Prudential Uncertainty Causes Time-Varying Risk Premiums

Andrew Y. Chen*
The Ohio State University

This version: November 16, 2012†

Abstract

Time-varying risk premiums are a natural consequence of prudent savings behavior. Prudence prescribes a countercyclical marginal propensity to consume which leads to countercyclical consumption volatility and risk premiums. This "prudential uncertainty" channel is amplified by external habit, which makes the investor feel poorer and act more prudently. A calibrated production economy with a slow-moving, external habit shows that this channel can quantitatively explain return and dividend predictability regressions. The model matches numerous other moments, including the mean and volatility of the equity premium, the mean and volatility of the risk-free rate, and the second moments of output, consumption, and investment.

*Fisher College of Business, The Ohio State University, 760 Fisher Hall, 2100 Neil Avenue, Columbus OH 43210. Tel: (614)-292-8708; email: chen_1678@fisher.osu.edu.

†I am indebted to Aubhik Khan, Xiaoji Lin, René Stulz, Julia Thomas, and Lu Zhang for their training, advice, and inspiration.

1. Introduction

The mysterious nature of asset price fluctuations is a central question of financial economics. The present value identity tells us that these fluctuations must correspond to fluctuations in expected cash flows or to fluctuations in expected returns. The data show that movements in asset prices have little relationship with the more tangible culprit: movements in expected cash flows. Rather, these fluctuations correspond to a more elusive phenomenon: movements in expected returns.[1]

I show that movements in expected returns are a natural consequence of prudent savings behavior (in the sense of Kimball [1990]). The channel I propose runs through consumption volatility. Prudent behavior affects consumption volatility because, in bad times, the proper conduct of future precautionary savings becomes especially uncertain. News that things will get even worse would lead investors to hunker down, increasing savings at the expense of consumption. Good news would relieve them of their concerns, freeing them from austerity and boosting consumption. While they await the news, future consumption is uncertain, risk premiums are high, and asset prices are low. This dynamic is absent in good times, when investors are wealthy and have little need for savings as protection against uncertainty. External habit preferences amplify this countercyclical risk and make it quantitatively significant. Intuitively, an investor with external habit judges his consumption by a reference level. If this reference level is high, the investor feels especially poor and precautionary savings dynamics kick in with more punch. I illustrate this 'prudential uncertainty' mechanism with theoretical results from a two-period model.

To examine the quantitative significance of this channel, I study a standard real business cycle model where prudential uncertainty is amplified by a slow-moving, external habit. Regressions on simulated data show that the model captures the nature of asset price fluctuations. Fluctuations in the price-dividend ratio have little relationship with future dividends, but have a tight relationship with future returns. The model matches a long list of asset price moments beyond the return and dividend predictability regressions. The model matches the mean, volatility, and persistence of the equity premium, the mean, volatility and persistence of the risk free rate, and the volatility and persistence of the price-dividend ratio. This asset pricing performance comes with good predictions about the real economy. The model produces data-like volatilities, autocorrelations, and cross-correlations of consumption, output, and investment. Taken together, these results show that the model has promise as a unified framework for understanding aggregate asset prices and the real economy.

[1] Cochrane [2011] argues that this is "the central organizing question of asset-pricing research." He also reviews the large literature linking valuations with expected cash flows and returns. This stylized fact was first documented by Shiller [1981] and LeRoy and Porter [1981].

In the quantitative model, the household has external habit preferences which are similar to those from Campbell and Cochrane [1999], but with an important difference: the habit sensitivity function $\lambda(s_t)$ is a constant. The habit sensitivity function controls the conditional volatility of habit, and the Campbell-Cochrane $\lambda(s_t)$ assumes countercyclical volatility. Changing $\lambda(s_t)$ to a constant implies that all cyclicality in the maximum Sharpe ratio must come from cyclicality in consumption volatility. I find that the model endogenously generates countercyclical consumption volatility which quantitatively mimics the Campbell-Cochrane $\lambda(s_t)$. One interpretation of the model's asset pricing performance is that prudential uncertainty endogenizes the Campbell-Cochrane mechanism.

An investor is prudent if the third derivative of the his value function is positive. Almost all models of asset prices feature investors which satisfy this criterion. Thus, the prudential uncertainty mechanism is active in most of the literature. For the channel to be quantitatively significant, however, there needs to be a force which amplifies the channel. This paper shows that external habit is one force which does the job. Other candidate amplification mechanisms include investor collateral constraints (Favilukis, Ludvigson, and Van Nieuwerburgh [2010]), countercyclical bond payments (Guvenen [2009]), and time-varying disaster risk (Gourio [2010]). Like external habit, all of these channels make investors feel poorer and act more prudently. Indeed, each of these papers generates time-varying risk premiums in settings with prudent investors. The theoretical results from this paper help shed light on the origins of time-varying risk premiums in these models.

The body of the paper is divided into two main sections. Section 2 presents the two-period model and theoretical results. The model is intended to provide a clear illustration of prudential uncertainty and favors simplicity over generality. It examines the consumption-savings problem of a prudent investor with an external and additive habit. He faces a wealth shock and can invest his wealth in a single risky asset. I prove that the investor's consumption volatility is countercyclical, that is, consumption volatility is decreasing in initial wealth. This result is a consequence of the shape of the consumption function. With positive habit, the investor's problem falls into the broad class of models which generates a strictly concave consumption function (Carroll and Kimball [1996]). This shape implies that, as argued long ago by Keynes [1936], the marginal propensity to consume (MPC) is greater at low levels of wealth. Since the MPC describes the sensitivity of consumption to shocks, this shape implies that consumption is more volatile when the investor is poor. I also prove that this countercyclicality is larger for investors with a higher level of habit. Mathematically, an additive habit acts just like a reduction in income. This makes the household feel less wealthy and increases the concavity of the consumption function. The Hansen and Jagannathan [1991] bound shows that this prudential uncertainty mechanism can lead to countercyclical risk premiums.

Section 3 examines the quantitative significance of the channel. This section presents a real business cycle model with a slow-moving external habit and convex capital adjustment costs. To focus on the prudential uncertainty channel, I eliminate all exogenous sources of time-varying volatility. Productivity is the only exogenous stochastic process, and it follows the standard, homoskedastic AR(1) form. This assumption, combined with the constant habit sensitivity, implies that a key driver of time-varying risk premiums is prudential uncertainty.

I calibrate the model's nine parameters to unconditional asset price and quantity moments. The calibration targets include only three asset price moments: the mean risk-free rate, the persistence of the risk-free rate, and the mean Sharpe ratio. Simulated data show that the model generates good predictions for a long list of other asset price moments. They also demonstrate that the time-varying risk premiums in the model are driven by a reasonable amount of time-varying consumption volatility. GARCH and non-parametric measures of heteroskedasticity show that the model generates less time-variation in consumption volatility than is seen in the data. As in the data, the price-dividend ratio predicts consumption volatility with a negative sign. This predictability is statistically significant, and its economic magnitude is roughly half the data's.

Relation to the literature The quantitative model belongs to the literature on asset pricing models with habit formation and production. Most models in this literature consider a short-lived, internal habit (Jermann [1998] and Boldrin, Christiano, and Fisher [2001]). In contrast, I consider a slow-moving, external habit, in the style of Campbell and Cochrane [1999]. Without this modification, the model would produce a counterfactually volatile risk-free rate. Lettau and Uhlig [2000] do consider a slow-moving, external habit in a production economy, but they assume costless adjustment of the capital stock. In contrast, I model convex capital adjustment costs. Without such costs, business cycle models imply a counterfactually smooth Tobin's Q (Boldrin, Christiano, and Fisher [1999]). Dew-Becker [2011] does consider a slow-moving, external habit, but his preferences are recursive with a high elasticity of intertemporal substitution. As a result, the forces driving is asset pricing results are very different than those in my model.

More broadly, the quantitative model fits into the literature on asset pricing in production economies. To my knowledge, none of the previous models in this literature is able to capture the relationship between the price-dividend ratio, returns, and dividend growth that is seen in the data. Guvenen [2009]'s limited participation model, Dew-Becker [2011]'s time-varying risk aversion model, and Favilukis, Ludvigson, and Van Nieuwerburgh [2010]'s housing collateral model can explain return predictability regressions, but they either do not examine or cannot explain the lack of dividend predictability in the data. Gourio [2009]'s time-varying disaster model, Favilukis and Lin [2011]'s infrequent renegotiation model, and Kuehn, Petrosky-Nadeau, and Zhang [2012]'s labor search model all generate time-varying risk premiums, but they do not discuss predictive regressions using the price-dividend ratio. All of the previously mentioned models feature prudent investors and so the prudential uncertainty channel should be active in each of these models. This paper's contribution is to clearly illustrate prudential uncertainty and to show how one friction in particular, external habit, can amplify the channel. Showing that prudential uncertainty is amplified by any of the previously mentioned frictions is a question for future research.

Regarding the qualitative results, a related paper is Posch [2011], who shows that non-linearities introduced by standard production technologies generate concave consumption and countercyclical risk premiums. My paper builds on his results by demonstrating the link between prudence, concave consumption, and consumption volatility. I also demonstrate that external habit amplifies this concavity and can generate a good fit to the data on asset prices.

The two-period model builds off of the buffer stock / precautionary savings literature (Zeldes [1989], Deaton [1991], Carroll [1994], Carroll [1997], among others). A key result of this literature is that prudence generally results in a consumption policy which is strictly concave in wealth (Carroll and Kimball [1996]). This paper extends the literature by examining the effects of this consumption policy on asset prices.

2. The Two-period Model

This section illustrates the prudential uncertainty mechanism with a two-period model. The model is a consumption-savings problem for a prudent investor with external habit who faces uncertainty in the return on savings. I prove two propositions and two corollaries. Proposition 1 and Corollary 1 show the consumption volatility is countercyclical. Proposition 2 and Corollary 2 show that habit amplifies this countercyclicality. The proofs are found in the Appendix.

Consider a two-period-lived investor with an external habit stock $H > 0$ and period utility

$$u(C - H) = \frac{(C - H)^{1-\gamma} - 1}{(1 - \gamma)}$$

with $\gamma > 0$. This specification of utility features prudence, that is, $-\frac{u'''(C-H)}{u''(C-H)} > 0$. This feature is required for the investor to display a precautionary savings motive (Kimball [1990]). He has no time-preference, and, for simplicity, I assume that his habit stock does not evolve over this time horizon.

At date 0, the investor has wealth W_0. Nothing occurs at date 0, but the date helps to serve as a reference point. Throughout this section, I describe a variable as 'countercyclical' if it is negatively related to W_0. At date 1, the investor receives a wealth shock ΔW_1, making his wealth $W_1 = W_0 + \Delta W_1$. He consumes C_1, and saves the rest of his wealth $W_1 - C_1$ in a risky asset. At date 2, the return on the risky asset R_2 is realized, and he consumes his remaining wealth $R_2(W_1 - C_1)$.

The model comes down to finding the investor's optimal consumption at date 1. His optimal consumption policy solves

$$C_1(W_1|H) = \arg\max_{\{C_1\}} u(C_1 - H) + \mathbb{E}_1[u(C_2 - H)] \tag{1}$$
$$s.t. \quad C_2 = R_2(W_1 - C_1)$$

Taking a Taylor of expansion of $C_1(W_1|H)$ around W_0, the volatility of C_1 is approximately

$$\begin{aligned} \sigma_0[C_1(W_1|H)] &\approx \sigma_0[C_1(W_0|H) + C_1'(W_0|H)\Delta W_1] \\ &= C_1'(W_0|H)\sigma_0[\Delta W_1] \end{aligned} \tag{2}$$

That is, consumption volatility is proportional to the marginal propensity to consume (MPC). Intuitively, the MPC captures the responsiveness of consumption to shocks. So the higher the MPC, the more responsive, and thus the more volatile, is consumption. For the remainder of this section, I assume that ΔW_1 is small enough so that equation (2) is a good approximation.

Unfortunately, even in a two-period consumption savings problem, the presence of prudence typically precludes closed-form solutions (Carroll [2001]). However, using the methods from Carroll and Kimball [1996], I can still prove that the solution exhibits countercyclical volatility and that habit amplifies this cyclicality.

Proposition 1. *The date 1 MPC is decreasing in wealth, that is,* $C_1''(W_1|H) < 0$.

The essence of the proof is that, as long as H is positive, the model falls into the broad class of models for which consumption is strictly concave (Carroll and Kimball [1996]). Carroll and Kimball do not show the strict concavity result for this model, but I show in the Appendix that their proofs can be extended to include this setting. The proof is rather technical, but, intuitively, a prudent investor will save for precautionary reasons in the presence of uncertainty. As the agent becomes wealthier, the uncertainty becomes less relevant, and this motive declines, creating convex savings and concave consumption.

Corollary 1. *The volatility of C_1 is countercyclical, that is,* $\frac{\partial}{\partial W_0}\sigma_0[C_1(W_1|H)] < 0$.

The proof is short and helps to illustrate the role of the MPC, so I state it here.

Proof. Because the MPC is decreasing in wealth, and the volatility of C_1 is proportional to the MPC (equation (2)), the volatility of C_1 is decreasing in wealth. $\square$

Proposition 1 and its corollary capture the "prudential uncertainty" mechanism. The intuition for this result stems from the effect of the date 1 wealth shock on the investor's desire for precautionary savings at date 1. A positive shock weakens this desire, decreasing savings and boosting consumption, while a negative shock encourages him to hunker down, with opposite effects. From the perspective of date 0, this uncertainty in the need for prudence at date 1 causes additional consumption volatility. Since the need for prudence intensifies at low levels of wealth, this 'prudential uncertainty' is countercyclical.

The Hansen-Jagannathan bound shows how Corollary 1 can lead to time-varying risk premiums. In standard models, the Hansen-Jagannathan bound is approximately

$$\max_{\{\text{All Assets}\}} \frac{\mathbb{E}_t[R_{t+1} - R_{f,t+1}]}{\sigma_t(R_{t+1})} \approx \sigma_t[\log M_{t+1}] \propto \sigma_t[\log C_{t+1}]$$

Where M_{t+1} is any investor's SDF and $R_{f,t+1} \equiv 1/E_t[M_{t+1}]$. The conditional maximum Sharpe ratio is approximately the volatility of marginal utility (the SDF) which in turn is related to the volatility of consumption. Thus, the countercyclical consumption volatility of Corollary 1 promotes countercylical risk premiums. This formula requires neither general equilibrium nor complete markets. Unfortunately, it cannot be related directly to the two-period model because the investor does not make a savings decision at date 0, and thus he does not have an Euler equation related to $\sigma_0(C_1)$. Introduction of this savings decision significantly complicates the analysis and without a compensating benefit in intuition.

Mechanically, prudential uncertainty is manifested as a strictly concave consumption function. Strictly concave consumption is an implication of a broad class of consumption-savings models (Carroll and Kimball [1996]). The magnitude of this channel, however, is typically too small to match the data on time-varying risk premiums (Posch [2011]). The following proposition and corollary show that habit amplifies this channel.

Proposition 2. *Habit increases the sensitivity of the MPC to wealth, that is, for any increase in habit* $\epsilon > 0, \quad C_1''(W_1|H+\epsilon) < C_1''(W_1|H)$

To understand this result, note that a investor with high habit has become accustomed to a high standard of living. He judges consumption not by its absolute level, but by how much it exceeds this standard. As a result, he acts with much more prudence than a investor who is accustomed to living in poverty. The intuition from Proposition 1 kicks in with more punch, and the MPC becomes even more sensitive to wealth. Mathematically, an external and additive habit acts as a reduction in income, which reduces wealth and amplifies prudential behavior.

Corollary 2. *As habit increases, the volatility of C_1 becomes more countercyclical. That is, for any set of initial wealth $\mathbb{W}_0$, and increase in habit $\epsilon > 0$,*

$$\max_{W_0 \in \mathbb{W}_0} \sigma_0[C_1(W_1|H+\epsilon)] - \min_{W_0 \in \mathbb{W}_0} \sigma_0[C_1(W_1|H+\epsilon)] > \max_{W_0 \in \mathbb{W}_0} \sigma_0[C_1(W_1|H)] - \min_{W_0 \in \mathbb{W}_0} \sigma_0[C_1(W_1|H)]$$

One can think of the set of initial wealth $\mathbb{W}_0$ as the volatility of the business cycle in the economy. Corollary 2 states that, over this business cycle, the range of consumption volatilities for the high habit investor (the LHS) is larger than the range of consumption volatilities for the low habit investor (the RHS). The proof simply applies the link between the MPC and consumption volatility (equation (2)) to Proposition 2. Because habit makes the MPC more countercyclical, habit also makes consumption volatility more countercyclical. An application of the Hansen-Jagannathan bound shows that habit can increase the magnitude of countercylical risk premiums.

Figures 1 and 2 illustrate the propositions of this section. The figures show results from numerical solutions of the two-period model. Figure 1 plots date 1 consumption and MPC as a function of wealth for various levels of habit. The left panel shows that, as long as habit is positive, consumption is strictly concave. The right panel shows that this concavity is reflected in an MPC which decreases in wealth (Proposition 1). Figure 1 also shows that habit intensifies this relationship (Proposition 2). As the line gets lighter, habit increases, and the slope of the MPC gets steeper. Intuitively, habit makes the investor feel poorer and act more prudently.

Figure 2 shows how these consumption policies are reflected in consumption volatility. As long as habit is positive, consumption volatility decreases in wealth (Corollary 1). This is the essence of the prudential uncertainty mechanism. Because prudence prescribes a countercyclical MPC, and because consumption volatility is proportional to the MPC (equation (2)), consumption volatility is countercyclical. The figure also shows that habit amplifies this countercyclicality (Corollary 2). As the line gets lighter, habit increases, and the range that is spanned by consumption volatility increases.

