Macroeconomic Evidence for Prudence

AJ A. Bostian and Christoph Heinzel*

April 20, 2010

Abstract

We investigate the presence of a "prudent" savings attitude at the macroeconomic level with a Real-Business-Cycle model. We extend the consumer side of the model with an "expo-power" utility function, a two-parameter specification that can generate an array of coefficients of risk aversion and prudence. We estimate the model on US macroeconomic data using the particle-filter maximum-likelihood method, and find evidence for weak decreasing relative prudence. We examine the "social" discount rate associated with our estimates. Although the discount effect is small in percentage-point terms, the value differentials over long horizons or at large scales can potentially be significant.

*Bostian: Australian School of Business, University of New South Wales. Heinzel: Toulouse School of Economics (LERNA). This manuscript was prepared for the 2009 Southern Economic Association annual meeting. It is not for citation purposes, but comments are certainly welcome. We would like to thank participants at the 2009 SEA meeting for helpful comments. This research was partially supported by the Commonwealth Environment Resources Fund (Australia), the Australian Research Council, and the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) Grant Agreement no. 230589.

1 Introduction

As formulated by Kimball (1990), "prudence" describes the intensity of the precautionary-savings motive due to risk. In an expected-utility framework, it can be captured by a positive third derivative of the utility function (*cf.* Leland 1968, Sandmo 1970, Drèze and Modigliani 1972). The effects of prudence have been investigated in a variety of theoretical contexts, such as the likely concavity of the consumption function (Carroll and Kimball 1996, 2001), the effect of background risk on risk-taking behavior (Gollier and Pratt 1996), asset demand arising from return predictability (Gollier 2004, 2008), and the term structure of interest rates (Gollier 2010). Importantly, under fairly general conditions, decreasing relative risk aversion – which arises when relative prudence is sufficiently stronger than relative risk aversion – causes the so-called "social discount rate" to decline over time (Gollier 2002a,b), leading to optimal savings paths that give relatively higher weight to long-term outcomes than would arise otherwise.

Although a multitude of theoretical results regarding prudence have been established, there have been few explicit attempts to empirically assess the curvature of marginal utility function that should be observed if prudence is present.[1] An early attempt by Dynan (1993) estimated a second-order Taylor approximation to the consumption Euler equation, and found a negligible coefficient of relative prudence.[2] Using varied methodologies and datasets, other studies have generally found higher, though mutually inconsistent values (*cf.*, Kuehlwein 1991, Merrigan and Normandin 1996, Eisenhauer 2000, Ventura and Eisenhauer 2006). Parker and Preston (2005) use a more detailed decomposition of the consumption Euler equation to study the relative importance of four proximate causes for movements in average consumption growth. They find movements in the precautionary term to impact on the variance of predictable movements in consumption in the same order of magnitude as movements in the real interest rate.

Carroll (2001) sharply criticized these reduced-form approaches for their inescapable identifica-

[1]Of course, third derivatives of the utility function are implicit in any structural model involving expected utility, so estimates of such models necessarily contain prudential implications. However, these third-derivative effects and their consumption-smoothing implications are rarely studied in their own right.

[2]Dynan used a power-utility specification in the Euler equation, which implies a coefficient of relative prudence equal to $1 + \rho$, where ρ is the coefficient of relative risk aversion. She found $1 + \rho \in [0.14, 0.166]$, which makes ρ inconsistent with macroeconomic evidence (*cf.* Meyer and Meyer 2005). Lee and Sawada (2007) re-estimated this model with liquidity constraints and found the still-incompatible values $1 + \rho \in [0.838, 1.094]$. These results illustrate the potential susceptibility of prudential estimates to specification and approximation error.

tion issues.[3] And indeed, more recently, precautionary-savings effects have instead been estimated via structural approaches.

Taking up the structural approach, we introduce prudential flexibility into a Real Business Cycle model using the "expo-power" utility function of Saha (1993), a two-parameter specification that can capture a wide array of risk attitudes. Depending on the parametrization, this specification can capture increasing, decreasing, and constant risk aversion and prudence.

Using the particle-filter maximum-likelihood method described in Fernández-Villaverde and Rubio-Ramírez (2005), we estimate our extended model on US macroeconomic data. We find evidence of decreasing relative risk aversion and decreasing relative prudence. Although the "decreasing" feature appears to be weak (in terms of the magnitudes of the coefficients), we can reject several alternative hypotheses involving constant relative risk aversion and prudence. Importantly, although the implicit social discount rate is at most a few tenths of a percentage point different from that implied by constant relative risk aversion, this small difference can potentially generate a large difference in expected value for economic activities with a long horizon.

The remainder of this paper is structured as follows. Section 2 describes our modified Real Business Cycle model, and how the parameters are likely to be identified. Section 3 describes our data and estimates, and provides some side-by-side graphs of the difference in trajectories relative to a log-utility counterfactual. Section 4 describes the how our results can be interpreted in the light of the social discount rate.

