Rational Ignorance in Long-run Risk Models

Stefano d'Addona[∗]
University of Roma Tre, Italy

Abstract

This paper documents an unpleasant feature of Epstein-Zin preferences in long-run risk models that are widespread in asset pricing: Agents with a preference for an early resolution of uncertainty achieve higher utility if they ignore information on the state of the business cycle.

Key words: recursive preferences; Epstein-Zin preferences; uncertainty aversion; information processing; time inconsistency

JEL classification: D83; D84; E32

1. Introduction

The separation of preference parameters governing the elasticity of intertemporal substitution and relative risk aversion permitted by the Epstein-Zin utility function (Epstein and Zin, 1989; Weil, 1989) has proven very fruitful in the asset pricing literature. Recent successful asset pricing models rely on calibrations of the utility function whereby the representative agent has both a high level of risk aversion and a high elasticity of intertemporal substitution. To name a few, Campbell and Vuolteenaho (2004) use such a calibration to explain stock market anomalies, Piazzesi and Schneider (2007) employ such a calibration to explain the average shape of the yield curve, and Lettau et al. (2008) utilize such a calibration to explain the run up in stock prices during the late 1990s. In a seminal paper, Bansal and Yaron (2004) forcefully demonstrate that in an exchange economy with a long-run risk component in consumption–that is, when the growth rate of the endowment good follows a trend stationary process with low conditional volatility and high unconditional volatility-such utility specifications produce both a low risk-free rate and a plausible risk premium for equity. Their paper has spawned a large body of research, which I will refer to as the long-run risk literature. With no pretension of a complete list, Kaltenbrunner and Lochstoer (2010) and Croce (2014) show how a

predictable component to consumption growth rates can arise in production economies. Colacito and Croce (2011) and Bansal and Shaliastovich (2013) study international linkages in open economies when each country has a small predictable component in its consumption growth rate trend. Hasseltoft (2012), Doh (2013), and Wu (2008) look at interest rate implications of long-run risk models. Constantinides and Ghosh (2011) and Rangvid et al. (2009) provide explicit estimation of long-run risk models. Bansal et al. (2012) and Cederburg and Hore (2008) analyze the extent to which a long-run risk model can explain predictability in the cross-secional data on security returns.

In this paper I go through the following thought experiment. I place an agent with Epstein-Zin preferences in a stylized complete markets endowment economy of the type analyzed in the long-run risk literature and give her the option to not incorporate any type of news when forming posterior beliefs about the current state of the trend consumption growth rate. If she chooses to do so, then her information set includes all the hyper-parameters of the economy and her current consumption level, but does not include any information that would help her determine the current level of the stochastic trend growth rate of the endowment good. I assume that her preference parameters are the same as those of the representative agent of the economy. Knowing that her information is coarser than that of other agents, she does not trade actively in a way that can be exploited by more informed agents. She keeps all her assets in the market portfolio and consumes the same as the representative agent. In other words, she holds only claims to the Lucas tree and consumes its fruit every period. The consumption profile of this agent mirrors that of the representative agent in the economy. The only way the coarser information set she uses influences her utility level is through the timing of information about future consumption.

I find that, for model parameters similar to those used in the asset pricing literature, the continuation value for the coarser information set is much higher than that for an agent whose information set also includes the current trend growth level of consumption. For a calibration that draws on Bansal and Yaron (2004), I find that utility gains from committing to using a coarser information set are equivalent to a 40% increase in lifetime consumption.