The results of this section may come as a surprise since many models in the finance literature produce consumption policies which are linear in wealth. These linear consumption policies are the result of the

Figure 1: Two-Period Model: Date 1 Consumption and MPC

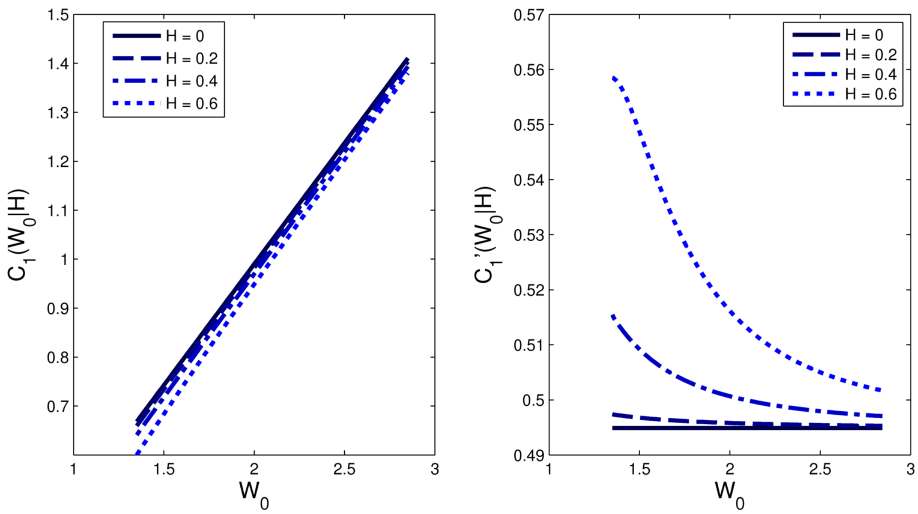

Figure 2: Two-Period Model: Date 0 Consumption Volatility

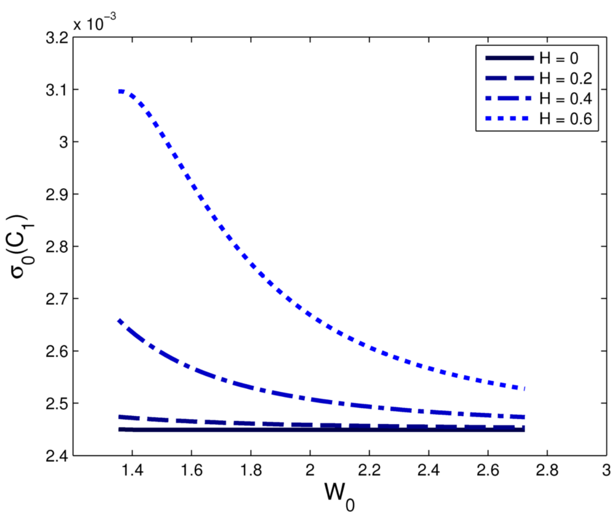

fact that much of this literature is set in continuous time (Sundaresan [1989], Constantinides [1990]) or relies on log-linear approximations (Campbell [1994], Lettau [2003], Kaltenbrunner and Lochstoer [2010]). Continuous time allows the investor to make an infinite number of trades within any trading period. This instantaneous trading allows for consumption policies which, if applied in a discrete time setting, would imply a positive probability that marginal utility becomes infinite (Brandt [2009]). It is exactly fear of hitting this condition of infinite marginal utility which seems to drive concave consumption behavior (Attanasio [1999]) and this behavior is critical for the prudential uncertainty mechanism. Log-linear approximations, on the other hand, implicitly assume that consumption policies are (log) linear in state variables. The non-linearities introduced by combining power utility functions with linear budget constraints are important for generating strictly concave consumption (Posch [2011]).

The results may also appear to conflict with the common intuition that habit encourages smooth consumption. This intuition is straightforward in internal habit models, where an increase in consumption has a direct effect of lowering utility in later periods (Sundaresan [1989]). In an external habit model, however, the investor by assumption does not take into account this indirect effect. Prices may encourage smooth consumption via general equilibrium effects, but the habit itself acts very similarly to a reduction in income which increases the marginal propensity to consume (Carroll and Kimball [1996]). I show in Section 3 that the decision-theoretic mechanisms of this section are have quantitatively significant effects in general equilibrium.

3. The Quantitative Model

This section examines the quantitative significance of the prudential uncertainty mechanism described in Section 2. I present the model (Section 3.1), describe the data and calibration (Section 3.2), and then describe the solution method and characterize the laws of motion (Section 3.3).

Section 3.4 discusses the main quantitative results. It shows that the model provides a good quantitative description of time-varying risk premiums as measured by regressions of returns and dividends on the price-dividend ratio. It also demonstrates how the prudential uncertainty channel mimics the Campbell and Cochrane [1999] $\lambda(s_t)$, and shows that the time-varying risk premiums are driven by a reasonable amount of time variation in consumption volatility. Section 3.5 shows results regarding unconditional moments. This section shows that the model matches a long list of asset price moments while being consistent with business cycle data. Section 3.6 examines comparative statics.

3.1. Model Setup

The quantitative model is a standard real business cycle model, but where prudential uncertainty is amplified by an external, slow-moving habit in the style of Campbell and Cochrane [1999]. There is a representative household and representative firm. Time is discrete, the horizon is infinite, and markets are complete. The setting is chosen to be a minimal deviation from the standard model in order to maintain clarity about the source of the quantitative results.

The model is designed to focus on the prudential uncertainty channel. I do not model uncertainty shocks (Bloom [2009]), slow-moving technological changes in volatility (Lettau, Ludvigson, and Wachter [2008]), government policy uncertainty (Pastor and Veronesi [2012]) or any other sources of time-varying uncertainty. While incorporating these other sources of time-varying uncertainty would help the model fit the data, it would make it more difficult to identify the quantitative significance of prudential uncertainty.

For the remainder of the paper, I denote log variables with lowercase, i.e. $z_t \equiv \log Z_t$.

3.1.1. Representative Household

There is a continuum of identical households with external habit formation preferences. Each household chooses its asset holdings to maximize

$$\mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \beta^t \frac{(C_t - H_t)^{1-\gamma} - 1}{1-\gamma} \right\} \tag{3}$$

Where H_t, the level of habit, is taken as external to the household. External habit is included to amplify the prudential uncertainty channel. For simplicity, the household does not value leisure and is endowed with a unit of labor. This utility function is the same as the one used in Section 2.

I specify the evolution of habit using surplus consumption, rather than the level of habit itself. That is, let

$$\hat{S}_t \equiv \frac{\hat{C}_t - H_t}{\hat{C}_t} \tag{4}$$

be the surplus consumption ratio, where the hats denote aggregates. This approach is chosen for comparability with the existing literature on external habit (Campbell and Cochrane [1999], Santos and Veronesi [2010], among others) and is not critical for the prudential uncertainty channel. Indeed, the analysis of the channel from Section 2 does not depend on the evolution of habit.

Surplus consumption evolves according to an autoregressive process

$$\hat{s}_{t+1} = (1-\rho_s)\bar{s} + \rho_s \hat{s}_t + \lambda(\hat{c}_{t+1} - \hat{c}_t) \tag{5}$$

where λ is a *constant*. The constant λ is not important for generating the prudential uncertainty mechanism, but it helps in isolating the channel. Equation (5) shows that λ controls the conditional volatility of the habit process. By choosing λ to be a constant, I force all time-varying volatility in my model to be endogenously generated. Though apparently innocuous, the constant λ represents a significant departure from Campbell and Cochrane [1999]. To see this, note that the conditional maximum Sharpe ratio is[2]

$$\max_{\{\text{all assets}\}} \left\{ \frac{\mathbb{E}_t(R_{t+1} - R^f_{t+1})}{\sigma_t(R_{t+1})} \right\} \approx \gamma(\lambda+1)\sigma_t(\Delta c_{t+1}) \tag{6}$$

Since λ is constant, the maximum Sharpe ratio can vary over time if and only if consumption volatility is time-varying. This is in direct contrast to Campbell and Cochrane [1999], where λ is assumed to be time-varying and consumption volatility is assumed to be constant. I discuss this issue further in the section on mimicking the Campbell-Cochrane $\lambda(s_t)$ (Section 3.4.2).[3]

A persistent habit is important for capturing the persistence in various asset prices. The data show that both the price-dividend ratio and the (ex-ante) risk-free rate are very persistent. As a result, the calibrated ρ_s will be close to one, meaning that habit depends on a very long history of consumption. This long dependence is in contrast with previous habit models in production economies (Jermann [1998] and Boldrin, Christiano, and Fisher [2001]) which assume that habit depends only on last quarter's consumption.

The external nature of habit produces the simple stochastic discount factor

$$M_{t,t+1} = \beta \left(\frac{\hat{C}_{t+1}}{\hat{C}_t} \frac{\hat{S}_{t+1}}{\hat{S}_t} \right)^{-\gamma} \tag{7}$$

The stochastic discount factor has the traditional consumption growth term $\left(\frac{\hat{C}_{t+1}}{\hat{C}_t}\right)^{-\gamma}$, but habit adds a term due to changes in habit $\left(\frac{\hat{S}_{t+1}}{\hat{S}_t}\right)^{-\gamma}$. This SDF is exactly the same as that in Campbell and

[2]The log-SDF is $m_{t+1} = \log\beta - \gamma\Delta s_{t+1} - \gamma\Delta c_{t+1}$ and the habit process is $\Delta s_{t+1} = -(1-\rho_s)(s_t - \bar{s}) + \lambda_t \Delta c_{t+1}$. Plug the habit process into the log-SDF, then assume that the SDF is log-normal, and we have $\frac{\sigma_t(M_{t+1})}{\mathbb{E}_t(M_{t+1})} \approx \sqrt{e^{\mathrm{Var}_t[m_{t+1}]} - 1} \approx \sigma_t[m_{t+1}] = \gamma[\lambda_t + 1]\sigma_t(\Delta c_{t+1})$.

[3]The constant λ also causes the issue that current habit can move negatively with current consumption. These preferences still capture the notion of habit in that H_{t+1} is comprised of a geometric average of the history of consumption. This issue is discussed further in the Appendix

Cochrane [1999] but is very different from the related EZ-habit model of Dew-Becker [2011]. Dew-Becker's SDF also has two components, but the two components are those generated by Epstein and Zin [1989] preferences: a consumption growth component and a term related to the return on wealth. In Dew-Becker's model, habit enters the SDF by affecting the curvature (power parameter) of the return on wealth component. This SDF is very different from the SDF in equation (7) and shows that the two models have very different mechanisms.

3.1.2. Representative firm

The production side of the economy is standard. The only feature is convex capital adjustment costs.

There is a unit measure of identical firms which produce consumption using capital K_t and labor N_t. Production is given by

$$\Pi(K_t, Z_t, N_t) \equiv A Z_t K_t^{\alpha} N_t^{1-\alpha} \tag{8}$$

Where α is capital's share of output, Z_t is productivity, and A is chosen so that the non-stochastic steady state capital stock is one. The choice of A does not have a material effect on the relevant model-simulated moments, but keeping the steady state capital stock near one helps with the accuracy of the numerical methods.

The Cobb-Douglas specification of production, combined with the fact that the household does not value leisure implies that wages are equal to the marginal product of labor

$$W_t = (1-\alpha) A Z_t \hat{K}_t^{\alpha} \tag{9}$$

Productivity Z_t follows the AR(1) process

$$z_{t+1} = \rho_x z_t + \sigma_z \epsilon_{z,t+1} \tag{10}$$

Where $\epsilon_{z,t+1}$ is a standard normal i.i.d. shock. Productivity is the only source of uncertainty in the model. Note that it has a constant conditional variance σ_z.

Capital accumulates according to the usual capital accumulation rule

$$K_{t+1} = I_t + (1-\delta) K_t \tag{11}$$

and firms face a convex capital adjustment cost

$$\Phi(I_t, K_t) = \frac{\phi}{2}\left(\frac{I_t}{K_t} - \delta\right)^2 K_t \tag{12}$$

This formulation of adjustment costs punishes the firm for deviating from the non-stochastic steady state investment rate of δ. I assume that the adjustment costs are a pure loss. They do not represent payments to labor. Adjustments costs are included because general equilibrium production economies produce a counterfactually smooth Tobin's Q unless one includes an investment friction. Quadratic costs are chosen for simplicity, but a richer model would incorporate investment frictions by modeling an investment good sector, as in Boldrin, Christiano, and Fisher [2001], or would feature heterogeneous plants with non-convex costs of adjustment, as in Khan and Thomas [2008].

The firm's objective is standard

$$\max_{\{I_t, K_{t+1}, N_t\}} \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} M_{0,t}[\Pi(K_t, Z_t, N_t) - W_t N_t - \Phi(I_t, K_t) - I_t] \right\} \tag{13}$$

It chooses investment, capital, and labor to maximize future dividends, discounted with the household's stochastic discount factor.

3.1.3. Equilibrium

Market clearing is standard:

$$\hat{C}_t + \hat{I}_t = Z_t \hat{K}_t^{\alpha} - \Phi(\hat{I}_t, \hat{K}_t) \tag{14}$$

Due to the consumption externality, the welfare theorems do not hold and the equilibrium cannot be easily described by a social planner's problem. Thus, I define equilibrium competitively. Since markets are complete, the household's side of the equilibrium boils down her SDF (7).

I define equilibrium recursively. The aggregate state variables are aggregate capital $\hat{K}$, surplus consumption $\hat{S}$, and productivity Z.

Definition. *Equilibrium is a firm decision rule for investment $I(K; \hat{K}, \hat{S}, Z)$, cum-dividend value $V(K; \hat{K}, \hat{S}, Z)$, aggregate consumption $\hat{C}(\hat{K}, \hat{S}, Z)$, and a law of motion of aggregate capital $\Gamma(\hat{K}, \hat{S}, Z)$ such that* *(i) Firms optimize: $I(K;\hat{K},\hat{S},Z)$ and $V(K;\hat{K},\hat{S},Z)$ solve*

$$V(K;\hat{K},\hat{S},Z) = \max_{\{N,I,K'\}} \Big\{ \Pi(K,Z,N) - W(\hat{K},\hat{S},Z)N - \Phi(I,K) - I \\ + \int_{-\infty}^{\infty} dF(\epsilon')M(\hat{K},\hat{S},Z;Z')V(K';\hat{K}',S',Z') \Big\} \tag{15}$$

subject to capital accumulation (11), competitive wages (9), the productivity process (10), the habit process (5), the SDF is equal to the household's IMRS (7), $\hat{K}' = \Gamma(\hat{K},\hat{S},Z)$, and where $F(\epsilon')$ is the standard normal CDF.

(ii) Markets clear and aggregates are consistent with individual behavior:

$$\hat{C}(\hat{K}',\hat{S}',Z') = \Pi(\hat{K},Z,1) - \Phi(I(\hat{K};\hat{K},\hat{S},Z),\hat{K}) - I(\hat{K};\hat{K},\hat{S},Z) \tag{16}$$

$$\Gamma(\hat{K},\hat{S},Z) = (1-\delta)\hat{K} + I(\hat{K};\hat{K},\hat{S},Z) \tag{17}$$

Since households and firms are identical, in equilibrium, $\hat{K} = K$, $\hat{S} = S$ and $\hat{C} = C$. Thus, in what follows, I drop the hats.

3.2. Data and Calibration

3.2.1. Data Sources and Definitions

The empirical moments use post-war (1948Q1-2010Q4) data from the CRSP, BEA, and Simona Cociuba's website. Some authors argue for using the longest sample available when evaluating consumption-based asset pricing models (e.g. Bansal, Kiku, and Yaron [2009]). However, in a production economy I must also address the data on aggregate investment and output. The nature of this data is significantly affected by using pre-war data. For example, the correlation between investment growth and output growth is 0.23 for the sample 1929-2011, vs 0.73 for the post-war sample. This consideration leads me to target only post-war data.

Consumption is real per capita non-durable goods and services consumption. This measure excludes volatile consumer durables such as automobiles, which theoretically produce a smooth consumption flow over the life of the durable. Aggregate equity is represented by the CRSP value-weighted index, adjusted for inflation with the consumer price index. Dividends are calculated with the assumption that all dividends are reinvested in the stock market. This method of aggregation preserves the Campbell and Shiller [1988] present value identity. The risk-free rate is the 30-day T-bill return, also adjusted with the consumer price index. Further details regarding the data are found in the Appendix.

Table 1: Calibration

The model is quarterly, and all parameter values are quarterly. Empirical figures are quarterly except where noted. The sample period is 1948Q1-2010Q4. The persistence of the ex-ante 30-day T-bill return is marked with a † to indicate that the figure is taken from Beeler and Campbell [2009] and corresponds to 1947Q2-2008Q4. Consumption is real non-durable goods and services consumption. The volatility of GDP and relative volatility of consumption are logged and HP-filtered with a smoothing parameter of 1600. Further details regarding the data are found in the Appendix

Parameter		Value	Empirical Target	
Long-Run Growth Moments				
α	Capital Share	0.35	Labor's Share of Output	0.65
δ	Depreciation Rate	0.016	Mean Investment Rate (Growth-Adjusted)	0.016
Unconditional Business Cycle Moments				
σ_z	Volatility of Productivity	0.0133	Volatility of GDP (%)	1.73
ρ_z	Persistence of Productivity	0.979	Persistence of Solow Residual	0.979
ϕ	Adjustment Cost	103	Relative Volatility of Consumption to GDP	0.48
Unconditional Asset Price Moments				
β	Time Preference	0.974	Mean 30-Day T-bill Return (annual)	0.89
ρ_s	Persistence of Surplus Consumption	0.966	Persistence of ex-ante 30-Day T-bill Return (annual)	$0.84^\dagger$
$\bar{S}$	Steady-State Surplus Consumption Ratio	0.065	Mean Sharpe Ratio of CRSP Index (annual)	0.44
Chosen Outside of the Model				
γ	Utility Curvature	2	(Chosen to match Campbell and Cochrane (1999))	

3.2.2. Calibration

Table 1 shows the calibration targets and the resulting parameter values. The model's nine parameters values are chosen, as much as possible, to fit unconditional moments of aggregate asset prices and quantities.