2 A Modified Real Business Cycle Model

The central feature of our model is Saha's expo-power utility function:

$$u(x) = \frac{1}{\alpha}\left[1 - \exp\left(-\frac{\alpha}{1-\rho}x^{1-\rho}\right)\right]$$

[3]Related discussions can be found in Gourinchas and Parker (2002a) and Ludvigson and Paxson (2001).

It is fairly easy to show that this specification nests exponential utility ($\rho = 0$) and power utility ($\alpha \to 0$). The coefficients of absolute and relative risk aversion are

$$ARA(x) = \frac{\alpha}{x^{\rho}} + \frac{\rho}{x} \quad RRA(x) = \frac{\alpha}{x^{\rho-1}} + \rho \tag{1}$$

and the coefficients of absolute and relative prudence are

$$AP(x) = ARA(x) + \frac{ARA'(x)}{ARA(x)} \quad RP(x) = RRA(x) + \frac{ARA'(x)}{ARA(x)}x \tag{2}$$

The presence of increasing or decreasing relative risk aversion is controlled by the combination of α and ρ. For example, $\alpha > 0$ and $\rho > 1$ generates decreasing relative risk aversion, with $\lim_{x\to\infty} RRA(x) = \rho$. And, $\alpha > 0$ and $\rho < 1$ generates increasing relative risk aversion, with $\lim_{x\to\infty} RRA(x) = \infty$.

These limiting features of expo-power utility under various parametrizations can suggest the type of estimates we might expect in macroeconomic data. When utility is additively-separable over time, it is well-known that the elasticity of intertemporal substitution ϵ is given by

$$\frac{1}{\epsilon_t} = RRA(x_t)$$

If there is no intertemporal elasticity to the consumption of x, then we ought to observe increasing relative risk aversion, because a positive α and a small ρ will cause ϵ to tend towards 0 as consumption increases over time. However, if there is some elasticity, then we ought to observe decreasing relative risk aversion. In this case, ϵ will depend upon the level of x, but will eventually converge to $1/\rho$ as x increases.[4]

In the Real Business Cycle model, a consumer's utility is a function of both consumption c

[4] Our estimates indicate that decreasing relative risk aversion (and prudence) is the best fit on US macroeconomic data. However, other estimates of the expo-power parameters have suggested the presence of increasing relative risk aversion. For example, Holt and Laury (2002) find a best fit of $\alpha = 0.029$ and $\rho = 0.269$ on data from lottery-choice experiments. This difference underscores the potential danger of too strongly associating a "coefficient of risk aversion" with a "risk attitude" in an intertemporal model. In the intertemporal setting, the coefficient of relative risk aversion is equivalent to the elasticity of intertemporal substitution. It is consumption-adjustment patterns that are being estimated, which may differ from a risk attitude. For example, Kimball and Weil (2009) analyze prudence in an intertemporal setting using a model that disentangles the two concepts, and find that there are potentially two separate effects that can operate simultaneously to determine the degree of prudence.

and leisure h. This concept is usually operationalized by putting utility over an aggregate index $x = \phi(c, h)$. When this occurs, we are usually most interested in the elasticity of intertemporal substitution with respect to consumption, given by

$$\frac{1}{\epsilon_t^c} = ARA\left(\phi\left(c_t, h_t\right)\right)\phi_c\left(c_t, h_t\right)c_t + \frac{\phi_{cc}\left(c_t, h_t\right)c_t}{\phi_c\left(c_t, h_t\right)} \tag{3}$$

Following Aruoba et al. (2006), we use a Cobb-Douglas aggregator function $\phi(c, h) = c^\theta h^{1-\theta}$, so that

$$\phi_c(c, h) = \left(\frac{h}{c}\right)^{1-\theta}$$

and

$$\frac{\phi_{cc}(c, h)\, c}{\phi_c(c, h)} = 1 - \theta$$

Equation (3) provides some insight into how the expo-power parameters can be identified. As the economy receives stochastic shocks, consumers will probably change their intertemporal consumption ratios in response (unless $\epsilon_t^c = 0$). The correlation between this response and the level of ϕ determines the values of α and ρ. Thus, conditional on θ, the parameters α and ρ can be identified by variation in intertemporal substitution elasticities.

The remainder of our model is fairly standard. Briefly, it consists three parts.