The numbers I find are not only surprising in their magnitude, but also in their direction. The parameterization I look is are such that the agent would be classified as having a *preference for early resolution of uncertainty* according to a common taxonomy. (See e.g. Kocherlakota, 1990 or Skiadas, 1998.) The presumed plausibility of utility functions generating a preference for early resolution of uncertainty seems to lend credence to Epstein-Zin preferences. It is worth noting that this paper discusses early resolution of uncertainty about consumption itself. The seminal article by Kreps and Porteus (1978) motivates the preference for early resolution over lotteries by stating that it is natural to prefer to know your income earlier so that you can better budget it for consumption purposes. In the endowment economy considered, equilibrium consumption of the representative agent is always going to be equal to the endowment stream, and so early resolution of uncertainty

does not provide any means for better budgeting since consumption is unaffected by it. Any preference for resolution of uncertainty in such economies must come directly from the way the distribution of possible consumption paths is aggregated to a certainty equivalent.

Figure 1 illustrates the above concept, which is taken from Kocherlakota (1990). An agent with a preference for early resolution of uncertainty would prefer tree (a) to tree (b): the two trees offer the same distribution of outcomes at each point in time, but in tree (a) time 2 consumption is revealed one period earlier. The phrase "preference for early resolution of uncertainty" seems to suggest that an agent would like to process any information on the current state of the economy, because it reduces uncertainty about her future consumption. For the parameters used in the long-run risk asset pricing literature, it is shown that always processing information is optimal in the sense of being a Nash strategy. However, if there is a persistent component in the consumption growth rate trend, then consumers can achieve an even higher utility level by committing to not processing information at any point in the future.

One way to understand our result is by noting that agents with relatively high risk aversion also dislike a positive correlation between current consumption growth and expected consumption growth (Piazzesi and Schneider, 2007). On the one hand, the consumer faces more consumption uncertainty when she relies on the coarser information set. On the other hand, relying on the coarser information set also shuts down any correlation between current and expected future consumption growth rates. In the simple economy studied, the second effect dominates, and so ignorant agents achieve, on average, a higher utility level.

Figure 1. Preferences for Resolution of Uncertainty

(a) Early resolution   (b) Late resolution
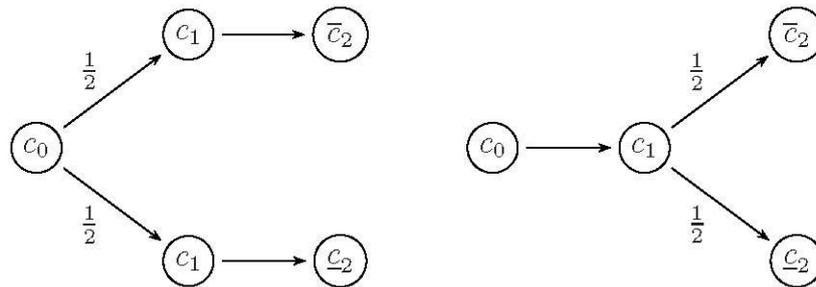
$$(a) \succ (b) \Longleftarrow \Rightarrow EIS > 1/CRRA$$
$$(a) \prec (b) \Longleftarrow \Rightarrow EIS < 1/CRRA$$

The rest of this paper runs as follows. Section 2 introduces the stylized long-run risk model used in the analysis. Section 3 analyzes the process of information acquisition by an agent facing the possibility to not incorporate any type of news about the current state of the trend consumption growth rate. Section 4 shows that learning the growth rate of the economy is a Nash strategy, but that ignorance can be supported as an equilibrium strategy when it yields a higher utility level. Section 5 quantifies the utility gains from ignorance using standard calibrations from the long-run risk asset pricing literature. Section 6 concludes.

2. A Stylized Economy with Long-run Risk

Our laboratory is a simple endowment economy where the growth rate of the log of the representative agent's consumption is the sum of an AR(1) component and a white noise shock. I base the set-up on Hansen et al. (2008), and the exposition closely follow theirs.

2.1 Endowment Process

Let ϵ_t and w_t be two series of i.i.d. standard normal innovation terms. Log consumption follows a random walk plus a time-varying drift. The first difference of the drift is:

$$c_{t+1} - c_t = \mu_c + x_t + \sigma_c \epsilon_{t+1}. \tag{1}$$

The log consumption growth rate trend at time t is a combination of a constant (μ_c) and a time-varying component x_t. Here, x follows an AR(1) process given by:

$$x_{t+1} = \kappa x_t + \sigma_x w_{t+1}, \tag{2}$$

where κ is the level of persistence for the consumption process.