I identify the five technological parameters values with moments of the real economy. Two of the technological parameters are chosen to fit the data moments on long-run growth. Specifically, the capital share $\alpha = 0.35$ and the depreciation rate $\delta = 0.016$ are chosen to fit the average labor share of output and the average growth-adjusted investment rate, respectively. The three remaining technology parameters are chosen to match the business cycle. I choose the volatility of productivity $\sigma_z = 0.0133$

to match the volatility of HP-filtered log GDP of 1.73. I choose the adjustment cost parameter $\phi = 103$ to match the relative volatility of non-durables and services consumption to GDP (logged and HP-filtered) of 0.48. This choice results in mean adjustment costs as a percentage of output of less that 1%. While the relative volatility of investment is a more intuitive identification for this parameter, there is much less consensus about the appropriate empirical measure of investment.[4] The ideal target for ρ_z is the autocorrelation of HP-filtered log GDP. Unfortunately, standard RBC models have difficulty reproducing the persistence of GDP growth (Cogley and Nason [1995]).[5] Thus, I use $\rho_z = 0.979$, the estimated persistence of the Solow residual with a fixed labor input.

The four preference parameter values are identified as much as possible using asset price data. Two of the preference parameters are chosen to match moments of the risk-free rate. Because time-preference is reflected in the risk-free rate, I set $\beta = 0.974$ to match the mean annual return on 30-day T-bills of 0.89. This value for β is rather low because the model features a non-trivial precautionary savings motive, and a low level of patience helps counteract that motive. Because the persistence of surplus consumption ρ_s describes the persistence of this motive, I choose $\rho_s = 0.966$ to match the annual autocorrelation of the (ex-ante) 30-day T-bill rate of 0.84.[6] The two remaining parameters, the utility curvature γ and the steady state surplus consumption ratio $\bar{S}$, jointly control risk aversion. As a result, it is difficult to identify them separately. To focus on the role of habit and to ease comparability with previous literature, I fix γ at Campbell and Cochrane [1999]'s value of 2. I then choose $\bar{S} = 0.065$ to match the mean annual Sharpe ratio of the CRSP index.

There are many more important dimensions of asset price data beyond these three targeted unconditional moments. The model is not calibrated to fit conditional moments regarding time-varying risk premiums. Nor is it calibrated to fit unconditional moments such as the mean equity premium, the volatility of the equity premium, the volatility of the price-dividend ratio, etc. These other moments can be considered over-identifying restrictions for examining the validity of the model.

3.3. Model Solution and Characterization

[4] Jermann [1998] uses fixed investment, which excludes government investment, and Papanikolaou [2011] uses non-residential fixed investment which excludes residential and government investment. On the other hand, Cooley and Prescott [1995] argue that the capital of the rbc model reflects capital used in government and household production sectors as well as inventories.

[5] One can reproduce the persistence of GDP growth by incorporating additional features, such as labor adjustment frictions as in Boldrin, Christiano, and Fisher [2001].

[6] I use Beeler and Campbell [2009]'s estimate for the autocorrelation of the ex-ante risk-free rate. They measure the ex-ante rate by forecasting the ex-post real return using past inflation and nominal yields.

3.3.1. Model Solution

I solve the model using a projection method (Judd [1992]). Specifically, I represent the law of motion for capital using cubic splines and then look for cubic spline coefficients which satisfy the firm's Euler equation. This method requires a good initial guess of the solution, and the external habit makes the standard RBC model an extremely bad guess. To overcome this difficulty, I use a homotopy method. I first solve the model with no habit, using the solution for the standard RBC model as an initial guess. I then gradually increase habit to full strength. This method is computationally intensive, and to help the computation I approximate the productivity process as a 15-point Markov chain using the Rouwenhorst method. Further details of the solution method are found in the Appendix.

Projection methods produce a global and non-linear solution. These properties are important for asset pricing models in general (Cochrane [2008b]), but are particularly important for capturing the mechanisms in this model. The main mechanism of this paper runs through prudence, which is related to the third derivative of the utility function. Moreover, prudence effects are particularly pronounced when the investor is threatened with infinite marginal utility (Attanasio [1999]).

3.3.2. Characterization of Equilibrium

Figure 3 shows the distribution of the model's state variables. The model has three aggregate states, capital, surplus consumption, and productivity, and this makes describing the distribution more complicated than in other asset pricing models. The univariate descriptions are fairly straightforward. Panel A shows the univariate distribution of each state variable. These distributions are constructed by computing histograms from simulated data. Each state variable is roughly log-normal on its own, and the productivity distribution displays the discreteness of the Markov chain approximation. The state variables interact, however. Panel B shows bivariate distributions as contour plots. This panel shows that all of the state variables are positively correlated. High capital, productivity, and surplus consumption all represent good times. This high correlation is particularly pronounced for surplus consumption and productivity, which can be seen in the lower right corner. Surplus consumption and productivity move together very closely. This equilibrium relationship is important for understanding the model's laws of motion which follow.

Figure 4 shows the law of motion for consumption. It shows four three-dimensional plots to provide a complete portrait of this function of three variables. The top two panels hold productivity fixed and show consumption as a function surplus consumption and capital. The bottom two panels hold capital fixed and show consumption as a function surplus consumption and productivity. Consumption is

Figure 3: Distribution of State Variables. Panel A shows simple histograms from simulated data. Panel B shows 2D histograms as contour plots.

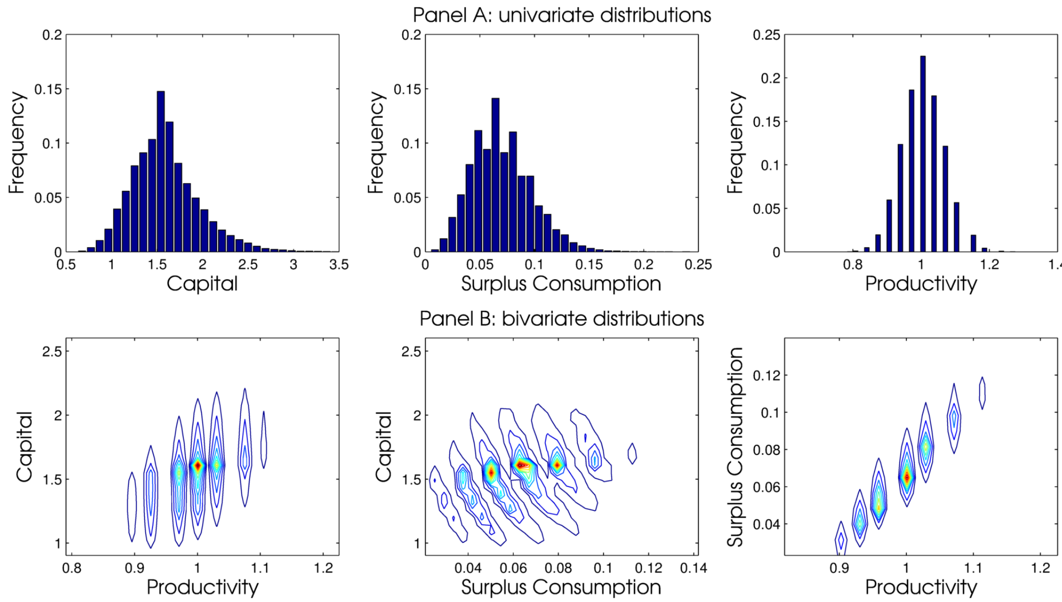

increasing and concave in capital and productivity, which is consistent with the prudential uncertainty mechanism. Consumption is decreasing and slightly convex in surplus consumption. This shape seems inconsistent with prudential uncertainty, however, surplus consumption and productivity are very highly correlated (Figure 3). As a result, the domain of the bottom panels which is explored in equilibrium runs from the bottom left to the upper right corner of the plots. Along this selection of the plot, consumption is strictly concave. Overall, Figure 4 shows that consumption is more sensitive to shocks when the household is poor.

Concave consumption results in countercyclical volatility, which is shown in Figure 5. This figure shows the law of motion for consumption volatility as a function of surplus consumption and productivity for two different levels of capital. I choose to focus on the state variables surplus consumption and productivity because most of the variation in consumption volatility occurs along these dimensions. Consumption volatility is high in bad times: either when surplus consumption is low or when productivity is low. This relationship is a natural result of prudential uncertainty. In bad times, future consumption and savings policies are particularly uncertain and this uncertainty leads to high volatility.

This countercyclical volatility is reflected in the Hansen-Jagannathan bound. Figure 6 plots this bound as a function of surplus consumption and productivity for various levels of capital. The

Figure 4: Law of Motion for Consumption.

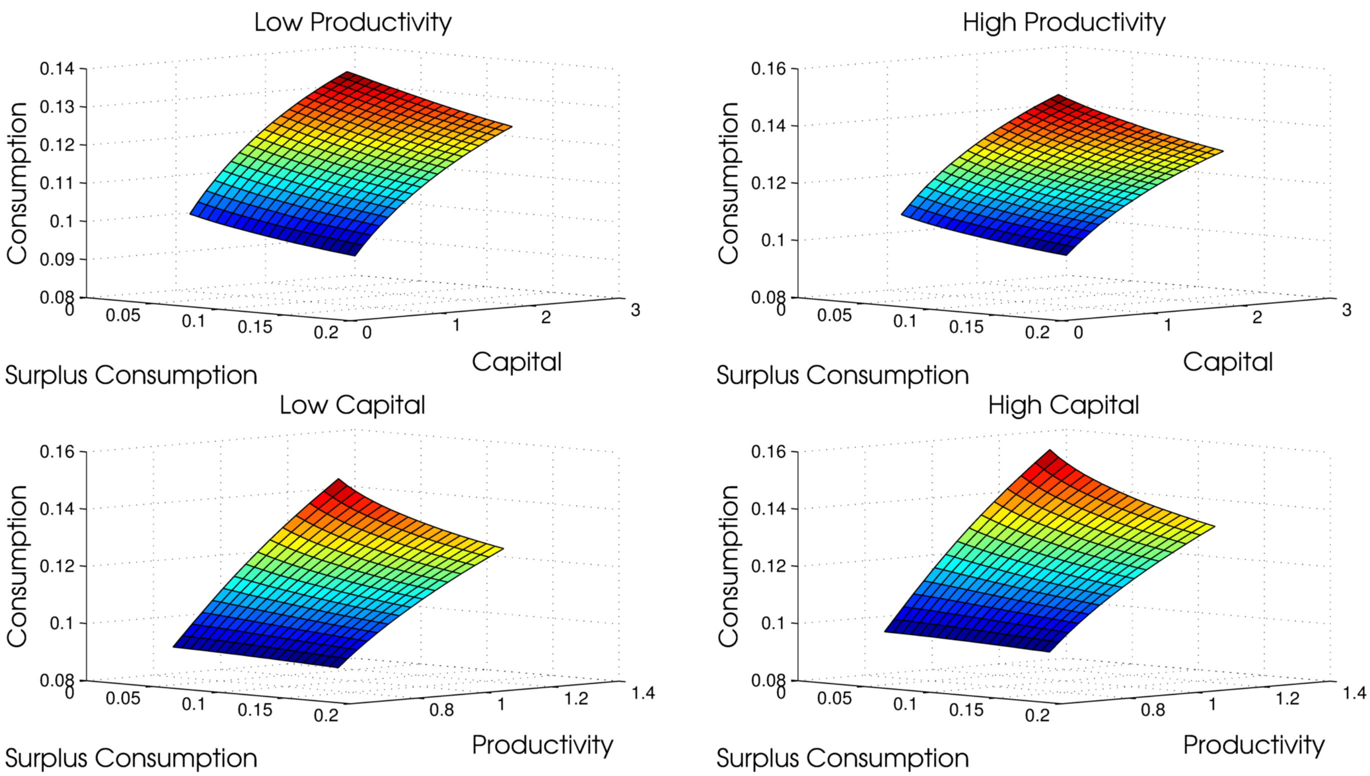

Figure 5: Law of Motion for Consumption Growth Volatility. Consumption growth volatility is the volatility of log first-differenced consumption.

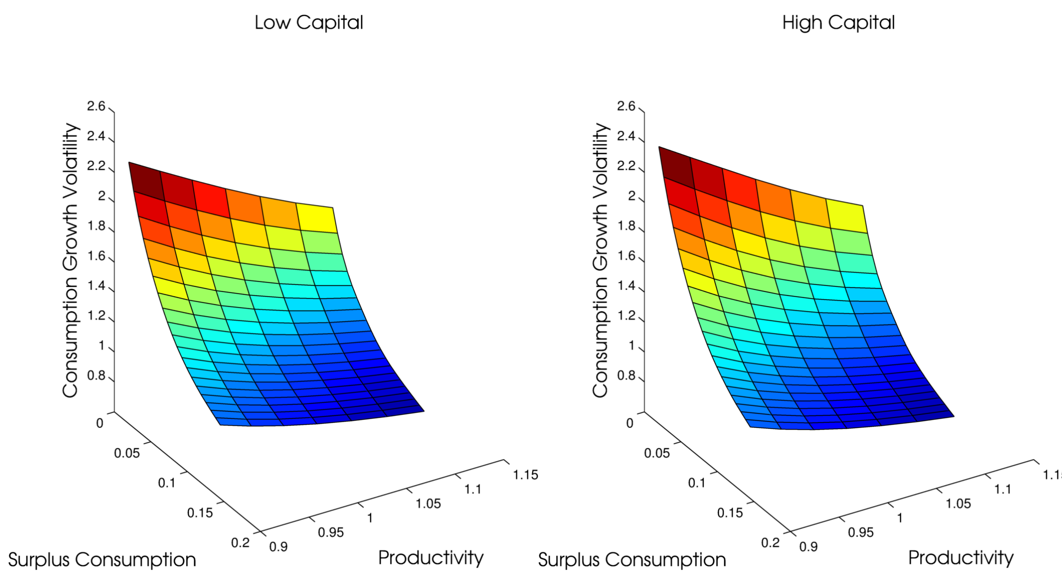

Figure 6: Law of Motion for the Hansen Jaganathan Bound.

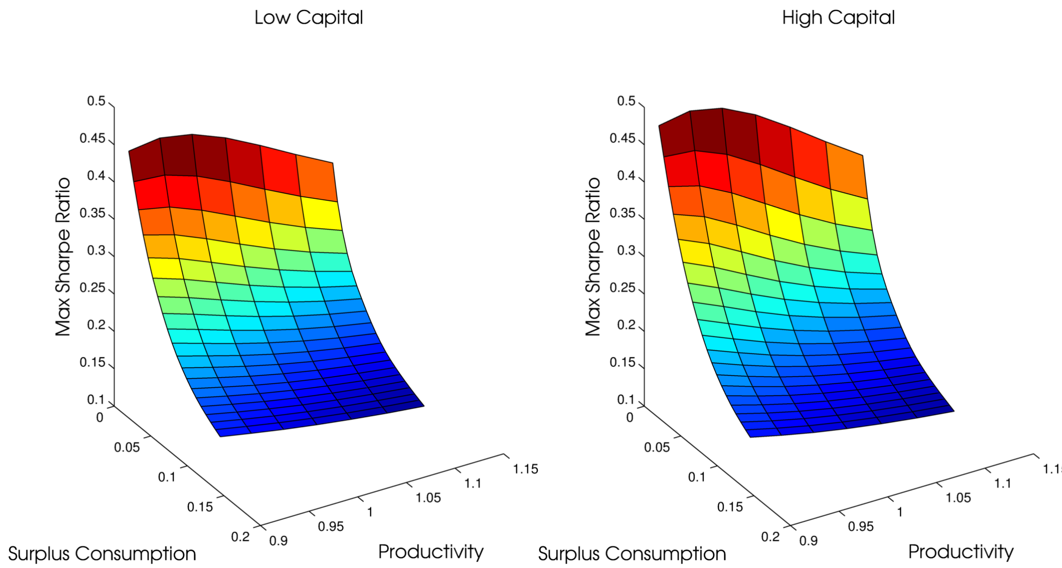

maximum Sharpe ratio declines in both surplus consumption and productivity. In bad times, when surplus consumption or productivity are low, consumption is especially volatile, making marginal utility especially volatile. This behavior results in investors demanding countercyclical returns for assets which are related to this countercyclical risk.

3.4. Time-Varying Risk Premiums

This section discusses the main quantitative results. It shows that the model generates quantitatively significant time-varying risk premiums, as reflected in regressions of returns and dividends on the price-dividend ratio. This section also compares the channel to Campbell and Cochrane [1999] and shows that time-varying risk premiums are driven by a reasonable amount of time-variation in consumption volatility.

All model results are from 1000 simulations of 504 quarters, and the first 252 quarters of each simulation are dropped to remove sensitivity to initial conditions. This simulation length is chosen to match the length of the empirical sample. In all of the quantitative results which follow, I measure moments for each of the 1000 simulations and report the mean, 5th percentile, median, and 95th percentile across the 1000 simulations.

I make no adjustments to account for un-modeled leverage. The price of equity is simply the present value of dividends from the representative firm, discounted with the household's SDF. Further details of the data methods are found in the Appendix.

3.4.1. Return and Dividend Predictability with the Price-Dividend Ratio

Table 2 shows the results of return and dividend predictability regressions on model-simulated data. Panel A shows regressions of future excess returns on today's price-dividend ratio for forecasting horizons between one and five years. Regression coefficients, standard errors, and R^2's, are close to the data values for all forecasting horizons. The coefficients on the price-dividend ratio are negative, and they are both economically and statistically significant. At the one-year horizon, the coefficient is -0.13 vs the data value of -0.12. To understand this economically, the volatility of the log price-dividend ratio is roughly 0.40 in both the model and the data. This means that a one standard deviation rise in the price-dividend ratio predicts a 4% reduction in the equity premium over the next year. The forecasting power increases with the forecasting horizon. R^2's for the five-year horizon are roughly 30% for both the model and the data. Overall, the model produces data-like fluctuations in expected returns.