1. Market clearing. At time t, the economy consists of a population P_t of identical customers, and a representative firm. Three markets operate: one each for capital K_t, labor L_t, and a consumption-investment numeraire Y_t. Denoting firm quantities by an uppercase letter and consumer quantities by the corresponding lowercase letter, market clearing at time t is defined by prices r_t (for capital) and w_t (for labor) such that

$$\begin{aligned} K_t &= P_t k_t \quad \text{(Capital Market Clears)} \\ L_t &= P_t l_t \quad \text{(Labor Market Clears)} \\ Y_t &= P_t y_t \quad \text{(Numeraire Clears)} \end{aligned}$$

2. Firm. At time t, the firm possesses a production technology F_t. It rents capital K_t at rate r_t

and hires labor L_t at rate w_t, and uses these inputs to produce Y_t units of the the numeraire good. Its profit-maximization problem is

$$\max_{\{K_t,L_t\}} E_0\left[\sum_{t=0}^{\infty} Y_t - r_t K_t - w_t L_t\right] \quad s.t.\ Y_t \le F_t\left(K_t, L_t\right)$$

As is usual in the Real Business Cycle formulation, we take the production technology F_t to be the Cobb-Douglas form $F_t\left(K_t, L_t\right) = A_t K_t^{\gamma} L_t^{1-\gamma}$. Uncertainty in the model arises due to fluctuations in the total factor productivity A_t. These fluctuations have the lognormal form $\ln\left(A_{t+1}\right) = \psi \ln\left(A_t\right) + \varepsilon_t^A$, where $\varepsilon_t^A \sim$ i.i.d. $N\left(0, \sigma_A^2\right)$.

3. Consumer. Each consumer enters the period with a personal capital stock k_t and a time endowment (normalized to 1 unit) that can be allocated between labor l_t and leisure h_t. The consumer rents capital at rate r_t, and receives wage w_t for labor. The consumer allocates this income to purchasing y_t units of the numeraire good, which it allocates between consumption c_t and capital investment i_t. Capital depreciates at a rate δ. The consumer's utility-maximization problem is[5]

$$\max_{\{c_t,l_t,k_{t+1}\}} E_o\left[\sum_{t=0}^{T} \beta^t u\left(\phi\left(c_t, h_t\right)\right)\right] \quad s.t. \quad \begin{cases} y_t = c_t + i_t \le r_t k_t + w_t l_t \\ i_t = \frac{P_{t+1}}{P_t} k_{t+1} - (1-\delta) k_t \\ l_t + h_t = 1 \end{cases}$$

At time t, there are three three first-order conditions on the consumer side

$$E_t\left[\beta \frac{u'\left(\phi\left(c_{t+1}, h_{t+1}\right)\right) \phi_c\left(c_{t+1}, h_{t+1}\right)}{u'\left(\phi\left(c_t, h_t\right)\right) \phi_c\left(c_t, h_t\right)} \cdot \frac{P_t}{P_{t+1}} \cdot \left(1 - \delta + r_{t+1}\right)\right] = 1 \tag{4}$$

$$\frac{\phi_h\left(c_t, h_t\right)}{\phi_c\left(c_t, h_t\right)} = w_t \tag{5}$$

$$r_t k_t + w_t l_t = c_t + i_t \tag{6}$$

[5] Our formulation of the investment identity takes into account future population growth. *From the perspective of time t,* a consumer's share of the time-$(t+1)$ capital stock is K_{t+1}/P_t. However, because we have defined $k_t = K_t/P_t\, \forall t$ (through the market-clearing condition on capital), clearly $K_{t+1}/P_t = \left(P_{t+1}/P_t\right) k_{t+1} \neq k_{t+1}$ (unless $P_{t+1} = P_t$). Because absolute levels are important in the analysis of prudent risk attitudes, we retain the possibility of (deterministic) population growth to get a more precise estimate of the individual capital stock.

There are also two first-order conditions on the firm side:

$$r_t = \gamma A_t \left(\frac{L_t}{K_t}\right)^{1-\gamma} \tag{7}$$

$$w_t = (1-\gamma)\, A_t \left(\frac{K_t}{L_t}\right)^{\gamma} \tag{8}$$

These equations provide insight into how the remaining parameters can be identified. The value of γ can be identified by changes in the capital/labor ratio over time (equations (7) and (8)), which feed into the consumer's total expenditure (equation (6)). Similarly, the values of ψ and σ_A can be identified by fluctuations in output which uniformly raise incomes in a manner uncorrelated with the capital/labor ratio. Finally, conditional on γ, the value of θ can be found by variation in the consumption/leisure ratio for various wage levels (equation (5)).

3 Estimates and Trajectories

Estimating the model from the previous section involves repetitions of a two-stage algorithm:

1. Solve the model's policy function for a given parametrization.

2. Given a solution to the policy function, approximate the likelihood function.

Fernández-Villaverde et al. (2006) note that the methodology in each stage has important implications for the accuracy of the overall solution. In particular, a second-order error in the policy-function solution generates a first-order error in the likelihood function. Errors in the policy function have also been shown to provide questionable welfare estimates (*e.g.*, the "spurious welfare reversals" found by Kim and Kim 2003).

Computational methodology is thus an important aspect of this paper, both from the perspective of obtaining more precise estimates, and from the ability to potentially undertake policy analysis with the model. We use the bootstrap particle filter described in Fernández-Villaverde and Rubio-Ramírez to approximate the likelihood function, and the Chebyshev-polynomial approximation described in Aruoba et al. to approximate the policy functions. The respective authors show that these are two of the most accurate methods in their respective stages. A more detailed description of our method can be found in Appendix A.