2.2 Preferences

All agents in the economy are ex-ante identical with preferences over consumption paths given by the recursion:

$$V_t = \left[(1-\beta)(C_t)^{1-\rho} + \beta \mathcal{R}_t (V_{t+1})^{1-\rho} \right]^{\frac{1}{1-\rho}}, \tag{3}$$

where ρ is equal to the reciprocal of the Elasticity of Intertemporal Substitution(EIS), and β is the subjective discount factor. The risk adjustment Rt is also of the constant elasticity of substitution type:

$$\mathcal{R}_t\left(V_{t+1}\right) = \mathrm{E}_t\left[\left(V_{t+1}\right)^{1-\theta}\right]^{\frac{1}{1-\theta}} \Big| \tag{4}$$

,

where θ is the coefficient of relative risk aversion. Given the process assumptions above, Vt is homogeneous of degree 1 in the level of consumption. Let v_t denote the logarithm of the continuation value normalized by the consumption level. I rewrite the recursion above as:

$$v_t = \frac{1}{1-\rho}\log\left[(1-\beta) + \beta\exp[(1-\rho)\mathcal{Q}_t(v_{t+1} + c_{t+1} - c_t)]\right], \tag{5}$$

where the operator Qt is:

$$\mathcal{Q}_t(v_{t+1}) = \frac{1}{1-\theta}\log E_t\left[\exp\left((1-\theta)v_{t+1}\right)\right]$$

I distinguish two main information sets for the consumer. Under the coarser information set $\mathcal{F}_t^I$, the consumer is endowed with information about all the model hyper-parameters and the current consumption level. The alternative information set $\mathcal{F}_t^N$ is a refinement of $\mathcal{F}_t^I$ where the consumer also knows the current level of the time-varying component xt. For analytical tractability, I focus on the case $\rho = 1$ as in Tallarini (2000). This assumption, in conjunction with the Gaussian shock processes that I assume, allows for simple closed-form solutions for the value function under the two information sets. The $\rho=1$ limit of recursion (5) is:

$$v_t = \frac{\beta}{1-\theta}\log E_t\left(\exp[(1-\theta)(v_{t+1} + c_{t+1} - c_t)]\right). \tag{6}$$

3. Optimal Information Acquisition

3.1 Alternative Value Functions

3.1.1 Updating Every Period (Nash)

I denote the log continuation value when the consumer observes x_t and expects to always learn x_t by v_t^N. As one will see below, always choosing to acquire information is a Nash equilibrium in a game that the agent plays against her future selves. In this case the continuation value from equation (6) is:

$$v_t^N = \mu_v + U_v x_t, \tag{7}$$

where

$$\mu_v = \frac{\beta}{1-\beta}\left(\mu_c + \frac{1-\theta}{2}\left(\sigma_c^2 + \frac{1}{(1-\kappa\beta)^2}\sigma_x^2\right)\right)$$
$$U_v = \frac{\kappa\beta}{1-\kappa\beta}. \tag{8}$$

The term μ_v is the unconditional expectation of the scaled log continuation value. It is from the discounted present value of the long-run consumption growth μ_c and a correction for the variance of the consumption growth rate that depends on the coefficient of relative risk aversion parameter θ. The coefficient Uv gives the discounted present value of the temporary increase in log-consumption growth induced by a unit change in the mean-reverting trend component x_t.