The Campbell and Shiller [1988] present value identity tells us that the price-dividend ratio must forecast returns, or dividends, or both. Thus, the lack of dividend predictability is an important indication of time-varying expected returns (Cochrane [2008a]). Panel B shows that the model does a fair job of capturing this flip side of return predictability. The model predicts no relationship between the price-dividend ratio and future dividend growth at the one year horizon, as is seen in the data. Both in the model and in the data, the relationship gets stronger with the time horizon. The model somewhat overstates the magnitude of the relationship at long horizons, but the coefficients never become statistically significant. Moreover, the coefficients, standard errors, and R^2's are all within the 95% confidence bounds implied by model simulations.

Return and dividend predictability regressions help identify conditional expected returns because they are not directly observable in the data. The model, however, allows us to compute 'true' expected returns from the laws of motion for the economy. The left panel of Figure 7 shows scatterplots of model-computed expected returns against the price-dividend ratio from simulated data. The right panel shows expected dividend growth. As in the regressions estimates of Table 2, the price-dividend ratio has a strong negative relationship with expected returns. Its relationship with dividend growth, however, is mild and ambiguous.

The figure also shows the non-linearity and skewness in the relationship between expected returns

Table 2: Predicting Returns and Dividends with the Price-Dividend Ratio

Figures are annual. r_t, p_t, and d_t are the log-returns, prices, and dividends from the CRSP value-weighted index. r_t^f is the log-return on 30-day T-bills. The model columns show means and percentiles from 1000 simulations of the same length as the empirical sample. Further details are found in the Appendix.

Panel A: Predicting excess returns

$$\sum_{j=1}^{L} r_{t+j} - r_{t+j}^f = \alpha + \beta(p_t - d_t) + \epsilon_{t+L}$$

	L	Data 1948Q1-2010Q4	Model mean	Model 5%	Model 50%	Model 95%
$\hat{\beta}$	1	-0.12	-0.13	-0.25	-0.12	-0.04
	3	-0.28	-0.32	-0.61	-0.31	-0.11
	5	-0.39	-0.47	-0.83	-0.46	-0.17
SE($\hat{\beta}$)	1	0.05	0.05	0.03	0.05	0.09
	3	0.08	0.09	0.05	0.09	0.15
	5	0.12	0.11	0.06	0.11	0.18
R^2	1	0.09	0.10	0.02	0.09	0.19
	3	0.20	0.25	0.06	0.25	0.45
	5	0.25	0.36	0.09	0.37	0.61

Panel B: Predicting dividend growth

$$\sum_{j=1}^{L} \Delta d_{t+j} = \alpha + \beta(p_t - d_t) + \epsilon_{t+L}$$

	L	Data 1948Q1-2010Q4	Model mean	Model 5%	Model 50%	Model 95%
$\hat{\beta}$	1	-0.03	-0.01	-0.05	-0.01	0.02
	3	0.00	0.06	-0.01	0.05	0.16
	5	0.04	0.11	0.02	0.10	0.27
SE($\hat{\beta}$)	1	0.03	0.03	0.01	0.02	0.04
	3	0.07	0.05	0.03	0.04	0.08
	5	0.10	0.06	0.03	0.06	0.10
R^2	1	0.01	0.00	0.00	0.00	0.02
	3	0.00	0.04	0.00	0.02	0.12
	5	0.01	0.09	0.00	0.07	0.26

Figure 7: Return and Dividend Predictability Scatterplots Expected returns and expected dividend growth are computed from the model's laws of motion. Figures are annualized.

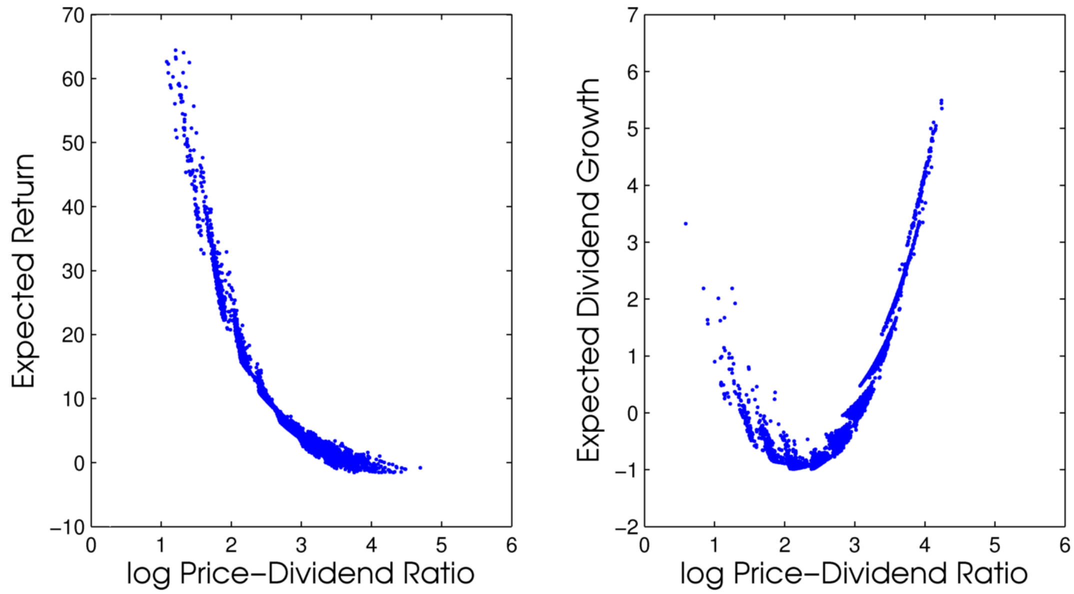

and asset valuations. Expected returns have a much steeper relationship with the log price-dividend ratio when the price-dividend ratio is particularly low. As the log price-dividend ratio increases, the relationship flattens, and expected returns never become very negative. The skewness is seen in the concentration of the dots from the scatterplot. The mean annual return is about 7.5%, indicating that the scatterplot is extremely concencentrated in the lower right portion of the expected return panel. The sparse dots in the upper left of the panel represent rare moments of crisis, when asset valuations are extremely low and expected returns are extremely high.

The results of this section show that the model does a good job of capturing the nature of fluctuations in asset prices. A boom in the stock market does not reflect an increase in corporate profits and dividends. Rather, the boom reflects a decline in expected returns. The model is an efficient markets model, and so the decline in expected returns reflects a decline in the equity risk premium. In particular, the decline in the risk premium is due to low consumption volatility which is driven by prudential uncertainty. This stylized fact is difficult to capture even in endowment economies (Bansal, Kiku, and Yaron [2009]). Capturing this relationship in a model with a non-trivial real side which must maintain consistency with business cycle moments is an even more challenging task (Favilukis, Ludvigson, and Van Nieuwerburgh [2010]).

3.4.2. Mimicking the Campbell-Cochrane $\lambda(s_t)$

The model produces data-like time-varying risk premiums because the prudential uncertainty mechanism generates time-varying consumption volatility which quantitatively mimics the Campbell and Cochrane [1999] sensitivity function $\lambda(s_t)$.

To see this, it helps to examine the log-normal approximation of the Hansen-Jagannathan bound (maximum Sharpe ratio):

$$\max_{\{\text{all assets}\}} \left\{ \frac{\mathbb{E}_t(R_{t+1} - R^f_{t+1})}{\sigma_t(R_{t+1})} \right\} \approx \gamma(\lambda + 1)\sigma_t(\Delta c_{t+1}) \tag{18}$$

The conditional maximum Sharpe ratio is the conditional volatility of consumption growth, multiplied by some preference parameters (γ is the utility curvature, and λ is the habit sensitivity). The preference parameters are constant. Thus the driver of time-varying risk premiums is the sole term on the right hand side which can vary over time: the conditional volatility of consumption. This conditional volatility is driven by a single, homoskedastic productivity shock. The lack of an external source for time-varying volatility indicates that prudential uncertainty is a key driver of variation in the conditional volatility of consumption growth.

Now compare the above expression to the Hansen-Jagannathan bound in Campbell and Cochrane [1999]:

$$\max_{\{\text{all assets}\}} \left\{ \frac{\mathbb{E}_t(R_{t+1} - R^f_{t+1})}{\sigma_t(R_{t+1})} \right\} \approx \gamma[\lambda(s_t) + 1]\sigma(\Delta c_{t+1}) \tag{19}$$

where

$$\lambda(s_t) = \begin{cases} \frac{1}{\bar{S}}\sqrt{1 - 2(s_t - \bar{s})} - 1 & \text{if} \quad s_t \le \bar{s} + \frac{1}{2}(1 - \bar{S}^2) \\ 0 & \text{if} \quad s_t \ge \bar{s} + \frac{1}{2}(1 - \bar{S}^2) \end{cases}$$

This expression is very similar in structure to the bound in my model (equation (18)), but the time dependence is reversed. $\sigma(\Delta c_{t+1})$ is assumed to be constant, but now the preference parameter $\lambda(s_t)$ changes over time. Thus, in Campbell-Cochrane, variation in risk premiums is driven by variation in the sensitivity of habit $\lambda(s_t)$.

By equating (19) and (18), I can create an 'implied consumption volatility' generated by the

Campbell-Cochrane $\lambda(s_t)$. Explicitly,

$$\text{Implied Consumption Volatility} = \left(\frac{\lambda(s_t)+1}{\lambda+1}\right)\sigma(\Delta c_{t+1}) \tag{20}$$

Implied consumption volatility is simply $\lambda(s_t)$, re-scaled by its steady state value λ, and multiplied by the unconditional volatility $\sigma(\Delta c_{t+1})$. Figure 8 compares the consumption volatility generated by my model to the Campbell-Cochrane implied consumption volatility. The blue dots are a scatterplot of consumption volatility against surplus consumption generated by simulated data. Consumption volatility declines in surplus consumption as a result of prudential uncertainty. On the same axis, I plot the Campbell-Cochrane implied consumption volatility with a green dashed line. Implied consumption volatility also declines in surplus consumption. Moreover, for much of the plot, implied consumption volatility runs right through the middle of the cloud of dots representing consumption volatility. The two channels are not only qualitatively similar, but they are quantitatively similar. This result shows that prudential uncertainty can be interpreted as a method for endogenizing the Campbell-Cochrane mechanism.

The two plots deviate for high levels of surplus consumption because Campbell and Cochrane's habit specification puts a maximum on surplus consumption. I impose no such restriction in my model. In untabulated results, I find that imposing this restriction improves the quantitative connection between these two channels.

Figures 9 and 10 provide a further comparison between my model and Campbell and Cochrane [1999]. Figure 9 shows scatterplots of asset prices against surplus consumption from simulated data. The left panel shows that the price-dividend ratio increases in surplus consumption, and the right panel shows that the expected equity premium decreases in surplus consumption. As in Campbell-Cochrane, times of low surplus consumption are risky times with high equity premiums and low asset valuations. Surplus consumption can be interpreted as a 'recession state variable,' that is, an indicator of bad times.

However, Figure 10 shows that consumption volatility is a key driver of asset price moments. The figure shows scatterplots of asset prices against consumption volatility from simulated data. The left panel shows that the price-dividend ratio has a negative relationship with consumption volatility, and the right panel shows that the expected equity premium has a very tight positive relationship with consumption volatility. This conditional relationship between volatility and asset prices is completely absent in the Campbell-Cochrane model, which assumes that consumption volatility is a constant.

Figure 8: Campbell-Cochrane $\lambda(s_t)$ vs time-varying volatility. Blue dots represent consumption volatility from model simulations. The green dashed line represents Campbell-Cochrane $\lambda(s_t)$ implied consumption volatility (equation (20)). Figures are annualized.

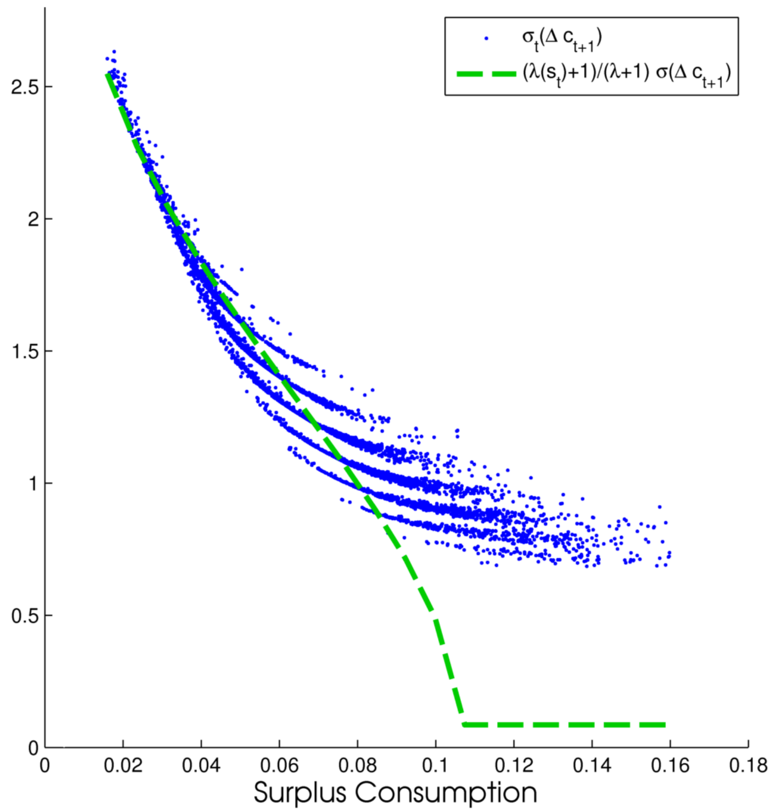

Figure 9: Surplus consumption and asset prices. Equity premium is computed from the model's laws of motion. Figures are annualized.

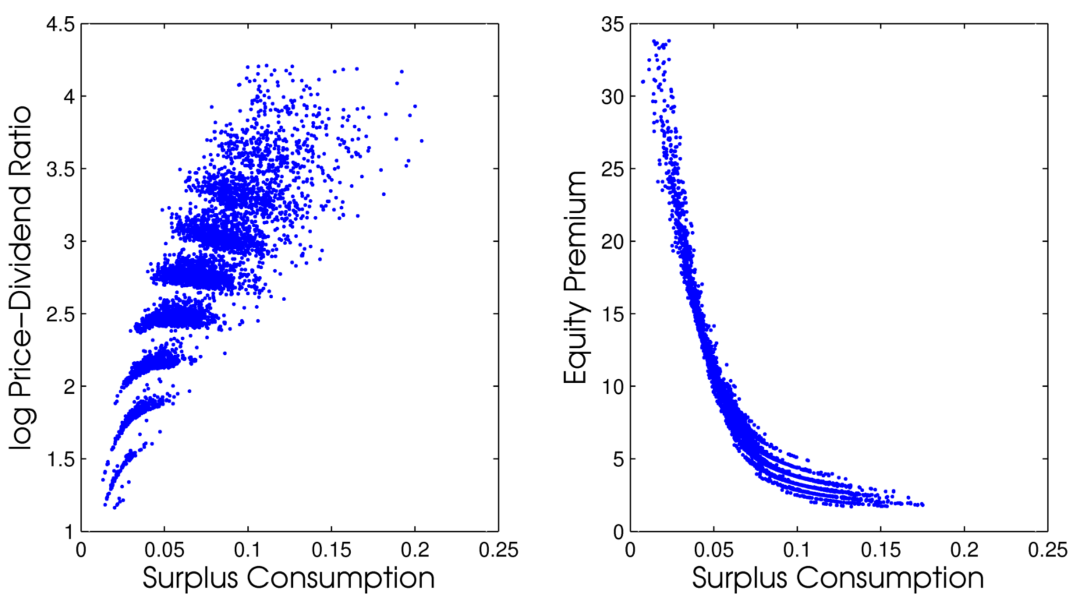

Figure 10: Consumption volatility and asset prices. Equity premium is computed from the model's laws of motion. Figures are annualized.

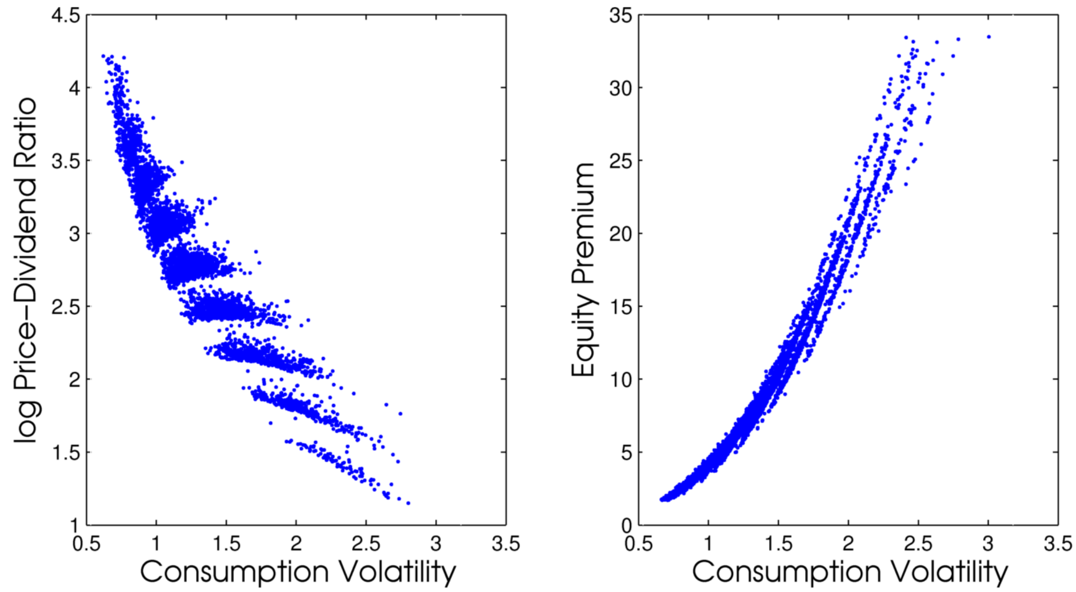

Indeed, equation (18) shows that it is variation in consumption volatility which drives the key asset pricing results in my model.