It is important to note that the results presented here are from a scaled-down version of the computational algorithm. Our method relies heavily on nested simulations, and the computational requirements grow exponentially as additional simulations are added in each level of the nest. Hence, our results are probably affected by small-sample simulation error. (This is the primary reason that we do not present standard errors, because the simulated likelihood function is not well-behaved in the neighborhood of the maximum we found.) However, the method can be scaled up through parallelization to a level that ameliorates this particular error.

Also, due to the ill-behaved likelihood function, we use a stochastic-search method to maximize the likelihood function instead of a derivative-based method. As a result, our reported values may not represent an exact maximum. We began the search at the parameter values estimated by Fernández-Villaverde and Rubio-Ramírez, who use data and methodology similar to ours.

Our dataset consists of seasonalized US quarterly data on consumption and investment, monthly data of labor hours, and yearly population data.[6] The resulting time coverage is 1964:Q1 to 2009:Q3. The monthly population data are first used to estimate a monthly population growth rate g using the simple growth model

$$\ln\left(Pop_t\right) = g + \ln\left(Pop_{t-1}\right) + \varepsilon_t^P$$

A time-series of population data for use in the model is then imputed using the recursion $P_t = e^{3g}P_{t-1}$. This imputed series is used to convert aggregates into per-capita values. [7]

Table 1 presents parameter estimates for the fitted model. The values $\alpha = 0.082$ and $\rho = 1.699$ suggest decreasing relative risk aversion and decreasing relative prudence. Relative to other

[6]The dataset identifiers are POP (population, monthly), PECC96 (consumption, quarterly), GPIDC96 (gross private domestic investment, quarterly), and AWHNONAG (weekly hours, monthly) from the Federal Reserve Bank of St. Louis FRED database. We translate POP and AWHNONAG into quarterly values by averaging the values for the 3 months in each quarter.

[7]We use this smoothed population series instead of the actual one because there are small fluctuations in population. Instead of incorporating population uncertainty into the model (with probably little effect), we simply use the growth rate g in each time period. The smoothed population series is always within 1% of the actual population series.

	US (quarterly)	Baseline Counterfactual
θ	0.472	(same)
β	0.997	(same)
α	0.082	0
ρ	1.699	1
γ	0.297	(same)
δ	0.008	(same)
ψ	0.974	(same)
σ_A	0.0003	(same)
ll	-4986.3	-5102.3

Table 1:

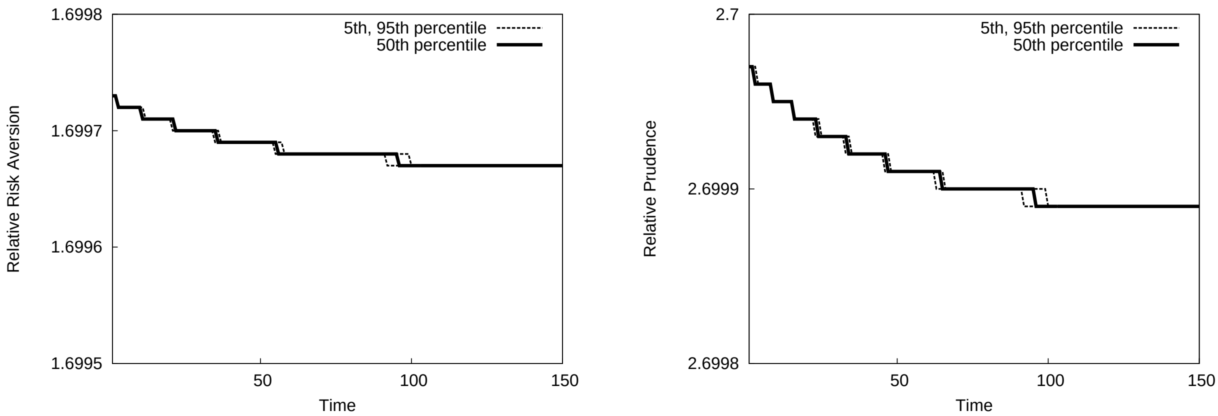

Figure 1:

estimates of the Real Business Cycle model, these estimates show a slightly lower capital share of output, and a lower standard deviation of productivity shocks. The alternative specification of log utility can be easily rejected with a likelihood-ratio test. And, in fact, many other power-utility specifications can also be rejected.

Referring to equation (1), the "decreasing" feature is primarily controlled by the value of α, which appears to be relatively small. In addition, the effect diminishes as consumption increases, and consumption is already relatively high in magnitude (thousands of dollars, with $\theta = 0.472$). Figure 1 plots the implicit change in relative risk aversion over time. Both relative risk aversion and relative prudence do fall over time, but the effect appears to be small.