3.1.2 Never Updating

I now turn to the agent's value function if she has no information on the current level of x and she can commit to never learning anything about x in the future. In the next section, I show how an equilibrium strategy can support this. When no information is revealed about x_t, the only variable in the agent's information set that changes over time is the current consumption level Ct. It follows that v_t is constant. I denote its value by v^I, where the superscript I reflects the relative ignorance of the consumer under this information set. From equation (6), it follows that v^I satisfies:

$$v^I = \frac{\beta}{(1-\beta)(1-\theta)}\log E\left[\exp[(1-\theta)(c_{t+1}-c_t]\right] \tag{9}$$

Unconditionally, $\Delta c_{t+1} \sim N\left(\mu_c, \sigma_c^2 + \sigma_x^2/(1-\kappa^2)\right)$. Solving for the expectation on the right-hand side of equation (9) gives:

$$v^I = \frac{\beta}{1-\beta}\left[\mu_c + \frac{1-\theta}{2}\left(\sigma_c^2 + \frac{1}{1-\kappa^2}\sigma_x^2\right)\right]. \tag{10}$$

The first term in equation (10) is the discounted present sum of future mean growth rates. The second term in the parenthesis of equation (10) is a risk adjustment that is proportional to the unconditional variance of consumption growth rates. With log utility ($\theta = 1$), this term is zero. When the coefficient of relative risk aversion is greater than 1 the risk correction lowers the continuation value.

3.1.3 Interpretation

Under the coarser information set, the unconditional and conditional variances of consumption growth rates are equal. For $\kappa \in (0, 1)$, this means that the consumer faces a higher conditional consumption volatility under the coarser information set. By itself, this will increase the perceived riskiness of the consumption path and gives the consumer an incentive to choose the finer information set.

Under the finer information set the conditional variance of trend innovations enters the consumer's continuation value with the scaling factor $1/(1-\beta\kappa)^2$. This reflects that any shock w_{t+1} to the trend growth rate is sticky. A shock w_{t+1} will increase consumption growth at $t+1+n$ by $\kappa^n w_{t+1}$. Relative to an increase in time $t+1$ consumption growth, consumption growth at $t+1+n$ is valued at β^n. The factor $\sum_{n=0}^{\infty} \beta^n \kappa^n = 1/(1-\beta\kappa)$ scales the effect of shocks to trend consumption to take account for the stickiness of the trend.

Under the coarser information set, the agent effectively finds herself living in an economy where consumption growth is a random walk with a drift. In this economy consumption growth rates are more volatile, which is reflected in the scaling factor $1/(1-\kappa^2)$ on the conditional variance of the trend growth rate of consumption. For $\theta > 1$, the consumer profits from the lower conditional variance, but suffers from the larger impact of innovations to x_{t+1} on her continuation value.

### 3.1.4 When Does Ignorance Pay Off?

It is only interesting for the consumer to opt for the coarser information set when her ex-ante continuation value is higher without information on the trend growth rate. This is the case when:

$$v^I = \mathcal{Q}(v^I) \geq \mathcal{Q}(v^N) = \mu_v + \frac{1-\theta}{2}(U_v)^2 \frac{\sigma_x^2}{1-\kappa^2} \tag{10}$$

Since v^I is constant, it is equal to its certainty equivalent (i.e. $v^I = Q(v^I)$). Here, v^N depends on the normally distributed trend growth rate, and so its certainty equivalent $Q(v^N)$ corrects for the influence of the trend growth rate through the last term on the right-hand side of the above equation. Figure 2 provides a graphical analysis of the agent's options: the shaded area in the figure gives combinations of β and κ, where v^N is higher than $Q(v^N)$. In other words, it gives parameter combinations for which an agent would prefer to commit to not learning the trend growth rate. For the high time discount factors used in the long-run risk literature (see Section 5), the figure indicates that the agent would prefer to commit to ignorance regardless of the value of the persistence parameter κ.