3.4.3. Time-variation in Consumption Volatility

This section shows that the model drives time-varying risk premiums with a reasonable amount of time-variation in consumption volatility. It also examines the quantitative link between asset prices and consumption volatility.

I use a number of proxies for conditional consumption volatility. All proxies are constructed by first fitting an AR(1) model to log consumption growth to remove an expected growth component:

$$\Delta c_{t+1} = b_0 + b_1 \Delta c_t + \epsilon^c_{t+1} \tag{21}$$

I then either estimate GARCH-type models on the residual ϵ^c_{t+1} or use the mean absolute residual as a non-parametric measure of conditional volatility. This procedure follows Bansal, Khatchatrian, and Yaron [2005] and Beeler and Campbell [2009].

Table 3 shows univariate measures of time-variation in consumption volatility. Panel A shows coefficients from a GJR-GARCH(1,1,1) model. I choose to focus on GJR-GARCH, rather than the simpler GARCH model, because the Glosten, Jagannathan, and Runkle [1993] GARCH variant allows negative shocks to have a different effect on volatility than positive shocks, as predicted by the model.[7] GJR-GARCH assumes that the consumption growth residual from the AR(1) model (equation (21)) follows

$$\begin{aligned} \epsilon^c_{t+1} &= \sigma_{c,t}\eta_{t+1} \\ \sigma^2_{c,t+1} &= \omega + \alpha\eta^2_t + \gamma I(\eta_t < 0)\eta^2_t + \beta\sigma^2_{c,t} \end{aligned} \tag{22}$$

where η_{t+1} is a standard normal shock. Tomorrow's conditional volatility $\sigma^2_{c,t+1}$ has a symmetric component $\alpha\eta^2_t$, an asymmetric component $\gamma I(\eta_t < 0)\eta^2_t$, and an autoregressive component $\beta\sigma^2_{c,t}$. Panel A shows that both in the model and the data, the $\hat{\gamma}$ coefficients are much larger than than $\hat{\alpha}$ coefficients, indicating that negative shocks have a much stronger effect on consumption volatility than positive shocks. This asymmetry is consistent with the prudential uncertainty mechanism. Consumption volatility rises in bad times because bad times create uncertainty about future consumption policy. In terms of magnitudes, the amount of time-varying volatility in the model is typically less than that in the data. This difference should not be surprising, since the model abstracts from other shifts in volatility such as Bloom [2009]'s technological uncertainty shocks, or Pastor and Veronesi [2012]'s government policy uncertainty shocks. The $\hat{\beta}$ coefficients show that the model produces very persistent consumption volatility. In terms of magnitudes, the model coefficient of 0.71 comes close to reaching the data value of 0.77.

For robustness, Panel B shows a non-parametric measure of time-varying consumption volatility. This measure is the autocorrelation of the absolute value of the residuals from the AR(1) model of consumption growth. Under the hypothesis of constant consumption volatility, this autocorrelation should be zero for all time-horizons. The model generates a mildly positive autocorrelation for all time horizons up to 20 quarters. This autocorrelation is much less than the point estimate for the autocorrelation in the data, but the point estimates are well within the 95% confidence intervals implied by model simulations. This result is due to the large amount of sampling uncertainty in these estimates. The huge confidence intervals speak to the difficulty in detecting persistent variation in a second moment of a relatively smooth time series using only 50 years of data. Overall, Table 3 shows that the model generates significant time-varying risk premiums with a reasonable amount of time-varying consumption volatility.

A number of papers have shown that conditioning on the price-dividend ratio helps pick up

[7] In untabulated results, I find that the GJR-GARCH(1,1,1) model does a much better job of capturing the model's true conditional volatility compared to GARCH(1,1) and the non-parametric specifications.

Table 3: Conditional Consumption Volatility Moments

Data and figures are quarterly. This table shows measures of heteroskedasticity from an AR(1) model of log consumption growth:

$$\Delta c_{t+1} = b_0 + b_1 \Delta c_t + \epsilon^c_{t+1}$$

The model columns show means and percentiles from 1000 simulations of the same length as the empirical sample. Further details are found in the Appendix.

Panel A: GJR-GARCH model

$$\epsilon^c_{t+1} = \sigma_{c,t}\eta_{t+1}$$
$$\sigma^2_{c,t+1} = \omega + \alpha\eta_t^2 + \gamma I(\eta_t < 0)\eta_t^2 + \beta\sigma^2_{c,t}$$

	Data	Model			
	1948Q1-2010Q4	mean	5%	50%	95%
$\hat{\omega}$	1.95E-06	1.04E-05	2.35E-07	2.62E-06	4.50E-05
$\hat{\alpha}$	0.06	0.01	0.00	0.00	0.10
$\hat{\gamma}$	0.15	0.07	-0.07	0.05	0.25
$\hat{\beta}$	0.77	0.71	0.00	0.90	0.98

Panel B: Autocorrelation of absolute residuals

$$\rho(|\hat{\epsilon}^c_t|, |\hat{\epsilon}^c_{t+L}|)$$

	Data	Model			
L	1948Q1-2010Q4	mean	5%	50%	95%
4	0.075	0.010	-0.089	0.002	0.141
8	0.072	0.006	-0.097	-0.001	0.131
12	-0.029	0.003	-0.100	-0.005	0.135
16	0.019	0.008	-0.093	0.000	0.138
20	0.037	0.001	-0.099	-0.004	0.128

time-variation in consumption volatility. These papers also show that the price-dividend ratio has a robust negative relationship with consumption volatility (Kandel and Stambaugh [1990], Bansal, Khatchatrian, and Yaron [2005], Bansal, Kiku, and Yaron [2009]). Table 4 examines the model's ability to replicate this fact.

Table 4 shows regressions of various proxies for consumption volatility on the price-dividend ratio. Panel A uses GJR-GARCH and GARCH estimates to proxy for consumption volatility. For both proxies, the model-produced coefficient on the price-dividend ratio is negative and statistically significant, as in the data. A high price-dividend ratio indicates a safe time of low consumption volatility. R^2's for the model are large and close to the data value of roughly 0.30. In terms of magnitudes, the table shows that the prudential uncertainty channel gets us about half-way to the data. Using the GJR-GARCH proxy, the coefficient on the price-dividend ratio is -0.23 and is roughly half of the data value of -0.38.

For robustness, Panel B uses a non-parametric measure of consumption volatility. This measure is the sum of the absolute value of residuals from the AR(1) model of consumption growth. I take logs to reduce the sensitivity to outliars, as in Bansal, Kiku, and Yaron [2009] and Beeler and Campbell [2009]. The slopes on the price-dividend ratio are all negative, as in the data. Consistent with the GJR-GARCH results, the model coefficients are roughly half of those from the data. Unfortunately, the model-produced standard errors are much larger than those from the data. This lack of precision in the regression estimates is reflected in the R^2's, which are much smaller than the data values. This dimension of asset prices and the real economy is very challenging to match. The endowment economy of Bansal, Kiku, and Yaron [2009], where consumption volatility is exogenously specified, generates R^2's similar to my model's results.

When interpreting Tables 4 and 3 one should keep in mind that both the model and the calibration exercise are intended to demonstrate *only* the prudential uncertainty mechanism. The inclusion of other, possibly technological, mechanisms could improve the fit displayed in Tables 4 and 3, but is beyond the scope of this paper.

3.5. Unconditional Asset Price and Business Cycle Moments

As is well-known, describing basic asset price and business cycle moments in a single model is a challenging task. Tables 5 and 6 show that the model has promise as a unified description of both. They show that the model captures a substantial list of unconditional asset price and business cycle moments beyond the five that it is calibrated to match.

Table 4: Regressions of Consumption Volatility on the Price-Dividend Ratio

Figures are annual. This table shows regressions of the form

$$\text{cvol}_t = \alpha + \beta(p_t - d_t) + \epsilon_t$$

Where cvol_t is a proxy for conditional consumption volatility, p_t is the log asset price, and d_t is the log dividend. To generate cvol_t, first, an AR(1) model is run on log consumption growth. Panel A estimates either a GARCH(1,1) or GJR-GARCH(1,1,1) model on the residuals, and then takes the log of the estimated consumption volatility. Panel B uses a non-parametric measure: $\text{cvol}_t(L) \equiv \log\left(\sum_{j=1}^{4L} |\epsilon^c_{t+j}|\right)$, where ϵ^c_{t+j} is the residual from the AR(1) model. Consumption data is quarterly and price-dividend ratio data is annual, which results in some abuse of notation. The model columns show means and percentiles from 1000 simulations of the same length as the empirical sample. Further details are found in the Appendix.

Panel A: GARCH-variant estimates of consumption volatility

		Data	Model			
		1948Q1-2010Q4	mean	5%	50%	95%
GJR-GARCH(1,1,1)	$\hat{\beta}$	-0.38	-0.23	-0.62	-0.19	0.02
	SE($\hat{\beta}$)	0.12	0.06	0.00	0.05	0.15
	R^2	0.32	0.31	-0.02	0.23	0.83
GARCH(1,1)	$\hat{\beta}$	-0.40	-0.15	-0.47	-0.10	0.02
	SE($\hat{\beta}$)	0.11	0.06	0.00	0.05	0.15
	R^2	0.38	0.24	-0.02	0.15	0.74

Panel B: Non-parametric estimates of consumption volatility
$\text{cvol}_t(L) \equiv \log\left(\sum_{j=1}^{4L} |\epsilon^c_{t+j}|\right)$

		Data	Model			
	L	1948Q1-2010Q4	mean	5%	50%	95%
$\hat{\beta}$	1	-0.63	-0.29	-1.21	-0.26	0.60
	3	-0.56	-0.30	-1.02	-0.31	0.45
	5	-0.51	-0.27	-0.88	-0.28	0.37
SE($\hat{\beta}$)	1	0.16	0.39	0.16	0.35	0.75
	3	0.12	0.25	0.08	0.21	0.57
	5	0.08	0.18	0.05	0.14	0.42
R^2	1	0.20	0.03	0.00	0.02	0.11
	3	0.34	0.07	0.00	0.04	0.23
	5	0.45	0.11	0.00	0.06	0.37

Table 5 shows unconditional moments of asset prices. The model is calibrated to fit three of these moments: the mean risk-free rate, the autocorrelation of the risk-free rate, and the mean market Sharpe ratio. The table shows that the model can fit these moments quite nicely. Beyond the calibration targets, the model also matches the basic moments of the equity premium. The model's mean equity premium of 6.62% is rather close to the data value of 6.00%. The calibration targets the mean Sharpe ratio, but this does not imply a large equity premium unless the model can generate enough volatility. A number of models generate a large Sharpe ratio but a small equity premium (e.g. Kaltenbrunner and Lochstoer [2010], Gourio [2010], Dew-Becker [2011]). Table 5 shows that the model does generate plenty of equity volatility. The model produces an equity premium volatility of 17.68%, which is extremely close to the data value of 17.47%. This large premium and volatility are all computed without adjusting for financial leverage.

Habit models that match the mean and volatility of the equity premium often do so at the expense of an excessively volatile risk-free rate (i.e. 11% in Jermann [1998] and 25% in Boldrin, Christiano, and Fisher [1999]). Table 5 shows that this model makes a substantial improvement in this dimension by producing a risk-free rate volatility of 3.34%, which is close to the data value of 1.82%. This low risk-free rate volatility is reflected in a reasonable term premium. The model produces a mean excess return on 10-year bonds of 2.03%, which is less than one third of the equity premium. This result indicates that the model's equity premium is distinct from the term premium, as in the data.

Prudential uncertainty is critical for the model's ability to produce a smooth risk-free rate. To see this, it helps to examine a log-normal approximation of the risk-free rate:[8]

$$\begin{aligned} r^f_{t+1} \approx & -\log\beta + \gamma(\lambda+1)\mathbb{E}_t(\Delta c_{t+1}) - \gamma(1-\rho_s)(s_t - \bar{s}) \\ & - (1/2)\gamma^2(\lambda+1)^2 \mathrm{Var}_t(\Delta c_{t+1}) \end{aligned} \tag{23}$$

The 2nd and 3rd terms are due to the elasticity of intertemporal substitution and tend to have a countercyclical effect on the risk-free rate. Intuitively, in bad times, investors want to borrow from the future in order to consume today. This motive pushes down the price on the risk-free bond and pushes up the risk-free rate. Habit models make marginal utility particularly volatile and increase the cyclicality of this motive. The 4th term represents a precautionary savings, or prudence effect. Prudential uncertainty causes this term to be procyclical, by making consumption volatility

[8] If m_{t+1} is normal,

$$r^f_{t+1} = -\log\left[\mathbb{E}_t(e^{m_{t+1}})\right] = -\mathbb{E}_t(m_{t+1}) - \frac{1}{2}\mathrm{Var}_t(m_{t+1})$$

Then just plug in the log SDF $m_{t+1} = \log\beta - \gamma\Delta s_{t+1} - \gamma\Delta c_{t+1}$ and habit process $\Delta s_{t+1} = -(1-\rho_s)(s_t - \bar{s}) + \lambda_t \Delta c_{t+1}$.

Table 5: Unconditional Asset Price Moments

All figures are annual. No adjustments are made to account for financial leverage. r, p, and d are the logs of returns, prices, and dividends from the CRSP value-weighted index. r_f is the log-return on 30 day T-bills. Capital letters represent levels rather than logs. Δc is the first-difference of log consumption growth. $\mathbb{E}$, σ, and AC1 represent the sample mean, standard deviation and first-order autocorrelation. Data moments marked with a † are taken from Beeler and Campbell [2009] and correspond to 1947Q2-2008Q4. The model columns show means and percentiles from 1000 simulations of the same length as the empirical sample. Further details are found in the Appendix.

	Data	Model			
	1948Q1-2010Q4	mean	5%	50%	95%
Calibrated Moments					
$\mathbb{E}[r_f]$ (%)	0.89	0.88	-3.61	0.48	6.79
$\mathbb{E}[R-R_f]/\sigma[R]$	0.44	0.43	0.27	0.42	0.61
$\text{AC1}[r_f]$	$0.84^\dagger$	0.85	0.65	0.87	0.99
Untargeted Moments					
$\mathbb{E}[r-r_f]$ (%)	6.00	6.62	4.51	6.62	8.63
$\sigma[r-rf]$ (%)	17.47	17.68	11.84	17.13	25.39
$\text{AC1}[r-r_f]$	-0.06	-0.07	-0.30	-0.07	0.16
$\sigma[r_f]$ (%)	$1.82^\dagger$	3.34	0.70	2.43	9.06
$\mathbb{E}[p-d]$	3.44	2.72	2.21	2.72	3.28
$\sigma[p-d]$	0.43	0.46	0.27	0.45	0.73
$\text{AC1}[p-d]$	0.95	0.90	0.79	0.91	0.97
$\mathbb{E}[r_{10yr}-r_f]$ (%)		2.03	1.03	1.99	3.22
$\sigma[r_{10yr}-r_f]$ (%)		1.42	0.52	1.30	2.78
$\text{AC1}[r_{10yr}-r_f]$		0.85	0.71	0.87	0.95

Figure 11: Decomposition of the Risk-free rate. Computed from model's laws of motion. Risk-free rate is in annualized %. Capital in all panels is fixed at the mean capital stock.

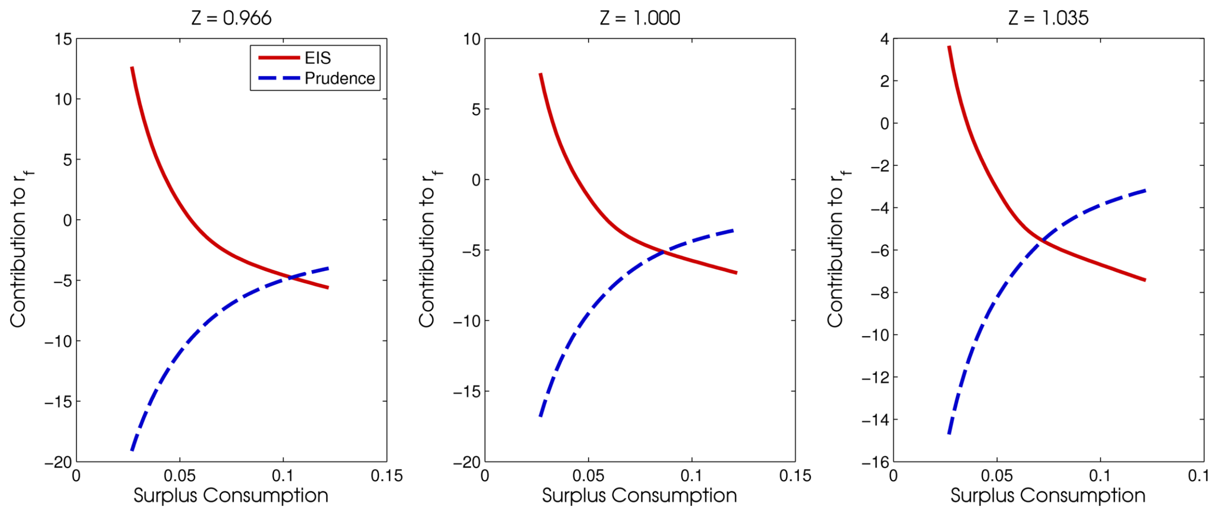

countercyclical. Intuitively, in bad times, investors want to save in order to protect themselves against heightened uncertainty. This motive pushes up the price on the risk-free bond and pushes down the risk-free rate.