Finally, some summary trajectories from this model are presented in Figure 2. The difference between these series and a log-utility baseline are provided for comparison. As one might expect, the effect of the increased prudence can be seen in the higher capital accumulation rate and lower

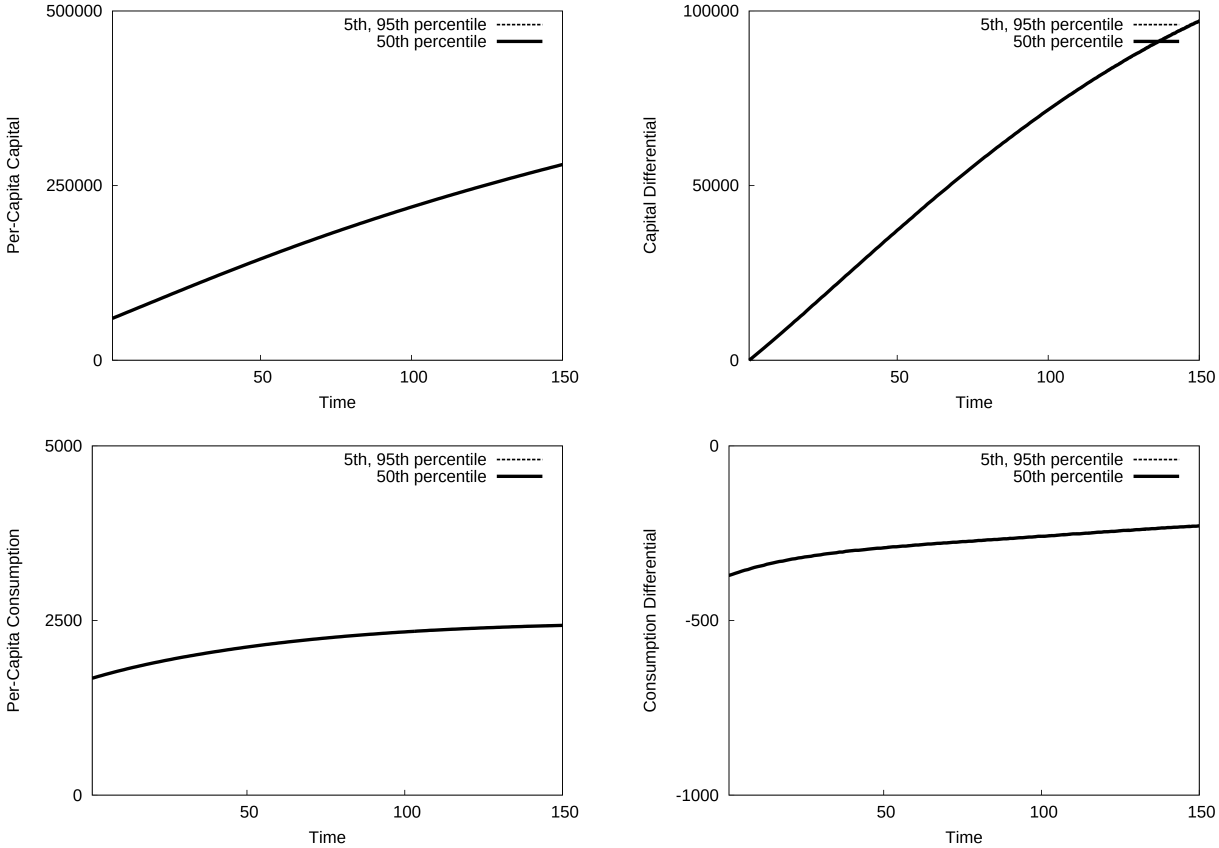

5th, 95th percentile
50th percentile
Per-Capita Capital
500000
250000
0
50
100
150
Time
Capital Differential
100000
50000
0
Per-Capita Consumption
5000
2500
Consumption Differential
-500
-1000

Figure 2:

consumption rate.

4 Policy Implications

It is difficult to assess the importance of declining relative prudence in our model by just looking at the magnitude of the coefficient. After all, utility is over a unitless aggregate index ϕ. Gollier (2002a) provides a means for us to assess the policy importance of our estimates via the social discount rate implied by our model. The social discount rate is related to the stochastic discount factor in equation (4):

$$E_t\left[\beta\frac{u'\left(\phi\left(c_{t+1},h_{t+1}\right)\right)\phi_c\left(c_{t+1},h_{t+1}\right)}{u'\left(\phi\left(c_t,h_t\right)\right)\phi_c\left(c_t,h_t\right)}\right]=\frac{1}{1+d_t}$$

Gollier shows that in the face of an uncertain future consumption stream prudence causes d_t to be smaller that under certainty, and that d_t will decrease (increase) with the time horizon if relative risk aversion is decreasing (increasing). By implication, future outcomes receive more (less) weight than they would otherwise have. This can have important implications for policy questions that involve economy-wide, long-horizon events (*e.g.*, climate change or retirement), because it is difficult to derive a discount rate for such outcomes. The social discount rate indicates that the appropriate discounting scheme is endogenous to the model, and is dependent upon the likely consumption path and the prevailing risk attitudes.