Figure 2. Parameter Regions where Information Lowers Utility when $\theta > 1$

The shaded area gives parameter values for which there is an expected utility loss from always learning the trend growth rate compared to the case of ignorance. The coefficient of relative risk aversion is fixed at a value of 10.
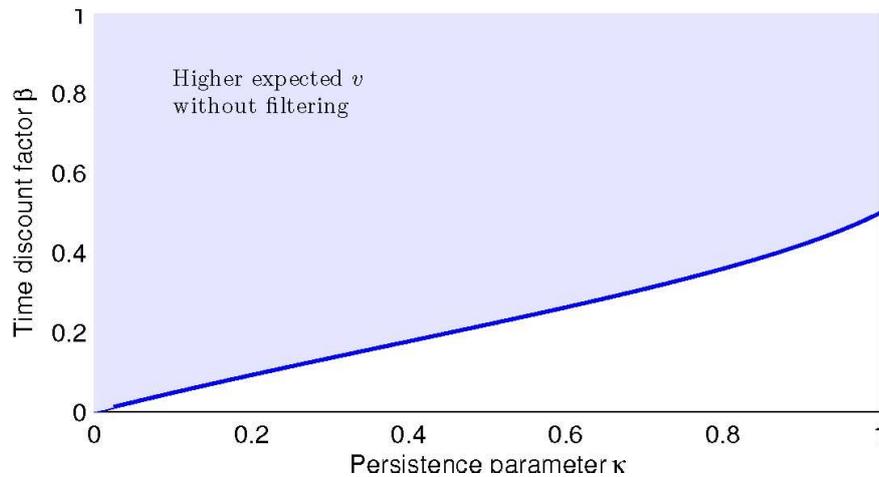
4. Implementability

For ignorance to be an equilibrium strategy, I need to show that it is feasible and individually rational.

4.1 Feasibility

Since we are in a complete markets endowment economy, one feasible investment strategy for the consumer is to invest all her wealth in a consumption claim (i.e. invest in a claim that pays dividends proportional to aggregate consumption) and every period consume its dividends. Because I am assuming that the preferences of the consumer are identical to those of the representative agent in the economy, this is the same consumption and portfolio choice she will achieve in equilibrium if she chooses to learn the trend growth rate every period.

4.2 Individual Rationality

Always Filtering is a Nash Equilibrium

Assume that the agent knows x_{t-1} and that she expects she will always include the current value of x in her information set in future periods. Her certainty equivalent if also she chooses to learn x_t this period is:

$$\mathcal{Q}(v_t \mid x_{t-1}) = \mathcal{Q}[\mu_v + U_v x_t \mid x_{t-1}]$$

$$= \mu_v + U_v \kappa x_{t-1} + \frac{1-\theta}{2} \frac{\beta^2 \kappa^2}{(1-\beta\kappa)^2} \sigma_x^2. \tag{11}$$

Here continuation value if she does not learn x_t is:

$$v_t = \frac{\beta}{1-\theta} \log E \left(\exp[(1-\theta)(v_{t+1} + c_{t+1} - c_t] \mid x_{t-1} \right)$$

$$= \mu_v + U_v \kappa x_{t-1} + \frac{1-\theta}{2} \frac{\beta \kappa^2}{(1-\kappa\beta)^2} \sigma_x^2. \tag{12}$$

The two expressions differ only in a factor βin the last term. For $\theta > 1$, the right-hand side of equation (12) is strictly lower than the certainty equivalent when the agent learns the state of the economy given by equation 11. This implies that the agent suffers a utility loss if she deviates by not updating her information on the trend growth rate in a single period.

Supporting Ignorance Strategy by Threat of a Nash

Consider the following strategy for the agent who has no information on the trend growth rate. As long as she has never learned the state of the economy in the past, she will never choose to learn it. She promises herself that, should she ever deviate from this strategy by learning the growth rate trend, she will always keep learning it in future periods. The threat is credible, since it amounts to playing a Nash strategy.

