$$

where $x_t = \frac{W_t}{Y_t}$ is the wealth–consumption ratio of the representative agent. Here, changes in σ_e affect the variance of the return on wealth mainly through the third term in (69). Learning creates a positive covariance between innovations in consumption growth and innovations in expected consumption growth m. If $\psi > 1$, then x is an increasing function (Proposition 1), and therefore learning makes the term $\mathrm{cov}_t(\ln C_t,\ \ln x_t)$ negative. Consequently, the variance of the return on wealth increases because of learning. Of course, if $\psi < 1$, the same argument implies that learning decreases the volatility of the return on aggregate wealth.

The optimal wealth–consumption ratio of the representative agent is increasing in the expected return on wealth if $\psi > 1$. In the linear production economy that I focus on in this paper, the volatility of the return on wealth is completely determined by the exogenous technology; therefore, learning does not affect the myopic demand component of the risk premium on aggregate wealth. In this case, the positive covariance between the return on aggregate wealth and the wealth–consumption ratio induced by learning works to reduce the volatility of the endogenously determined consumption process. In the pure exchange economy above, the volatility of the consumption growth rate is exogenous. The positive covariance between the realized consumption growth rate and the wealth–consumption ratio resulting from learning works to raise the volatility of the return on aggregate wealth and hence increases the myopic demand component of the equity premium.

Second, learning increases the hedging demand component of the equity premium if $\gamma > 1 (\gamma < 1)$. Again, using a log-linear approximation around $\psi = 1$,

$$\frac{\partial}{\partial \sigma_e}\left[HD_t\right] \approx \frac{1}{\psi}\left(\gamma - 1\right)\frac{\left[a\left(2/\psi - 1\right) + \varpi_1\right]}{\left(a + \varpi_1\right)^2}\frac{\partial Q_E}{\partial \sigma_e}. \quad (69)$$

Therefore, for ψ close to one, the hedging demand component of the equity premium is increasing in σ_e if $\gamma > 1$. The intuition for this is the same as discussed earlier in Section II.B.

In the case of CRRA preferences, $\gamma = \frac{1}{\psi}$, the effect on the myopic demand component of the equity premium always dominates. (One can see this by letting $\gamma = \frac{1}{\psi}$ and comparing (67) with (69).) This is why Veronesi (2000) obtains the result that lower information quality decreases (increases) the equity premium if $\gamma > 1 (\gamma < 1)$. In the general case of Kreps–Porteus utility, the direction of the effect of learning on the myopic demand component of the equity premium depends on the IES parameter, not the RRA parameter. In this sense, Veronesi's (2000) result is driven by the agent's attitude toward intertemporal substitution.

V. Conclusion

I study the asset pricing implications of information quality about the long-run growth rate of the economy in a simple production economy. I show that lower information quality increases the risk premium on aggregate wealth if the representative agent's risk aversion is higher than one, and reduces the volatility of equilibrium consumption growth if IES is larger than one. The effects of learning are quantitatively significant and lead to substantial improvement upon the Bansal and Yaron (2004) model in terms of its predictions about the volatility of the return on aggregate wealth and the wealth–consumption ratio.

Several remarks are in order. First, the assumption of the linear production technology is made for analytical convenience. Under this assumption, the direction in which information quality affects the equity premium depends only on the RRA parameter, and the sign of the effect of information quality on the volatility of consumption growth is determined completely by the IES parameter. In pure exchange economies, or production economies with adjustment costs, the direction of the effect of information depends on both parameters. However, the basic intuition developed in this paper will still be useful in these more general settings. In particular, when both RRA and IES are larger than one, learning is likely to increase the risk premium on aggregate wealth, since the effect of learning on the hedging demand component and that on the myopic demand component of the equity premium work in the same direction. The simple setting of linear technology allows for a complete separation of the roles of RRA and IES.

Second, all asset pricing implications of the model can be obtained by assuming exogenously the conditional distribution of the relevant state variables. Learning endogenizes the conditional distributions and imposes discipline on the choice of the distributions.

Finally, the choice of the linear production technology clearly has its limitations. As is well known, incorporating production presents additional challenges for general equilibrium models to generate a realistic equity premium (Rouwenhorst (1995)). One reason is that the physical capital in the data is very smooth. In order to generate large fluctuations in the equity price, one has to rely on adjustment costs as in Jermann (1998) or on real rigidity as in Boldrin, Christiano, and Fisher (2001). In the absence of investment frictions and a labor market, my model is not rich enough to address the empirical evidence on investment dynamics and other stylized business cycle facts, but this topic would be an interesting one for future research.

REFERENCES

Attanasio, Orazio P., and Guglielmo Weber, 1989, Intertemporal substitution, risk aversion and the Euler equation for consumption, *Economic Journal* 99, 59–73.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2007, Risks for the long run: Estimation and inference, Working paper, Duke University and University of Pennsylvania.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Boldrin, Michele, Lawrence J. Christiano, and Jonas D. M. Fisher, 2001, Habit persistence, asset returns, and the business cycle, *American Economic Review* 91, 149–166.

Breeden, Douglas T., 1979, An intertemporal asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* 7, 265–296.