Figure 11 plots the decomposition from equation (23) against surplus consumption for various levels of productivity Z. The solid red 'EIS' lines represent the countercyclical elasticity of intertemporal substitution effects. The dashed blue 'prudence' lines represent the procyclical precautionary savings effects. The EIS effects increase in surplus consumption while the prudence effects decrease in surplus consumption. Quantitatively, the two effects balance each other out quite nicely.

This smoothing effect on the risk-free rate is closely related to prudential uncertainty's ability to mimic the Campbell-Cochrane $\lambda(s_t)$ function. As in the Campbell-Cochrane model, low surplus consumption indicates a bad time when investors want to borrow in order to consume today. Bad times are also times of high precautionary concerns, however, and these concerns keep investors from borrowing and smooth the risk-free rate. Precautionary concerns in Campbell-Cochrane are driven by countercyclical preferences, a countercyclical $\lambda(s_t)$ function. In my model, they're driven by a countercyclical consumption volatility, which in turn is the result of prudential uncertainty.

Beyond the first two moments of the equity premium and the risk-free rate, the model also matches key features of the price-dividend ratio, the persistence of asset prices, and second moments of aggregate quantities. The model generates a volatility of the log price-dividend ratio of 0.46, which is close to the data value of 0.43. This moment has been difficult to capture even in endowment economies (Bansal, Kiku, and Yaron [2009]). Regarding persistence, the model replicates the mild negative autocorrelation of the equity premium that is seen in the data, and reproduces the high persistence of the price-dividend ratio and risk-free rate. The model's annual autocorrelation of the price-dividend ratio is 0.90, which is close to the data value of 0.95. The model's annual autocorrelation of the risk-free rate is 0.85, which is very close to the data value of 0.84, as intended by the calibration.

Table 6 shows business cycle moments. The table examines Hodrick-Prescott filtered quantities in order to focus on business-cycle frequency fluctuations. The model hits the volatility of output and relative volatility of consumption to output seen in the data, as intended by the calibration. Investment is more volatile than output, as in the data. The model-generated relative volatility of investment of 3.56 is higher than the data value of 2.26, but this excessive investment volatility is a common problem with business cycle models.[9] Like the data, the model produces strong co-movement between output, consumption, and investment. Output, consumption, and investment are highly persistent, and are nearly as persistent as the data. First-differenced log consumption has a low annual volatility of 1.13%, which is slightly lower than the data value of 1.28%. Lastly, the table also shows that average adjustment costs are small at less than 1% of output.

3.6. Comparative Statics

The quantitative results come from calibrated parameter values, but there is some uncertainty regarding the parameter values for non-traditional modeling elements. This section investigates the effect of changing these parameter values.

Table 7 shows key moments from these comparative statics exercises. Each column examines moments generated by models where only one of the parameter values is changed from the calibration of Table 1. Three different parameter changes are examined: lower persistence of habit, weaker steady state habit, and lower capital adjustment costs. The magnitude of the perturbations are chosen to be the smallest change that produces a clearly recognizable deviation from the calibrated results. This approach helps isolate the direct effect of changing a parameter value from its interaction with other elements of the model.

Lower persistence of habit The third column of Table 7 examines a model where the persistence of habit ρ_s is lowered from the calibrated value of 0.966 to 0.800. This parameter is important because the strong persistence of habit is one way in which this model deviates from previous asset pricing

[9] The data moment for investment volatility is low compared to most measurements because it includes government investment, as advocated by Cooley and Prescott [1995].

Table 6: Basic Business Cycle Moments

Figures are quarterly, except where noted. The subscript hp indicates that the moment is calculated from logged and HP-filtered data with a smoothing parameter of 1600. Δ indicates first-differences. The model columns show means and percentiles from 1000 simulations of the same length as the empirical sample. Consumption is real per capita non-durable goods and services consumption. Further details are found in the Appendix.

	Data	Model			
	1948Q1-2010Q4	mean	5%	50%	95%
Calibrated Moments					
$\sigma(y_{hp})(\%)$	1.73	1.74	1.33	1.72	2.23
$\sigma(c_{hp})/\sigma(y_{hp})$	0.48	0.49	0.41	0.49	0.57
Untargeted Moments					
$\sigma(i_{hp})/\sigma(y_{hp})$	2.26	3.56	3.24	3.57	3.85
$\rho(y_{hp}, c_{hp})$	0.80	0.99	0.98	0.99	1.00
$\rho(y_{hp}, i_{hp})$	0.89	1.00	1.00	1.00	1.00
AC1(y_{hp})	0.85	0.71	0.62	0.71	0.78
AC1(c_{hp})	0.86	0.70	0.61	0.71	0.78
AC1(i_{hp})	0.85	0.70	0.62	0.71	0.78
$\sigma[\Delta c]$ (annual %)	1.28	1.13	0.78	1.10	1.54
$AC1[\Delta c]$ (annual)	0.45	0.24	0.02	0.24	0.47
$\mathbb{E}[\text{Adjustment Cost}/Y](\%)$		0.60	0.26	0.55	1.14
$\mathbb{E}[\text{Adjustment Cost}/I](\%)$		3.75	1.44	3.26	7.41

models with habit and production. The high persistence means that habit today depends on a very long history of consumption. In contrast, habit in Jermann [1998] and Boldrin, Christiano, and Fisher [2001] depends only on the last quarter's consumption.

The table shows that the most direct effect of lowering the persistence of habit is that the model no longer matches the persistence of asset prices. The calibrated value is chosen to fit the annual autocorrelation of the risk-free rate of 0.85. With the lower persistence of habit, the autocorrelation of the risk-free rate falls to 0.66. The same drop in persistence can be seen in the persistence of the price-dividend ratio, which falls from 0.90 to 0.77. Intuitively, ρ_s captures the persistence of preferences, which is reflected in the persistence of asset prices.

The lower persistence of habit worsens the model fit in many other areas. Notably, the volatility of the risk free rate more than doubles from the calibrated value of 3.34% to 8.96%. Because the risk-free rate is a convex function of surplus consumption, this increases the mean risk-free rate from the calibrated value of 0.88% to 3.71%. Figure 12 illustrates why. The figure plots the risk-free rate decomposition of equation (23) for the low persistence of habit model. The figure shows that the EIS effect on the risk-free rate is still countercyclical and the prudence effect is still procyclical. However, the two channels no longer cancel each other out quantitatively. The EIS effect becomes more countercyclical and the prudence effect becomes less procyclical as compared to the calibrated model (Figure 11). Both of these effects can be traced to the lower persistence of habit. A low persistence of habit means that habit will strongly mean revert tomorrow. Regarding intertemporal substitution, this means that, in bad times, there is a pronounced desire to borrow in order to consume today. This boosts the countercyclicality of the EIS effect. Regarding prudence, this means that prudential uncertainty is weakened. Habit will strongly mean revert tomorrow, meaning tomorrow the household has less need for prudence.

This weakening of prudential uncertainty is reflected in a reduction in the amount of time-varying risk premiums. The third column of Table 7 shows that the amount of time-variation in risk premiums drops significantly. R^2's from regressions of future returns on the log price-dividend ratio drop by more than half. At the 5-year horizon, the R^2 drops from the calibrated value of 0.35 to the 0.15.

Weaker steady state habit The fourth column of Table 7 raises steady state surplus consumption $\bar{S}$ from the calibrated value of 0.065 to 0.12. This comparative static brings the model closer to CRRA utility. Roughly speaking, $\bar{S} = 1.00$ is an economy with no habit, so this model is still very far from the standard model.

Table 7: Comparative Statics

Figures are annual, excepting $\sigma(c_{hp})/\sigma(y_{hp})$, which is quarterly. 'Calibrated' represents parameter values from Table 1. All other columns use the calibrated values but with one parameter changed. 'Lower persistence of habit' sets $\rho_s = 0.80$. 'Weaker steady state habit' sets $\bar{S} = 0.12$. 'Lower adjustment costs' sets $\phi = 40$. Data moments marked with a † are taken from Beeler and Campbell [2009] and correspond to 1947Q2-2008Q4. Details of the data are found in the Appendix.

		Data 1948Q1-2010Q4	Calibrated	Lower persistence of habit	Weaker steady state habit	Lower adjustment costs
Calibrated Moments						
$\mathbb{E}[r_f]$ (%)		0.89	0.88	3.71	6.27	5.44
$\mathbb{E}[R - R_f]/\sigma[R]$		0.44	0.43	0.35	0.29	0.31
AC1$[r_f]$		$0.84^\dagger$	0.85	0.66	0.87	0.94
$\sigma(c_{hp})/\sigma(y_{hp})$		0.48	0.49	0.44	0.64	0.37
Untargeted Moments						
$\mathbb{E}[r - r_f]$ (%)		6.00	6.62	6.30	3.36	2.14
$\sigma[r - rf]$ (%)		17.47	17.68	23.40	14.39	7.36
AC1$[r - r_f]$		-0.06	-0.07	-0.11	-0.04	-0.03
$\sigma[r_f]$ (%)		2.25	3.34	8.96	4.58	2.23
$\mathbb{E}[p - d]$		3.44	2.72	2.32	2.36	2.66
$\sigma[p - d]$		0.43	0.46	0.36	0.32	0.37
AC1$[p - d]$		0.95	0.90	0.77	0.89	0.93
R^2 from	1-year	0.09	0.10	0.06	0.04	0.05
forecasting r	3-year	0.20	0.25	0.12	0.11	0.15
with $p - d$	5-year	0.25	0.36	0.15	0.16	0.23

Figure 12: Decomposition of the Risk-free rate: low persistence of habit. $\rho_s = 0.800$. All other parameter values are from the calibration in Table 1. Risk-free rate is in annualized %. Capital in all panels is fixed at the mean capital stock.

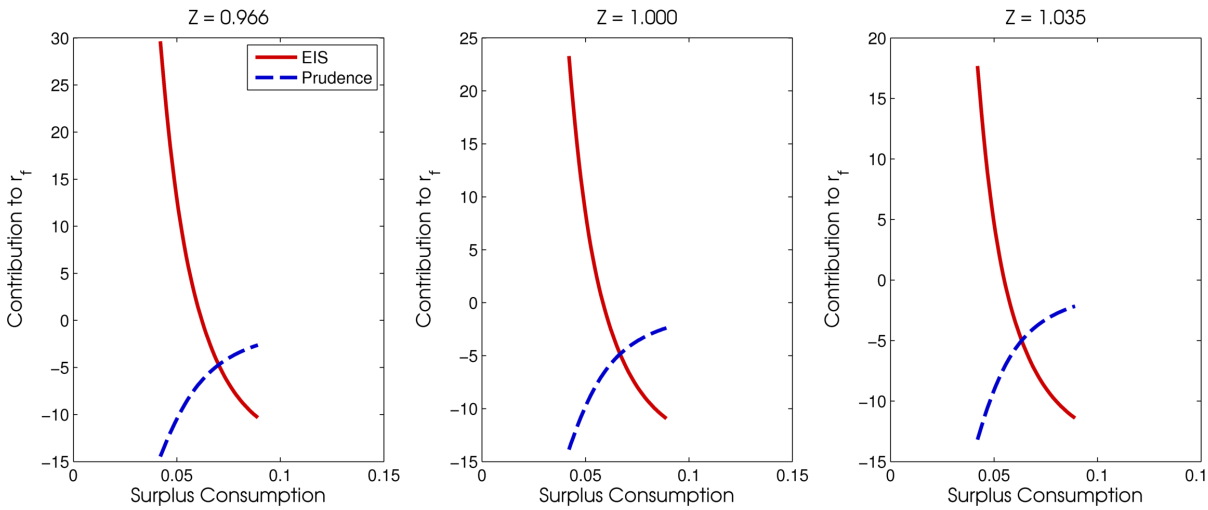

Weakening steady state habit lowers the overall risk aversion of the household, and thus has the direct effect of lowering the Sharpe ratio from the calibrated value of 0.43 to 0.29. This reduced Sharpe ratio has the obvious effect of reducing the equity premium from the calibrated value of 6.62% to 3.36%. The Sharpe ratio does not drop as much as one might expect, however, since the volatility of consumption increases. The relative volatility of consumption to output increases from the calibrated value of 0.49 to 0.64. This increase in consumption volatility is due to a weaker general equilibrium consumption smoothing effect. External habit induces strong consumption smoothing (Lettau and Uhlig [2000]). This is a general equilibrium effect which runs through asset prices and firm investment. External habit results in a low elasticity of intertemporal substitution and this preference for smooth consumption is reflected in asset prices. The equilibrium asset prices, in turn, reward the firm for providing smooth dividends. Although the result is similar, this general equilibrium effect is distinct from the consumption smoothing effects of internal habit models (for example Sundaresan [1989]).

Weakening steady state habit has the additional effect of weakening prudential effects. With weaker habit, there is less need for precautionary savings, and the risk-free rate jumps from its calibrated value of 0.88% to 6.27%. This weakening of prudential effects can be seen in the return forecasting regressions too. The R^2's from regressions of future returns on the price-dividend ratio all drop by more than half.

Lower adjustment costs The final column of Table 7 examines the effect of lowering capital adjustment costs. The adjustment cost parameter ϕ is lowered from the calibrated value of 103 to

40. This comparative static helps compare my model with Lettau and Uhlig [2000], who also examine external habit in a production economy. My model deviates from theirs in two important ways. The first is that I include capital adjustment costs, while Lettau and Uhlig assume costless adjustment. The second is that I use a non-linear, global solution method, while Lettau and Uhlig linearize the model around the non-stochastic steady state. This comparative static shows that capital adjustment costs make a large impact on the model results and that lowering the costs brings my model closer to Lettau and Uhlig [2000].

The table shows that lowering adjustment costs reduces the relative volatility of consumption to output falls from the calibrated value of 0.49 to 0.37. The intuition for this result is as follows. External habit encourages the firm to provide smooth dividends through general equilibrium effects. To smooth dividends, the firm invests in response to positive shocks and disinvests in response to negative shocks. Capital adjustment costs restrain this motive and, when calibrated, the model can match the data value on consumption volatility. Reducing capital adjustment costs, then, decreases consumption volatility.

This decrease in consumption risk is reflected in a lower Sharpe ratio, which falls from 0.44 to 0.31 and a lower equity premium which falls from 6.62% to 2.14%. With lower consumption risk, the precautionary motive is weakened and the risk-free rate rises from the calibrated value of 0.88% to 5.44%. This decrease in the precautionary motive is also seen in lower R^2's from return forecasting regressions.

4. Conclusion

This paper presents a new "prudential uncertainty" channel for generating time-varying expected returns. Prudent savings dynamics lead to countercyclical consumption volatility. Because equity exposes investors to countercyclical consumption risk, investors demand countercyclical equity returns. This channel is amplified by external habit which makes the household feel poorer and act more prudently. With external habit as an amplifier, this channel can quantitatively explain time-varying risk premiums, as embodied in return and dividend predictability regressions.

The quantitative model shows promise as a "fundamental" model of asset prices, that is, a model of the real economy which helps us understand asset prices. The model matches a long list of asset pricing facts while producing business cycle moments which match the data. One interpretation of this result is that the prudential uncertainty channel behaves very much like the Campbell and Cochrane [1999] sensitivity function $\lambda(s_t)$. As a result, this channel allows many of the insights of Campbell and Cochrane [1999] to be extended into a production economy.

The model puts an emphasis on the role of production in asset pricing theory. In production economies, the investor's preferences and beliefs affect not only affect asset prices, but macroeconomic quantities as well. As a result, production economies break the equivalence between distorted beliefs and state-dependent preferences that exists in endowment economies (Cochrane [2011]). This model shows that an external habit model with production provides a good fit for numerous asset price and macroeconomic quantity moments. Whether an "equivalent" model with distorted beliefs can generate the same fit is an open question.

The quantitative model uses a variant of Campbell and Cochrane [1999] preferences, but this is not the only way to generate an economically significant prudential uncertainty channel. The two-period model demonstrates that any mechanism which makes the investor 'feel poorer' would also amplify the channel. Borrowing constraints (Favilukis, Ludvigson, and Van Nieuwerburgh [2010], Carroll and Kimball [2001]) and investor leverage are natural alternatives for amplifying the channel.

A Appendix

A1. Proofs

Proof of Proposition 1. [10] A change of variables shows that this model is equivalent to a standard consumption-savings problem. Shift consumption by assigning $C^* = C_1 - H$. Then the date 1 consumption rule can be written as

$$C_1(W_1|H) = C^*(W_1 - H, -H) + H \tag{24}$$

Where $C^*(W, Y)$ solves a simple consumption-savings problem with wealth W and certain future income of Y:

$$C^*(W, Y) \equiv \arg\max_C u(C) + \mathbb{E}\{u[R(W - C) + Y]\} \tag{25}$$

and for ease of notation I suppress the subscript 2 on R. It turns out that for $Y \neq 0$, R random, and CRRA utility, $C^*(W, Y)$ is strictly convex in W. That is, with CRRA utility, rate of return randomness, and the introduction of any (even constant) future income is a sufficient condition for generating strict convexity of the consumption function. This is not one of the sufficient conditions shown in Carroll and Kimball [1996], so I will show that it is sufficient in what follows.

The FOC of the shifted problem (25) is

$$u'(C^*(W, Y)) = \phi'(W - C^*(W, Y)) \tag{26}$$

Where, for convenience, I've defined the function

$$\phi(S) \equiv \mathbb{E}\{u[RS + Y]\}$$

Taking $\frac{\partial}{\partial W}$ of the FOC and rearranging gives

$$\frac{\partial}{\partial W} C^*(W, Y) = \frac{\phi''}{u'' + \phi''}$$

[10] I am grateful to Pok-Sang Lam for teaching me his version of the Carroll and Kimball [1996] proof.