Her continuation value conditional on never updating in the future is v^I. Her expected scaled continuation value if she deviates is:

$$\mathcal{Q}[v_t^N] = \mu_v + \frac{1-\theta}{2} U_v^2 \frac{\sigma_x^2}{1-\kappa^2}.$$

As long as the preference and process parameters belong to the shaded area in Figure 2, the agent will never choose to deviate, and so never learning x_t is an equilibrium strategy.

5. Numerical Results

In this section I quantify the utility gains that an agent could achieve by committing to ignorance using two parameterizations taken from successful asset pricing models: one taken employed by Bansal and Yaron (2004) and one employed by Hansen (2007).

I measure the utility gains from ignorance by solving for the percentage change in consumption level an agent, who is forced to play the Nash strategy of always learning the trend growth rate of consumption, would need to make her equally happy ex-ante as an agent who is allowed to commit to ignorance.

For $\rho = 1$ the utility gain using the closed-form solutions provided in equations 7 and 10. In particular, I subtract from v the certainty equivalent $\mathcal{Q}(v_t^N)$.

For $\rho \neq 1$, I use a Gaussian quadrature with 300 nodes to approximate the law of motion for the trend growth rate and solve for v_t on the nodes of the quadrature. Here I compute the certainty equivalent by applying the operator Q to the values of vt on the grid using the ergodic state probabilities implied by the discretized law of motion. (See Tauchen and Hussey (1991) for a discussion of this method.)

Table 1. Utility Gains

Reported below are the estimated utility gains for an agent who is not processing the available information in the analyzed experimental economy. The first column is based on the calibration introduced by Bansal and Yaron (2004), where the parameters for the utility function are as follows: the coefficient of relative risk aversion is set to 10, EIS is set to 1.5, and the discount factor β is set to 0.998. For this case, I complete the gains with the closed solution case of EIS =1. The last column is calculated with the set of parameters specified in Hansen (2007), which has a coefficient of relative risk aversion of 2 and an EIS set to 1. All reported gains are in percentage points.

	Bansal and Yaron (2004)	Hansen (2007)
Process parameters:		
μ_c	0.0015	0.0056
σ_c	0.0078	0.0054
σ_x	0.0003	0.0005
κ	0.9720	0.9800
Utility parameters:		
β	0.998	0.998
$1/\rho$	1.5	1.0
θ	10	2
Utility Gains:		
EIS $\neq$ 1	42.986	
EIS = 1	35.021	12.573

To match asset prices, all the proposed parametrizations share a high level of

persistence (κ) for the consumption process. Such high levels of κ generate large utility gains from committing to ignore the state of the trend growth rate, because it magnifies both the effect of the information of the trend growth rate on the conditional variance of consumption growth and the larger impact of innovations to x_{t+1} on her continuation value.

Table 1 reports the quantitative results. The first column gives the gain from commiting to ignorance for the calibration used by Bansal and Yaron (2004), which has an elasticity of intertemporal substitution ($1/\rho$) of 1.5. Results are striking, with a 43% increase in lifetime consumption obtained by ignoring the trend growth rate. The last line of the table gives the same figure in the case when the elasticity of intertemporal substitution is 1. The utility gains from committing to ignore the information on the trend growth rate are still sizable at 35%. The second column of Table 1 gives results for the parametrization used by Hansen (2007). He sets the risk aversion parameter to 2 and the elasticity of substitution to 1. By itself, reducing θ from 10 to 2 cuts the utility gains from committing to ignorance by 88%, but the higher standard deviation of innovations to the trend growth rate still produces a utility gain equivalent to a 12% increase in lifetime consumption from committing to ignorance.

6. Conclusion

In this paper I have documented an unknown feature of the family of recursive preferences known as Epstein-Zin preferences that arise in long-run risk models used heavily in asset pricing. I have shown that an agent can achieve large utility gains after committing to ignoring information on the state of the trend growth rate.

The feature documented is surprising, because the preference parameters used are known to produce a preference for early resolution of uncertainty. This study's model of the agent's decision problem as a repeated game against her future selves shows that such a commitment to ignorance can be implemented and supported as an equilibrium strategy.