Brennan, Michael J., and Yihong Xia, 2001, Stock price volatility and equity premium, *Journal of Monetary Economics* 47, 249–283.

Brevik, Frode, and Stefano D'Addona, 2007, Information quality and stock returns revisited, Working paper, University of St. Gallen and University of Rome.

Browning, Michael J., Lars Peter Hansen, and James J. Heckman, 1999, Micro data and general equilibrium models, Vol. 1A of *Handbook of Macroeconomics* (Elsevier Science Publishers, Amsterdam), 543–633.

Campbell, John Y., 1996, Understanding risk and return, *Journal of Political Economy* 104, 298–345.

Campbell, John Y., 1999, Asset prices, consumption, and the business cycle, Vol. 1C of *Handbook of Macroeconomics* (Elsevier Science Publishers, Amsterdam), 1231–1303.

Campbell, John Y., George Chacko, Jorge Rodriguez, and Luis M. Viceira, 2004, Strategic asset allocation in a continuous-time VAR model, *Journal of Economic Dynamics and Control* 28, 2195–2214.

Campbell, John Y., and Robert J. Shiller, 1989, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

Cochrane, John, 2005, *Asset Pricing* (Princeton University Press, Princeton, NJ).

Croce, Mariano M., Martin Lettau, and Sydney C. Ludvigson, 2007, Investor information, long-run risk, and the duration of risky cash flows, Working paper 1291. National Bureau of Economic Research.

Detemple, Jerome B., 1986, Asset pricing in a production economy with incomplete information, *Journal of Finance* 41, 383–391.

Dothan, Michael U., and David Feldman, 1986, Equilibrium interest rates and multiperiod bonds in a partially observable economy, *Journal of Finance* 41, 369–382.

Duffie, Darrell, 2001, *Dynamic Asset Pricing Theory* (Princeton University Press, Princeton, NJ).

Duffie, Darrell, and Larry G. Epstein, 1992a, Asset pricing with stochastic differential utility, *Review of Financial Studies* 5, 411–436.

Duffie, Darrell, and Larry G. Epstein, 1992b, Stochastic differential utility, *Econometrica* 60, 353–394.

Duffie, Darrell, and Pierre-Louis Lions, 1992, PDE solutions of stochastic differential utility, *Journal of Mathematical Economics* 21, 577–606.

Epstein, Larry G., and Stanley E. Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Gennotle, Gerod, 1986, Optimal portfolio choice under incomplete information, *Journal of Finance* 41, 733–746.

Gollier, Christian, and Edward E. Schlee, 2006, Information and the equity premium, Working paper, University of Toulouse and Arizona State University.

Hall, Robert E., 1988, Intertemporal substitution in consumption, *Journal of Political Economy* 96, 339–357.

Hansen, Lars Peter, and Thomas J. Sargent, 2009, Fragile beliefs and the price of model uncertainty, Working paper, University of Chicago and New York University.

Hansen, Lars Peter, Thomas J. Sargent, and Thomas D. Tallarini Jr., 1999, Robust permanent income and pricing, *Review of Economic Studies* 66, 873–907.

Hansen, Lars Peter, and Kenneth Singleton J., 1982, Generalized instrumental variables estimation of nonlinear rational expectations models, *Econometrica* 50, 1269–1286.

Jermann, Urban J., 1998, Asset pricing in production economies, *Journal of Monetary Economics* 41, 257–275.

Kreps, David M., and Evan L. Porteus, 1978, Temporal resolution of uncertainty and dynamic choice theory, *Econometrica* 46, 185–200.

Kushner, Harold J., and Paul G. Dupuis, 2001, *Numerical Methods for Stochastic Control Problems in Continuous Time* (Springer-Verlag, New York).

Liptser, Robert S., and Albert N. Shiryaev, 2001, *Statistics of Random Processes* (Springer-Verlag, New York).

Lucas, Robert E. Jr., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.

Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan, 2008, The wealth-consumption ratio, Working paper, University of California, New York University, and Boston University.

Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.

Merton, Robert C., 1971, Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory* 3, 373–413.

Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–887.

Rouwenhorst, K. Geert, 1995, Asset pricing implications of equilibrium business cycle models, in T. F. Cooley ed.: *Frontiers of Business Cycle Research* (Princeton University Press, Princeton, NJ).

Schroder, Mark, and Costis Skiadas, 1999, Optimal consumption and portfolio selection with stochastic differential utility, *Journal of Economic Theory* 89, 68–126.

Tallarini, Thomas D., 2000, Risk sensitive real business cycles, *Journal of Monetary Economics* 45, 507–532.

Veronesi, Pietro, 1999, Stock market overreaction to bad news in good times: A rational expectations equilibrium model, *Review of Financial Studies* 12, 975–1007.

Veronesi, Pietro, 2000, How does information quality affect stock returns? *Journal of Finance* 55, 807–837.

Vissing-Jorgensen, Annette, 2002, Limited asset market participation and the elasticity of intertemporal substitution, *Journal of Political Economy* 110, 825–853.

Weil, Philippe, 1989, The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* 24, 401–422.