Figure 3 plots the implicit social discount rate from our model with an alternative log-utility specification. The increased prudence in our model does indeed generate a declining discount rate, and the discount rate is higher than under the alternate specification. The differential is about 0.5 percentage points per quarter at the beginning of the sample, and declines to about 0.01 percentage points by the end.

The effects of decreasing relative prudence appear relatively small in percentage-point terms. However, consider that a differential of 0.5% in the growth of $1 yields about a $1.75 difference in value after 100 quarters, and a differential of 0.01% yields about a $0.02 difference. When evaluating policies that involve time horizons such as these or longer, and that involve billions or

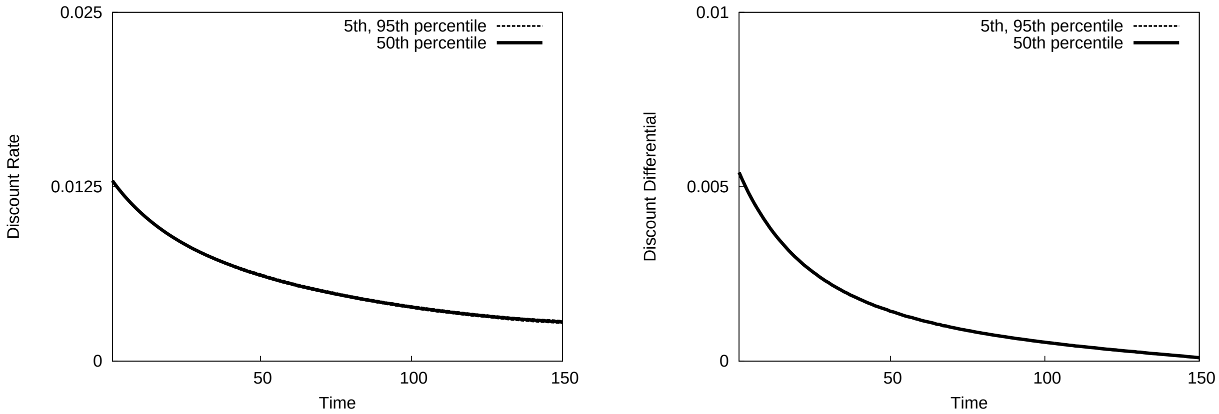

Figure 3:

trillions of dollars, such a valuation differential can translate into a very large amount. In such a scenario, care is needed in choosing a specification for d_t, and our results suggest that a non-trivial formulation of prudence can play a role in this choice.

References

S. Boraøgan Aruoba, Jesús Fernández-Villaverde, and Juan F. Rubio-Ramírez. Comparing solution methods for dynamic equilibrium economies. *Journal of Economic Dynamics and Control*, 30 (12):2477–2508, 2006.

Christopher D. Carroll. Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation). *Advances in Macroeconomics*, 1(1), 2001. Available at: http://www.bepress.com/bejm/advances/vol1/iss1/art6/.

Christopher D. Carroll and Miles S. Kimball. On the Concavity of the Consumption Function. *Econometrica*, 64(4):981–992, 1996.

Christopher D. Carroll and Miles S. Kimball. Liquidity Constraints and Precautionary Saving. 2001. NBER Working Paper 8496.

Jacques H. Drèze and Franco Modigliani. Consumption Decisions under Uncertainty. *Journal of Economic Theory*, 5(3):308–335, 1972.

Karen E. Dynan. How Prudent Are Consumers? *Journal of Political Economy*, 101(6):1104–1113, 1993.

Joseph G. Eisenhauer. Estimating Prudence. *Eastern Economic Journal*, 24(6):379–392, 2000.

Jesús Fernández-Villaverde and Juan F. Rubio-Ramírez. Estimating Dynamic Equilibrium Economies: Linear Versus Nonlinear Likelihood. *Journal of Applied Econometrics*, 20(7):891–910, 2005.

Jesús Fernández-Villaverde, Juan F. Rubio-Ramírez, and Manuel F. Santos. Convergence Properties of the Likelihood of Computed Dynamic Models. *Econometrica*, 74(1):93–119, 2006.

Christian Gollier. Discounting an uncertain future. *Journal of Public Economics*, 85(2):149–166, 2002a.

Christian Gollier. Time Horizon and the Discount Rate. *Journal of Economic Theory*, 107(2): 463–473, 2002b.

Christian Gollier. Optimal Portfolio Risk with First-Order and Second-Order Predictability. *Contributions to Theoretical Economics*, 4(1), 2004. Available at: http://www.bepress.com/bejte/contributions/vol4/iss1/art4.

Christian Gollier. Understanding saving and portfolio choices with predictable changes in asset returns. *Journal of Mathematical Economics*, 44(5-6):445–458, 2008.

Christian Gollier. *Pricing the future: The economics of discounting and sustainable development.* 2010. Unpublished manuscript, to appear with Princeton University Press, Princeton, NJ.