References

Bansal, R. and A. Yaron, (2004), "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59(4), 1481-1509.

Bansal, R. and I. Shaliastovich, (2012), "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," *Review of Financial Studies*, 26(1), 1-33.

Bansal, R., D. Kiku, and A. Yaron, (2012), "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices," *Critical Finance Review*, 1(1), 183-221.

Campbell, J. Y. and T. Vuolteenaho, (2004), "Bad Beta, Good Beta," *American Economic Review*, 94(5), 1249-1275.

Cederburg, S. and S. Hore, (2008), "Long-Run Risk in the Cross Section," *Working*

Paper, University of Iowa.

Colacito, R. and M. M. Croce, (2011), "Risks for the Long Run and the Real Exchange Rate," *Journal of Political Economy*, 119(12), 153-181.

Constantinides, G. M. and A. Ghosh, (2011), "Asset Pricing Tests with Long-Run Risks in Consumption Growth," *Review of Asset Pricing Studies*, 1(1), 96-136.

Croce, M. M., (2014), "Long-Run Productivity Risk: A New Hope for Production-based Asset Pricing?" *Journal of Monetary Economics*, 66, 13-31.

Doh, T., (2013), "Long-Run Risks in the Term Structure of Interest Rates: Estimation," *Journal of Applied Econometrics*, 28(3), 478-497.

Epstein, L. G. and S. Zin, (1989), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption Growth and Asset Returns I: A Theoretical Framework," *Econometrica*, 57(4), 937-969.

Hansen, L. P., (2007), "Beliefs, Doubts and Learning: Valuing Macroeconomic Risk," *American Economic Review*, 97(2), 1-30.

Hansen, L. P., J. Heaton, and N. Li, (2008), "Consumption Strikes Back?: Measuring Long-Run Risk," *Journal of Political Economy*, 116(2), 260-302.

Hasseltoft, H., (2012), "Stocks, Bonds, and Long-Run Consumption Risks," *Journal of Financial and Quantitative Analysis*, 47(2), 309-332.

Kaltenbrunner, G. and L. A. Lochstoer, (2010), "Long-Run Risk through Consumption Smoothing," *Review of Financial Studies*, 23(8), 3190-3224.

Kocherlakota, N. R, (1990), "Disentangling the Coefficient of Relative Risk Aversion from the Elasticity of Intertemporal Substitution: An Irrelevance Result," *Journal of Finance*, 45(1), 175-190.

Kreps, D. M. and E. L. Porteus, (1978), "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46(1), 185-200.

Lettau, M., S. Ludvigson, and J. Wachter, (2008), "The Declining Equity Premium: What Role Does Macroeconomic Risk Play?" *Review of Financial Studies*, 21(4), 1653-1687.

Piazzesi, M. and M. Schneider, (2007), "Equilibrium Yield Curves," in *NBER Macroeconomics Annual 2006*, Volume 21, National Bureau of Economic Research, Inc., 389-472.

Rangvid, J., M. Schmeling, and A. Schrimpf, (2009), "Global Asset Pricing: Is There a Role for Long-run Consumption Risk?" *Working Paper*, Copenhagen Business School.

Skiadas, C., (1998), "Recursive Utility and Preferences for Information," *Economic Theory*, 12(2), 293-312.

Tallarini, T. D., (2000), "Risk-sensitive Real Business Cycles," *Journal of Monetary Economics*, 45(3), 507-532.

Tauchen, G. and R. Hussey, (1991), "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, 59(2), 371-396.

Weil, P., (1989), "The Equity Premium Puzzle and the Risk Free Rate Puzzle," *Journal of Monetary Economics*, 24(3), 401-421.

Wu, S., (2008), "Long-Run Consumption Risk and the Real Yield Curve," *Working Paper*, University of Kansas 2008.