Take another $\frac{\partial}{\partial W}$, do some serious algebra, and we get an expression for the convexity of C^*:

$$\frac{\partial^2}{\partial W^2} C^*(W,Y) = \left[\frac{(u'')^2 \, (\phi'')^2}{u' \times [u'' + \phi'']^3} \right] \left[\frac{\phi' \phi'''}{(\phi'')^2} - \frac{u' u'''}{(u'')^2} \right] \tag{27}$$

The first bracket is negative simply because $u' > 0$ and $u'' < 0$. To show that the second bracket is (strictly) positive, first note that, due to the CRRA specification, $\frac{u'u'''}{(u'')^2} = 1 + \frac{1}{\gamma}$. I will now show that, due to the non-zero future income Y, $\frac{\phi'\phi'''}{(\phi'')^2} > 1 + \frac{1}{\gamma}$. This is an extension of Carroll and Kimball [1996]'s Lemma 4.

Proving $\frac{\phi'\phi'''}{(\phi'')^2} > 1 + \frac{1}{\gamma}$ requires the following technical Lemma.

Lemma 1. *Let Φ_i for $i = 1, ..., N$ be 2×2 symmetric matrices with the following properties:*

- *the diagonals of each Φ_i are positive*
- *the off-diagonals of each Φ_i are all negative*
- *for every i, $|\Phi_i| = 0$*
- $\left| \sum_{i=1}^{N} \Phi_i \right| = 0$

Then for each pair i, j there is some constant k such that

$$\Phi_i = k \Phi_j$$

Proof. I will first show this for the case where $N = 2$. I'll then use the $N = 2$ results to prove the general case. For ease of notation assign $\Phi_1 \equiv \begin{pmatrix} p & q \\ q & r \end{pmatrix}$ and $\Phi_2 \equiv \begin{pmatrix} x & y \\ y & z \end{pmatrix}$. With some algebra, one can show that

$$|\Phi_1 + \Phi_2| = |\Phi_1| + |\Phi_2| + [\sqrt{pz} - \sqrt{xr}]^2 + 2[\sqrt{prxz} - qy] \tag{28}$$

A few facts will let us simplify this expression dramatically. First $|\Phi_1| = |\Phi_2| = |\Phi_1 + \Phi_2| = 0$, so those terms all drop out. Then note that, since $|\Phi_1| = |\Phi_2| = 0$, we have $pr = q^2$ and $xz = y^2$, and we can rewrite

$$\sqrt{prxz} = \sqrt{q^2 y^2} = qy$$

and so the last term in equation (28) also drops out. Thus equation (28) implies that $pz = xr$, or

$$\frac{p}{x} = \frac{r}{z} = k \tag{29}$$

where k is the conjectured constant of proportionality. We just need to show that $\frac{q}{y} = \frac{p}{x} = k$. To show this, plug $pr = q^2$ and $xz = y^2$ into (29) and we have

$$\frac{p}{x} = \frac{q^2/p}{y^2/x} \Rightarrow \frac{p}{x} = \frac{q}{y}$$

This completes the $N = 2$ case.

To show the general case, first note that if Φ_i and Φ_j satisfy the requirements of the lemma, then $\Phi_i + \Phi_j$ also satisfies those requirements. Thus I can apply the $N = 2$ results to the general case, where one matrix is Φ_i and the other matrix is $\sum_{j \neq i} \Phi_i$. Moreover, note that if there is a k such that $\Phi_i = k \sum_{j \neq i} \Phi_i$, then there is m such that $\sum_j \Phi_j = m\Phi_i$. Apply this to all i and get the desired result $\Phi_i = k\Phi_j$. $\square$

Now, back to proving the proposition. I want to show that $\frac{\phi'\phi'''}{(\phi'')^2} > 1 + \frac{1}{\gamma}$. Suppose, for contradiction, that $\frac{\phi'\phi'''}{(\phi'')^2} \leq 1 + \frac{1}{\gamma}$. I can write this expression using the determinant of a 2×2 matrix by defining

$$\Phi \equiv \mathbb{E}\begin{bmatrix} \phi' & \sqrt{1+\frac{1}{\gamma}}\phi'' \\ \sqrt{1+\frac{1}{\gamma}}\phi'' & \phi''' \end{bmatrix} = \mathbb{E}\begin{bmatrix} Ru'(z) & \sqrt{1+\frac{1}{\gamma}}R^2u''(z) \\ \sqrt{1+\frac{1}{\gamma}}R^2u''(z) & R^3u'''(z) \end{bmatrix}$$

Where, for ease of notation, $z \equiv R(W - C^*(W,Y)) + Y$. This expression can now be written compactly as

$$|\Phi| \leq 0$$

Note that Φ is the weighted sum of many component matricies $\begin{bmatrix} Ru'(z) & \sqrt{1+\frac{1}{\gamma}}R^2u''(z) \\ \sqrt{1+\frac{1}{\gamma}}R^2u''(z) & R^3u'''(z) \end{bmatrix}$, and that due to the CRRA specification of u, the determinant of each component matrix is zero. Thus Φ is positive semidefinite, so $|\Phi| \geq 0$. But our assumption for contradiction says Φ is negative semidefinite, and so it must be that $|\Phi| = 0$.

Now I use Lemma 1. The lemma states that if $|\Phi| = 0$, all of the component matrices must be proportional to one another. This means that for any states i and j the ratio of the diagonal terms of the corresponding matrices is equal, that is, for any i and j,

$$\begin{aligned}
&\frac{R_i}{R_j}\frac{u_i'}{u_j'} = \left(\frac{R_i}{R_j}\right)^3 \frac{u_i'''}{u_j'''} \\
\Rightarrow \quad &\left(\frac{R_iS+Y}{R_jS+Y}\right)^{-\gamma} = \left(\frac{R_i}{R_j}\right)^2 \left(\frac{R_iS+Y}{R_jS+Y}\right)^{-\gamma-2} \\
\Rightarrow \quad &\frac{R_iS+Y}{R_jS+Y} = \frac{R_i}{R_j} \\
\Rightarrow \quad &S + \frac{Y}{R_i} = S + \frac{Y}{R_j} \\
\Rightarrow \quad &R_i = R_j
\end{aligned}$$

which is a contradiction, since R is random. Note that the presence of a nonzero income Y is critical because otherwise, I could not move from the fourth line to the fifth line in the equations above.

Therefore, $\frac{\phi'\phi'''}{(\phi'')^2} > 1 + \frac{1}{\gamma}$, and by equation (27), $C^*(W,Y)$ is strictly concave in W, and by equation (24), $C_1''(W|H) > 0$. □

Proof of Proposition 2. Proving this proposition makes use of three lemmas. The first Lemma simply states that consumption increases in both current wealth and future income.

Lemma 2. *$\frac{\partial}{\partial W}C(W,Y) > 0$ and $\frac{\partial}{\partial Y}C(W,Y) > 0$*

Proof. To show $\frac{\partial}{\partial W}C(W,Y) > 0$, take $\frac{\partial}{\partial W}$ of the FOC (26):

$$\frac{\partial}{\partial W}C^*(W,Y) = \frac{\phi''}{u'' + \phi''}$$

This expression is positive because $u'' < 0$ and $\phi'' < 0$.

To show $\frac{\partial}{\partial Y}C(W,Y) > 0$, take $\frac{\partial}{\partial Y}$ of the FOC (26):

$$\begin{aligned}
u''\frac{\partial C}{\partial Y} &= \phi'' \times \left(-\frac{\partial C}{\partial Y}\right) + \frac{\partial \phi'}{\partial Y} \\
\Rightarrow \frac{\partial C}{\partial Y} &= \frac{\frac{\partial \phi'}{\partial Y}}{u'' + \phi''} \\
\Rightarrow \frac{\partial C}{\partial Y} &= \frac{\mathbb{E}[Ru''(z)]}{u'' + \mathbb{E}[R^2u''(z)]}
\end{aligned}$$

Since $u'' < 0$, both the numerator and denominator are negative, so $\frac{\partial}{\partial Y}C(W,Y) > 0$ □

The second Lemma states that habit increases the MPC. The proof of this lemma uses Lemma 2.

Lemma 3. *The MPC increases in H, that is, $\frac{\partial}{\partial H}C_1'(W_0|H) > 0$.*

Proof of Lemma. Taking $\frac{\partial}{\partial W}$ of equation (24),

$$C_1'(W|H) = \frac{\partial}{\partial W}C^*(W-H,-H)$$

and then take $\frac{\partial}{\partial H}$

$$\frac{\partial}{\partial H}C_1'(W|H) = -\frac{\partial^2}{\partial W^2}C^*(W-H,-H) - \frac{\partial^2}{\partial W \partial Y}C^*(W-H,-H)$$

From Proposition 1, we know that $\frac{\partial^2}{\partial W^2}C^*(W-H,-H) < 0$. It remains to show that $\frac{\partial^2}{\partial W \partial Y}C^*(W-H,-H) \leq 0$.

Note that for any W, Y, there is some change in wealth $\Delta(Y)$ such that the change in wealth compensates for having zero future income:

$$C^*(W,Y) = C^*(W+\Delta(Y),0)$$

Taking derivatives,

$$\begin{aligned}
\frac{\partial}{\partial W}C^*(W,Y) &= \frac{\partial}{\partial W}C^*(W+\Delta(Y),0) \\
\frac{\partial^2}{\partial W \partial Y}C^*(W,Y) &= \frac{\partial^2}{\partial W^2}C^*(W+\Delta(Y),0)\Delta'(Y)
\end{aligned}$$

By Carroll and Kimball [1996], the first term $\frac{\partial^2}{\partial W^2}C^*(W+\Delta(Y),0) \leq 0$. To see that the second term $\Delta'(Y) > 0$, note that

$$\frac{\partial}{\partial Y}C^*(W+\Delta(Y),0) = \frac{\partial}{\partial W}C^*(W+\Delta(Y),0)\Delta'(Y)$$

By Lemma 2, $\frac{\partial}{\partial Y}C^* > 0$ and $\frac{\partial}{\partial W}C^* > 0$. Thus, $\Delta'(Y) \geq 0$, which shows that $\frac{\partial^2}{\partial W \partial Y}C^*(W,Y) \leq 0$. $\square$

The final Lemma says that habit has no effect on the MPC for the extremely wealthy.

Lemma 4. $\forall H, \quad \lim_{W\to\infty} C_1'(W|H) = C^*(1,0)$

Proof. Due to homogeneity of the CRRA u, we can rescale the optimization problem (25)

$$
\begin{aligned}
C^*(W,Y) &= W\left\{\arg\max_{C/W} u(C/W) + \mathbb{E}\left\{u[R(1-C/W)+Y/W]\right\}\right\} \\
&= WC^*(1, Y/W)
\end{aligned}
$$

Thus, the MPC can be written

$$
\begin{aligned}
C_1'(W|H) &= \frac{\partial}{\partial W} C^*(W-H,-H) \\
&= \frac{\partial}{\partial W}\left[WC^*\left(1,\frac{-H}{W-H}\right)\right] \\
&= C^*\left(1,\frac{-H}{W-H}\right) + W\frac{H}{(W-H)^2}\frac{\partial}{\partial Y}C^*\left(1,\frac{-H}{W-H}\right)
\end{aligned}
$$

Taking limits

$$
\lim_{W\to\infty} C_1'(W|H) = C^*(1,0)
$$

Note that the RHS does not depend on H. $\square$

Now to prove the proposition, suppose for contradiction that the proposition is false. Then $\exists \epsilon > 0$ such that $C_1''(W|H+\epsilon) - C_1''(W|H) \geq 0$. Taking integrals,

$$
\begin{aligned}
&\int_{\hat{W}}^{\infty} dW\left[C_1''(W|H+\epsilon) - C_1''(W|H)\right] \\
&\qquad = \lim_{W\to\infty}\left[C_1'(W|H+\epsilon) - C_1'(W|H)\right] - \left[C_1'(\hat{W}|H+\epsilon) - C_1'(\hat{W}|H)\right] \geq 0
\end{aligned}
$$

But by Lemma 4, $\lim_{W\to\infty}\left[C_1'(W|H+\epsilon) - C_1'(W|H)\right] = 0$. We now have

$$
C_1'(\hat{W}|H) \geq C_1'(\hat{W}|H+\epsilon)
$$

which contradicts the Lemma 3. Thus, $\forall\epsilon > 0$, $C_1''(W|H+\epsilon) < C_1''(W|H)$. $\square$

Proof of Corollary 2. Let $\underline{W}_0$ and $\overline{W}_0$ be the minimum and maximum of $\mathbb{W}_0$, respectively.

$$\begin{aligned}
&\max_{W_0 \in \mathbb{W}_0} \sigma_0[C_1(W_1|H)] - \min_{W_0 \in \mathbb{W}_0} \sigma_0[C_1(W_1|H)] \\
&\qquad = \sigma_0(\Delta W_1)\left[\max_{W_0 \in \mathbb{W}_0} C_1'(W_0|H) - \min_{W_0 \in \mathbb{W}_0} C_1'(W_0|H)\right] \\
&\qquad = \sigma_0(\Delta W_1)\left[C_1'(\underline{W}_0|H) - C_1'(\overline{W}_0|H)\right] \\
&\qquad = \sigma_0(\Delta W_1)\left[\int_{\overline{W}_0}^{\underline{W}_0} dW C_1''(W|H)\right] \\
&\qquad = -\sigma_0(\Delta W_1)\left[\int_{\underline{W}_0}^{\overline{W}_0} dW C_1''(W|H)\right]
\end{aligned}$$

Where the 2nd line uses equation (2), the third line uses $C_1''(W|H) < 0$. Then the difference between volatility ranges is

$$\begin{aligned}
&\max_{W_0 \in \mathbb{W}_0} \sigma_0[C_1(W_1|H+\epsilon)] - \min_{W_0 \in \mathbb{W}_0} \sigma_0[C_1(W_1|H+\epsilon)] - \max_{W_0 \in \mathbb{W}_0} \sigma_0[C_1(W_1|H)] - \min_{W_0 \in \mathbb{W}_0} \sigma_0[C_1(W_1|H)] \\
&\qquad = -\sigma_0(\Delta W_1)\int_{\underline{W}_0}^{\overline{W}_0} dW \left[C_1''(W|H+\epsilon) - C_1''(W|H)\right]
\end{aligned}$$

Proposition 2 tells us that $C_1''(W|H+\epsilon) - C_1''(W|H) < 0$. Thus the LHS is positive. □

A2. The Interpretation of Habit

Here I describe how deviating from Campbell and Cochrane [1999]'s specification of λ raises some questions regarding the interpretation of habit in the model. The key issue is that the constant λ of my model can make habit decrease in response to an increase in consumption. This violates some traditional notions of habit.

This issue can be illustrated by taking the derivative of log habit h_{t+1} with respect to log consumption c_{t+1}:

$$\frac{dh_{t+1}}{dc_{t+1}} = 1 - \frac{\lambda}{S_{t+1}^{-1} - 1}$$

Thus if S_{t+1} is large enough, $\frac{dh_{t+1}}{dc_{t+1}}$ will be negative.

The preferences of this paper still preserve the standard notion of habit in that H_t is a geometric average of previous consumption. This can be seen by following the analysis of Campbell and Cochrane [1999] found in Campbell [2003]. I can log-linearize the log surplus consumption ratio around the steady state:

$$s_t = \log[1 - \exp(h_t - c_t)] \approx \kappa - \lambda^{-1}(h_t - c_t) \tag{30}$$

Plugging this into the definition of the habit process (5), I find the link between habit and historical consumption

$$h_{t+1} \approx (\text{Constants}) + \rho_s h_t + (1 - \rho_s)c_t \tag{31}$$

$$= (\text{Constants}) + (1 - \rho_s)\sum_{j=0}^{\infty} \rho_s^j c_{t-j} \tag{32}$$

This informal demonstration is verified by simulated data. In the simulated data, habit is highly correlated with consumption. The contemporaneous correlation is 0.978 and the correlation with lagged consumption is 0.983. Habit growth and consumption growth are moderately correlated. The correlation between Δh_t and Δc_t is 0.411.

That habit should move non-negatively with consumption everywhere is not required if one entertains a very slow-moving, historical average of consumption as responsible for our current reference point for consumption. Moreover, Campbell and Cochrane [1999]'s specification is also vulnerable to this this issue. Ljungqvist and Uhlig [2009] show that while habit moves positively with small movements in consumption, it can move negatively with large movements.

The issue illustrated in this section is related to Campbell and Cochrane [1999]'s three requirements on $\lambda(s_t)$. They require (i) the risk-free rate is constant, (ii) habit is predetermined at the steady state surplus consumption, and (iii) habit is predetermined near the steady state. The first assumption is not critical for making habit move non-negatively with consumption. In my model, (ii) is satisfied, but (iii) is not. (iii), in combination with Campbell and Cochrane [1999]'s specification for $\lambda(s_t)$ results in $\frac{dh_{t+1}}{dc_{t+1}} \geq 0$ for all s_t.

A3. Solution Method Details

General Method I approximate the autoregressive process for productivity z_t with a 15 point Markov Chain using the Rouwenhorst method. I solve the model using a projection method, that is, I approximate the law-of-motion for capital $\Gamma(\hat{K}, S, Z_i)$ with some basis functions and then solve for basis function parameters that satisfy the firm's Euler equation. The solution program makes extensive use of the Miranda and Fackler [2001] CompEcon toolbox.

I approximate Γ in the $\hat{K}$ and S directions using a two-dimensional cubic spline. The spline is of

10th degree in the $\hat{K}$ direction and 10th degree in the S direction. The spline breakpoints are log-spaced in both the $\hat{K}$ and S directions. The endpoints of the spline's domain are 0.4 to 4.25 times the non-stochastic steady state capital stock in the $\hat{K}$ direction, and 0.02 to 4.5 times the non-stochastic steady state surplus consumption ratio in the S direction.