Christian Gollier and John W. Pratt. Risk Vulnerability and the Tempering Effect of Background Risk. *Econometrica*, 64(5):1109–1123, 1996.

Pierre-Olivier Gourinchas and Jonathan A. Parker. The Empirical Importance of Precautionary Saving. *American Economic Review*, 91(2):406–412, 2002a.

Charles A. Holt and Susan K. Laury. Risk Aversion and Incentive Effects. *American Economic Review*, 92(5):1644–1655, 2002.

Jinill Kim and Sanghyun Henry Kim. Spurious welfare reversals in international business cycle models. *Journal of International Economics*, 60(2):471–500, 2003.

Miles Kimball and Philippe Weil. Precautionary Saving and Consumption Smoothing across Time and Possibilities. *Journal of Money, Credit and Banking*, 41(2-3):245–284, 2009.

Miles S. Kimball. Precautionary Saving in the Small and in the Large. *Econometrica*, 58(1):53–73, 1990.

Michael Kuehlwein. A test for the presence of precautionary saving. *Economics Letters*, 37(4): 471–475, 1991.

Jeong-Joon Lee and Yasuyuki Sawada. The degree of precautionary saving: A reexamination. *Economics Letters*, 96(2):196–201, 2007.

Hayne E. Leland. Saving and Uncertainty: The Precautionary Demand for Saving. *Quarterly Journal of Economics*, 82(3):465–473, 1968.

Sydney Ludvigson and Christina H. Paxson. Approximation Bias in Linearized Euler Equations. *Review of Economics and Statistics*, 83(2):242–256, 2001.

Philip Merrigan and Michel Normandin. Precautionary Saving Motives: An Assessment from UK Time Series of Cross-Sections. *Economic Journal*, 106(438):1193–1208, 1996.

Donald J. Meyer and Jack Meyer. Relative Risk Aversion: What Do We Know? *Journal of Risk and Uncertainty*, 31(3):243–262, 2005.

Jonathan A. Parker and Bruce Preston. Precautionary Saving and Consumption Fluctuations. *American Economic Review*, 95(4):1119–1143, 2005.

Atanu Saha. Expo-Power Utility: A 'Flexible' Form for Absolute and Relative Risk Aversion. *American Journal of Agricultural Economics*, 75(4):905–913, 1993.

Agnar Sandmo. The Effect of Uncertainty on Saving Decisions. *Review of Economic Studies*, 30 (3):353–360, 1970.

George Tauchen. Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions. *Economics Letters*, 20(2):177–181, 1986.

Luigi Ventura and Joseph G. Eisenhauer. Prudence and Precautionary Saving. *Journal of Economics and Finance*, 30(2):155–168, 2006.

A Computational Methods

Estimating our model involves two algorithms: one to approximate the policy functions, and one to approximate the likelihood function.

To approximate the policy functions, we use the Chebyshev-polynomial method described in Aruoba et al.. Briefly, this method involves computing a functional approximation to the labor policy function $l(A, k, P)$. Note that once the state variables A, K, and P are available, as well as the labor value l, then r and w can be immediately computed using equations (7) and (8). Equation

(5) can then be used to recover c (using a nonlinear root-finding algorithm), and equation (6) can be used to recover i. The form of our functional approximation is

$$l(A,k,P) \simeq \sum_{m_A=0}^{M_A} \sum_{m_k=0}^{M_k} \sum_{m_P=0}^{M_P} \omega_{m_A,m_k,m_P} \varphi_{m_A}(A) \varphi_{m_k}(k) \varphi_{m_P}(P)$$

where φ_m denotes the Chebyshev polynomial of degree m, and $\{\omega_{m_A,m_k,m_P}\}_{m_A=0,m_k=0,m_P=0}^{M_A,M_k,M_P}$ are weights to be estimated.[8]

To find these weights, the free first-order condition 4 is used as an estimating equation. Trial points

$$\{A_i, k_i, P_i\}_{i=1}^{(M_A+1)(M_k+1)(M_P=1)}$$

are generated from the nodes of the (M_A+1)-, (M_k+1)- and (M_P+1)-degree polynomials, by taking all possible combinations of the node points. Each trial point can be individually applied to equation 4, for a total of $(M_A+1)(M_k+1)(M_P+1)$ equations. This exactly-identified system can be solved for the weights by a nonlinear root-finding algorithm.[9]

To approximate the likelihood function, we use the bootstrap particle-filter method described in Fernández-Villaverde and Rubio-Ramírez. For a particular parametrization of the model, this method approximates the likelihood using a "swarm" of "particles." In our case, a particle $\left(\varepsilon^{A,t,n}, S_0^n\right)$ is composed of a history of shocks $\varepsilon^{A,t,n} = \left\{\varepsilon_0^{A,n}, \ldots, \varepsilon_t^{A,n}\right\}$ and an initial state $S_0^n = \{A_0^n, k_0^n, P_0^n\}$. Given a swarm (*i.e.*, draws) of N such particles from time $t-1$, the likelihood contribution for time t can be computed as follows:

1. Draw a "transition swarm" $\left\{\varepsilon^{A,t|t-1,n}, S_0^n\right\}_{n=1}^N$ by sampling from

$$p\left(\varepsilon^{A,t}|y^t, S_0\right) = f_A\left(\varepsilon_t^A\right) p\left(\varepsilon^{t-1}|c^{t-1}, l^{t-1}, i^{t-1}, S_0\right)$$

[8]Because the Chebyshev polynomials are defined over $[-1,1]$, A, k, and P must be first mapped into this domain. We choose a linear mapping such that the bounds are unlikely to be reached in our dataset: $A \in [0.75, 1.25]$, $k \in \left[5000, 1\times10^6\right]$, $P = \left[100\times10^6, 400\times10^6\right]$.

[9]Because equation 4 contains an expectation, we also construct an integral approximation using an integration-by-simulation strategy with 100 simulations of ε_{t+1}^A. For stability, we preserve the underlying standard normal values across iterations.

The time-$(t-1)$ swarm $\left\{\varepsilon^{A,t-1,n}, S_0^n\right\}_{n=1}^{N}$ is already a sample from $p\left(\varepsilon^{t-1}|c^{t-1}, l^{t-1}, i^{t-1}, S_0\right)$. Thus, the transition swarm can be drawn by sampling N particles from the entering swarm, and adding to each one a draw from f_A.

2. Assign weights $\{q_t^n\}_{n=1}^N$ to each particle in the transition swarm, where

$$q_t^n = \frac{p\left(c_t^n, l_t^n, i_t^n|\varepsilon^{A,t-1,n}, S_0^n\right)}{\sum_{m=1}^N p\left(c_t^m, l_t^m, i_t^m|\varepsilon^{A,t-1,m} S_0^m\right)}$$

In our case, the density $p\left(c_t^n, l_t^n, i_t^n|\varepsilon^{A,t-1,n}, S_0^n\right)$ is simply the likelihood of the "measurement errors" corresponding to that particular particle:

$$p\left(c_t^n, l_t^n, i_t^n|\varepsilon^{A,t-1,n}, S_0^n\right) = f_c\left(\varepsilon_t^{c,n}\right) f_l\left(\varepsilon_t^{l,n}\right) f_i\left(\varepsilon_t^{i,n}\right)$$

The measurement errors denote the discrepancy between the model prediction and the empirical observation. Letting $\bar{c}_t$, $\bar{l}_t$, and $\bar{\imath}_t - \delta k_t$ denote the empirical analogs of the model outputs,[10] the distribution of these errors is

- $\varepsilon_t^c = c_t - \bar{c}_t \sim$ i.i.d $N\left(0, \sigma_c^2\right)$
- $\varepsilon_t^l = l_t - \bar{l}_t \sim$ i.i.d. $N\left(0, \sigma_l^2\right)$
- $\varepsilon_t^i = i_t - \left(\bar{\imath}_t - \delta k_t\right) \sim$ i.i.d. $N\left(0, \sigma_i^2\right)$

3. The time-t likelihood contribution is

$$p\left(c_t, l_t, i_t|\varepsilon^{A,t-1}, S_0\right) \simeq \frac{1}{N}\sum_{n=1}^{N} p\left(c_t^n, l_t^n, i_t^n|\varepsilon^{A,t-1,n}, S_0^n\right)$$

4. Draw a new swarm $\left\{\varepsilon^{A,t,n}, S_0^n\right\}_{n=1}^{N}$ by sampling from the transition swarm according to the probability density defined by $\{q_t^n\}_{n=1}^N$. This new swarm is a draw from $p\left(\varepsilon^t|c^t, l^t, i^t, S_0\right)$.

Because the state variables A and k are unobservable, initializing the swarm is an issue. Note that with P_0, A_0, and a solution to the policy functions in-hand, it is possible to recover k_t using

[10] Because our investment data $\bar{\imath}_t$ reflect only gross investments, the appropriate comparison is between $\bar{\imath}_t$ and $i_t + \delta k_t$.

the free first-order condition 4 (using a nonlinear root-finding algorithm). Hence, we draw A_0 using the finite-state approximation of an autoregressive process described in Tauchen (1986). To initialize the swarm, we sample N times from a 15-state distribution, and obtain the associated k_0 for each. Even though the swarm will initially be populated with several identical (A_0, k_0) pairs, each will receive a different set of simulated shocks, so each particle will convey unique stochastic information after the initial period.

The results presented here make use of the polynomials up to degree $M_A = 5$, $M_k = 9$, and $M_P = 9$, for a total of $(M_A + 1)(M_k + 1)(M_P + 1) = 600$ weights. We also use $N = 100$ initial swarm draws. The literature mentioned previously suggests that N should be about 10,000 to avoid simulation error. Our results here are probably based upon too many samples near the mode of the swarm distribution, and too few samples from the tails. Scaling up is not of major concern, but it will likely require parallelization of the algorithm.