In a projection method, one must define what it means to satisfy the Euler equation. I use the collocation method, which specifies that the Euler equation should hold exactly at a set of points (collocation nodes) in the $\hat{K}$ and S domain. I choose these nodes to be the standard nodes for splines using knot averaging. I consider the algorithm converged if, across these collocation nodes, the maximum absolute Euler equation residual, expressed as $1 - E[M'R_I']$, is less than 1.0E-8. I search for spline coefficients which satisfy this condition by using Broyden's method (a Quasi-Newton method).

On the collocation nodes, the Euler equation error is less than 1.0-8, but the solution cannot be so exact for the entire state space. I find that along the 1,000 simulations of 504 quarters, the mean absolute euler equation error is 0.0004. This value is the same when I examine the pricing kernel equation for the return on equity and the return on the consumption claim. The maximum error along during these 504,000 quarters is 0.0017. I find that increasing the degree to 14th in the S direction or to 14th in the $\hat{K}$ direction has no material impact on the quantitative results.

Solving for the Evolution of Surplus Consumption An unusual aspect of this model is that it contains a state variable which is not predetermined but is not exogenous. Surplus consumption tomorrow is not known today and it depends on consumption tomorrow, which is endogenous, and this makes solving for the surplus consumption process rather difficult.

To see this clearly, it helps to write the evolution of surplus consumption (5) as functions of state variables

$$s' = (1 - \rho_s)\bar{s} + \rho_s s + \lambda[c(\hat{K}', S', Z_j) - c(\hat{K}, S, Z_j)] \tag{33}$$

Where Z_i and Z_j are discrete productivity states. Since $\hat{K}'$ is predetermined, s' is a function of four variables $s'(\hat{K}, S, Z_i, Z_j)$. Note that surplus consumption tomorrow s' appears on both sides of this equation. Additionally, consumption is built from a 2D cubic spline and so this equation cannot be solved for s' analytically. To deal with this issue, I solve equation (33) numerically. This calculation is very computationally intensive because I must solve (33) numerically at every collocation node for every potential productivity shock within every iteration of the big Broyden's method which is solving for the coefficients of the approximation of $\hat{\Gamma}$. The fact that there are three state variables makes the number of collocation nodes which need to be dealt with very large. To speed up computation, I solve for many of the s' values simultaneously using Broyden's method.

Homotopy Method Projection methods require a good initial guess of the spline coefficients. The typical initial guess is based off a planner's problem (which is guaranteed to converge with value function iteration), however, the external habit in the model makes the model equilibrium very far from the planner's problem. To overcome this issue, I use a homotopy method. Specifically, I modify the firm's problem so that it discounts future profits using the SDF

$$M_{t,t+1} = \beta \left[\frac{\hat{C}_{t+1}}{\hat{C}_t} \left(\frac{\hat{S}_{t+1}}{\hat{S}_t} \right)^{\chi} \right]^{\gamma}$$

$\chi = 0$ corresponds to a model with no habit, and $\chi = 1$ corresponds to the full model. I begin by solving the model for $\chi = 0$, and then slowly increase χ, using the coefficients from the previous χ as the initial guess for the current χ.

Euler Equation To be explicit about the firm's Euler equation, let $\pi_Z(Z_i, Z_j)$ be the transition matrix for the discretized productivity process. The firm's problem is to find capital policy $K' = G(K; \hat{K}, S, Z_i)$ to solve

$$\begin{aligned} V(K; \hat{K}, S, Z_i) = \max_{K', I, N} \Big\{ &\Pi(K, Z_i, N) - W(\hat{K}, S, Z_i)N - \Phi(I, K) - I \\ &+ \sum_{Z_j} \pi_Z(Z_i, Z_j) M(\hat{K}, S, Z_i; Z_j) V(K'; \hat{K}', S', Z_j) \Big\} \end{aligned}$$

subject to

$$K' = I + (1 - \delta)K$$

The FOC for investment and the envelope condition are:

$$\begin{aligned} 1 + D_1\Phi(I, K) &= \sum_{Z_j} \pi_Z(Z_i, Z_j) M(\hat{K}, S, Z_i; Z_j) D_1 V(K'; \hat{K}', S', Z_j) \\ D_1 V(K; \hat{K}, S, Z_i) &= D_1\Pi(K, N, Z_i) + (1 + D_1\Phi(I, K))(1 - \delta) - D_2\Phi(I, K) \end{aligned}$$

which together produce the Euler equation

$$1 + D_1\Phi(I,K) = \sum_{Z_j} \pi_Z(Z_i,Z_j)M(\hat{K},S,Z_i;Z_j)[D_1\Pi(K',Z_j,N') + (1 + D_1\Phi(I_j,K'))(1-\delta) - D_2\Phi(I_j,K')]$$

Impose the fact that the household does not value leisure and consistency, and we have

$$1 = \sum_{Z_j} \pi_Z(Z_i,Z_j)M(\hat{K},S,Z_i;Z_j)\hat{R}_I(\hat{K},S,Z_i;Z_j) \tag{34}$$

$$\hat{R}_I(\hat{K},S,Z_i;Z_j) \equiv \frac{1 + D_1\Pi(\hat{K}',Z_j,1) + (1 + D_1\Phi(\hat{I}_j,\hat{K}'))(1-\delta) - D_2\Phi(\hat{I}_j,\hat{K}')}{1 + D_1\Phi(\hat{I},\hat{K})}$$

Where

$$M(\hat{K},S,Z_i,Zj) = \beta\left(\frac{C_j}{C}\frac{S_j}{S}\right)^{-\gamma} \tag{35}$$

$$C_j = \Pi(\hat{K}',Z_j,1) - \Phi(\hat{I}_j,\hat{K}') - \hat{I}_j$$

$$C = \Pi(\hat{K},Z_i,1) - \Phi(\hat{I},\hat{K}) - \hat{I}$$

$$\hat{K}' = \hat{G}(\hat{K},S,Z_i)$$

$$\hat{I} = \hat{G}(\hat{K},S,Z_i) - (1-\delta)\hat{K}$$

$$\hat{I}_j = \hat{G}(\hat{K}',S',Z_j) - (1-\delta)\hat{K}'$$

and the evolution of surplus consumption satisfies

$$s_j = (1-\rho_s)\bar{s} + \rho_s s + \lambda(c_j - c) \tag{36}$$

The projection algorithm looks for cubic spline coefficients which solve equations (34), (35), and (36).

Solving for asset prices To find asset prices, the firm's Bellman equation, with optimal values plugged in, is:

$$\begin{aligned} V(K;\hat{K},S,Z_i) =& \Pi(K,Z_i,1) - W(\hat{K},S,Z_i) - \Phi(I,K) - I \\ &+ \sum_{Z_j} \pi_Z(Z_i,Z_j)M(\hat{K},S,Z_i;Z_j)V(K';\hat{K}',S',Z_j) \end{aligned}$$

In equilibrium, $W(\hat{K}, S, Z_i) = (1-\alpha)AZ_i\hat{K}^\alpha$, $K = \hat{K}$, $I = \hat{I}$, and $K' = \hat{K}'$, so

$$\begin{aligned} V(\hat{K}; \hat{K}, S, Z_i) =& \alpha A Z_i \hat{K}^\alpha - \Phi(\hat{I}, \hat{K}) - \hat{I} \\ &+ \sum_{Z_j} \pi_Z(Z_i, Z_j) M(\hat{K}, S, Z_i; Z_j) V(\hat{K}'; \hat{K}', S', Z_j) \end{aligned}$$

Let's define $\hat{V}(\hat{K}, S, Z_i) \equiv V(\hat{K}; \hat{K}, S, Z_i)$. The above equation suggests that $\hat{V}(\hat{K}, S, Z_i)$ can be found by repeatedly applying the above equation as an operator (using the law of motion for capital $\hat{G}(\hat{K}, S, Z_i)$).

A4. Data and Simulation Details

The data span 1948Q1 thru 2010Q4. Macroeconomic quantities are from the BEA NIPA, BEA fixed asset tables, and Simona Cociuba's website. All data is real and per-capita. Output is simply GDP. Consumption is nondurable consumption plus services consumption. Investment and capital are fixed investment plus nondurable consumption plus government investment. All quantities are constructed by dividing nominal figures (Table 1.1.5 or Table 3.9.5) by the appropriate price index (Table 1.1.4 or Table 3.9.4), and then dividing by the population (from Table 2.1).

Most studies exclude government investment from their definition of investment, citing the fact that their models do not include government. However, this logic would also imply that one should remove government purchases from GDP, which is not typically done. I choose to keep government purchases in GDP and include government investment in investment as this approach preserves the idea that the data aggregates come from a single general equilibrium system. Moreover, this approach is closer to the spirit of Cooley and Prescott [1995].[11]

Asset price data is taken from CRSP. The equity is the CRSP value-weighted index, and the risk-free rate is the return on the 30-day T-bill. Price-dividend ratios are computed annually using the stock market reinvestment assumption. As discussed in Cochrane [2011], this aggregation method preserves the Campbell-Shiller identity, which is useful for identifying the source of asset price fluctuations. All asset prices are deflated using the CPI, also obtained from CRSP.

Model moments are computed using 1000 simulations of 504 quarters, with the first 252 quarters dropped. Each simulation is initialized with K equal to the median capital stock from a simulation of 200,000 quarters, S equal to $\bar{S}$, and Z_i equal to the median productivity level. Scatterplots are

[11] Cooley and Prescott [1995] say that "Our economy is very abstract: it contains no government sector, no household production sector, no foreign sector and no explicit treatment of inventories. Accordingly, the model economy's capital stock, K, includes capital used in all of these sectors plus the stock of inventories. Similar, output, Y, includes the output produced by all of this capital" (page 17).

generated by randomly selecting 5,000 quarters among the 1000 simulations (ignoring the first 252 quarters of each simulation).

A5. Adjustment Costs and Accounting

Many papers specify adjustment costs in the following manner:

$$K' = (1-\delta)K + \phi(I/K)K$$

(e.g. Jermann [1998], Gourio [2009], Kaltenbrunner and Lochstoer [2010], Guvenen [2009], among others). This formulation preserves the traditional Cobb-Douglas formulation of output:

$$Y = ZK^{\alpha}N^{1-\alpha} = C + I$$

This formulation, however, deviates from the standard accounting treatment of investment and capital. The standard treatment specifies that end-of-period capital is beginning-of-period capital plus investment less depreciation. With geometric depreciation, this translates into the standard, adjustment-cost-free formulation of capital accumulation:

$$K' = (1-\delta)K + I$$

I choose to preserve this accounting identity. As a result, capital adjustment costs are pushed into output:

$$Y = ZK^{\alpha}N^{1-\alpha} - [\text{Adj Cost}] = C + I$$

Fortunately, both choices result in the same capital and consumption allocations. For example, Gourio [2009] uses

$$K^{*\prime} = (1-\delta)K + \phi^*(I^*/K)K$$
$$\phi^*(x) = x - \frac{\eta}{2}(x-\delta)^2$$

From these expressions, we have capital evolution and consumption

$$K^{*\prime} = (1-\delta)K + I^* - \frac{\eta}{2}\left(I^*/K - \delta\right)^2 K$$
$$C = ZK^\alpha N^{1-\alpha} - I^*$$

which is identical to my formulation with

$$I \equiv I^* - \frac{\eta}{2}\left(I^*/K - \delta\right)^2 K$$

References

O.P. Attanasio. Consumption. *Handbook of macroeconomics*, 1:741–812, 1999.

R. Bansal, V. Khatchatrian, and A. Yaron. Interpretable asset markets? *European Economic Review*, 49(3):531–560, 2005.

R. Bansal, D. Kiku, and A. Yaron. An empirical evaluation of the long-run risks model for asset prices. Technical report, National Bureau of Economic Research, 2009.

J. Beeler and J.Y. Campbell. The long-run risks model and aggregate asset prices: an empirical assessment. Technical report, National Bureau of Economic Research, 2009.

N. Bloom. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, 2009.

M. Boldrin, L.J. Christiano, and J.D.M. Fisher. Habit persistence, asset returns and the business cycle. 1999.

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

M.W. Brandt. Portfolio choice problems. *Handbook of financial econometrics*, 1(1), 2009.

J.Y. Campbell. Inspecting the mechanism: An analytical approach to the stochastic growth model. *Journal of Monetary Economics*, 33(3):463–506, 1994.

J.Y. Campbell. Consumption-based asset pricing. *Handbook of the Economics of Finance*, 1:803–887, 2003.

J.Y. Campbell and J.H. Cochrane. By Force of Habit: A Consumption Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy*, 107(2):205–251, 1999.

J.Y. Campbell and R.J. Shiller. The dividend-price ratio and expectations of future dividends and discount factors. *Review of financial studies*, 1(3):195–228, 1988.

C.D. Carroll. How does future income affect current consumption? *The Quarterly Journal of Economics*, 109(1):111–147, 1994.

C.D. Carroll. Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly Journal of Economics*, 112(1):1–55, 1997.

C.D. Carroll. A theory of the consumption function, with and without liquidity constraints. *The Journal of Economic Perspectives*, 15(3):23–45, 2001.

C.D. Carroll and M.S. Kimball. On the concavity of the consumption function. *Econometrica*, 64(4): 981–992, 1996.

C.D. Carroll and M.S. Kimball. Liquidity constraints and precautionary saving. Technical report, National Bureau of Economic Research, 2001.

J.H. Cochrane. The dog that did not bark: A defense of return predictability. *Review of Financial Studies*, 21(4):1533–1575, 2008a.

J.H. Cochrane. Financial markets and the real economy. 2008b.

J.H. Cochrane. Presidential address: Discount rates. *The Journal of Finance*, 66(4):1047–1108, 2011.

T. Cogley and J.M. Nason. Output dynamics in real-business-cycle models. *The American Economic Review*, pages 492–511, 1995.

G.M. Constantinides. Habit formation: A resolution of the equity premium puzzle. *Journal of political Economy*, pages 519–543, 1990.

T.F. Cooley and E.C. Prescott. Economic growth and business cycles. *Frontiers of business cycle research*, pages 1–38, 1995.

A. Deaton. Saving and liquidity constraints. *Econometrica*, 59(5):1221–1248, 1991.

I. Dew-Becker. A model of time-varying risk premia with habits and production. 2011.

L.G. Epstein and S.E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica: Journal of the Econometric Society*, pages 937–969, 1989.

J. Favilukis and X. Lin. Infrequent renegotiation of wages: A solution to several asset pricing puzzles. Technical report, Working paper London School of Economics and Ohio State University, 2011.

J. Favilukis, S.C. Ludvigson, and S. Van Nieuwerburgh. The macroeconomic effects of housing wealth, housing finance, and limited risk-sharing in general equilibrium. Technical report, National Bureau of Economic Research, 2010.

L.R. Glosten, R. Jagannathan, and D.E. Runkle. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of finance*, pages 1779–1801, 1993.

F. Gourio. Disasters risk and business cycles. Technical report, National Bureau of Economic Research, 2009.

F. Gourio. Disaster risk and business cycles. 2010.

F. Guvenen. A parsimonious macroeconomic model for asset pricing. *Econometrica*, 77(6):1711–1750, 2009. ISSN 1468-0262.

L.P. Hansen and R. Jagannathan. Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–62, 1991.

U.J. Jermann. Asset pricing in production economies. *Journal of Monetary Economics*, 41(2):257–275, 1998.

K.L. Judd. Projection methods for solving aggregate growth models. *Journal of Economic Theory*, 58 (2):410–452, 1992.

G. Kaltenbrunner and L.A. Lochstoer. Long-run risk through consumption smoothing. *Review of Financial Studies*, 23(8):3190–3224, 2010.

S. Kandel and R.F. Stambaugh. Expectations and volatility of consumption and asset returns. *Review of Financial Studies*, 3(2):207–232, 1990.

J.M. Keynes. *The general theory of interest, employment and money*. London: Macmillan, 1936.

A. Khan and J.K. Thomas. Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics. *Econometrica*, 76(2):395–436, 2008. ISSN 1468-0262.

M.S. Kimball. Precautionary saving in the small and in the large. *Econometrica: Journal of the Econometric Society*, pages 53–73, 1990.

L.A. Kuehn, N. Petrosky-Nadeau, and L. Zhang. An equilibrium asset pricing model with labor market search. Technical report, National Bureau of Economic Research, 2012.

S.F. LeRoy and R.D. Porter. The present-value relation: Tests based on implied variance bounds. *Econometrica: Journal of the Econometric Society*, pages 555–574, 1981.

M. Lettau. Inspecting the mechanism: Closed-form solutions for asset prices in real business cycle models*. *The Economic Journal*, 113(489):550–575, 2003.

M. Lettau and H. Uhlig. Can habit formation be reconciled with business cycle facts? *Review of Economic Dynamics*, 3(1):79–99, 2000.

M. Lettau, S.C. Ludvigson, and J.A. Wachter. The declining equity premium: What role does macroeconomic risk play? *Review of Financial Studies*, 21(4):1653–1687, 2008.

L. Ljungqvist and H. Uhlig. Optimal endowment destruction under campbell-cochrane habit formation. Technical report, National Bureau of Economic Research, 2009.

M. Miranda and P. Fackler. Applied computational economics, 2001.

D. Papanikolaou. Investment shocks and asset prices. *Journal of Political Economy*, 119(4):639–685, 2011.

L. Pastor and P. Veronesi. Uncertainty about government policy and stock prices. *The Journal of Finance*, 67(4):1219–1264, 2012.

O. Posch. Risk premia in general equilibrium. *Journal of Economic Dynamics and Control*, 35(9): 1557–1576, 2011.

T. Santos and P. Veronesi. Habit formation, the cross section of stock returns and the cash-flow risk puzzle. *Journal of Financial Economics*, 98(2):385–413, 2010.

RJ Shiller. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71(3):421–436, 1981.

S.M. Sundaresan. Intertemporally dependent preferences and the volatility of consumption and wealth. *Review of financial Studies*, 2(1):73–89, 1989.

S.P. Zeldes. Optimal consumption with stochastic income: Deviations from certainty equivalence. *The Quarterly Journal of Economics*, 104(2):275–298, 1989.