Robust Equilibrium Yield Curves*

Isaac Kleshchelski†, Nicolas Vincent‡

Abstract

This paper studies the quantitative implications of the interaction between uncertainty aversion and stochastic volatility for key asset pricing phenomena. We present an equilibrium term structure model in which output growth is conditionally heteroskedastic. The agent does not know the true model of the economy and chooses optimal policies that are robust to model misspecification. The choice of robust policies greatly amplifies the effect of conditional heteroskedasticity. We estimate the model by exploiting closed-form moment conditions and find that it is consistent with key empirical regularities of both the bond and equity markets. We also show that the set of models the robust representative agent entertains is 'small'.

*We are especially indebted to Martin Eichenbaum and Sergio Rebelo for their continuous advice. We also thank Torben Andersen, Gadi Barlevi, Marco Bassetto, Craig Burnside, Lars Hansen, Zhiguo He, Ravi Jagannathan, Arvind Krishnamurthy, Jonathan Parker, Giorgio Primiceri, Costis Skiadas, Annette Vissing-Jorgensen, Ernst Schaumburg and seminar participants at Barclays Global Investors, Columbia, Duke, Emory, Federal Reserve Bank of Chicago, Federal Reserve Bank of New York, Federal Reserve Board of Governors, Northwestern, UC Berkeley, UCLA, UCSD, University of Houston, University of Minnesota, Washington University.

†Washington University in St.Louis, Olin Business School, 1 Brookings Drive, St. Louis, MO 63130. Email: kleshchelski@wustl.edu

‡HEC Montréal and CIRPÉE, 3000 chemin de la Côte-Sainte-Catherine, Montréal, QC, Canada H3T 2A7. Email: nicolas.vincent@hec.ca

1 Introduction

This paper studies the implications of the interaction between robust control and stochastic volatility for key asset pricing phenomena for both the bond and equity markets. We estimate a parsimonious one-factor asset pricing model and show that robustness, or fear of model misspecification, coupled with state-dependent volatility provides an empirically plausible characterization of the level and volatility of the equity premium, the risk free rate, and the cross-section of real yields on treasury bonds. We also show that robustness offers a novel way of reconciling the shape of the term structure of interest rates with the persistence of yields. Finally, we quantify the level of robustness encoded in agents' behavior.

Our framework is a continuous-time, Lucas (1978)-type, real asset pricing model in which a representative agent is averse to both risk and ambiguity. The presence of ambiguity stems from the agent's incomplete information about the economy's data generating process (DGP). Introducing ambiguity aversion into our framework allows us to reinterpret an important fraction of the market price of risk as the market price of model (or Knightian) uncertainty. We model ambiguity aversion using robust control techniques as in Anderson et al. (2003).[1] In our model, the representative agent distrusts the reference model and optimally chooses a distorted DGP. His consumption and portfolio decisions are then based on this distorted distribution. Ambiguity aversion gives rise to endogenous pessimistic assessments of the future. Being pessimistic, the agent tilts his subjective distribution towards states in which marginal utility is high.

Once we introduce stochastic volatility into the model, a positive surge in volatility results in a more diffuse distribution of future consumption growth. The agent seeks policies that

[1] Behavioral puzzles such as the Ellsberg paradox (Ellsberg (1961)) led to the axiomatization of the maxmin decision making by Gilboa and Schmeidler (1989). Robust control is one way of modeling Knightian uncertainty. For a comprehensive treatment of robustness see Hansen and Sargent (2007a). Examples of the use of robust control in economics and finance include Anderson et al. (2003), Cagetti et al. (2002), , Hansen and Sargent (2010), Hansen et al. (2006), Liu et al. (2005), Maenhout (2004), Routledge and Zin (2001), Uppal and Wang (2003). An alternative approach to modeling ambiguity allows agents to have multiple priors. See, for example, Epstein and Schneider (2003), Epstein and Wang (1994).

can reasonably guard against 'bad' realizations of the consumption process. The interaction between robustness and stochastic volatility introduces a state dependent distortion to the drift of consumption growth, and therefore, to the drift in the agent's intertemporal marginal rate of substitution. As we show, this state dependent distortion generates sharp implications for numerous asset pricing phenomena.

We contribute to the literature along three main dimensions. First, our parsimonious framework allows us to focus on a small number of free parameters and to derive closed-form equilibrium conditions instead of relying on approximations. Second, we bring the model to the data by fitting analytical moment conditions using GMM. To be consistent with the real structure of our framework, we are careful to estimate the model using inflation-adjusted data. Third, our estimation exercise is disciplined by the use of cross-equation restrictions across *both* the bond *and* equity markets to improve the identification of structural parameters in our model and the estimation of the market price of risk and uncertainty.[2]

Our main findings are as follows. First, we show that our model, calibrated with a unitary degree of risk aversion and elasticity of intertemporal substitution (EIS), can reproduce both the high and volatile equity premium and the low and stable risk free rate observed in the data. Previous studies, such as Mehra and Prescott (1985) and Weil (1989), show that explaining the behavior of the equity premium requires implausibly high levels of risk aversion. Ambiguity aversion greatly amplifies the effect of stochastic volatility in consumption growth and generates an uncertainty premium that helps to alleviate the difficulties encountered in these previous studies. In fact, we show that in the absence of ambiguity aversion, plausible degrees of stochastic volatility in consumption growth do not generate sufficient variation in the stochastic discount factor. Since there is no benchmark value for the degree of ambiguity aversion, we use detection error probabilities to show that the degree of robustness required to fit the data is reasonable. In other words, we show empirically that the set of models the

[2]See Campbell (2000) for a discussion of the importance of exploiting restrictions across markets to discipline the estimation of term structure models.

robust representative agent entertains is small. By this we mean that it is statistically difficult to distinguish between models in this set.

Second, our model can account for the means of the cross-section of real bond yields. In particular, we can replicate the upward sloping unconditional yield curve observed in the data. This result highlights a novel interpretation of the uncertainty premium generated by robustness. On empirical grounds, we assume that the conditional variance of output growth, and hence consumption growth, is stationary and positively correlated with the consumption growth process. This positive correlation implies that when marginal utility is high the conditional variance of consumption growth is low. Consequently, a downward bias in the subjective conditional expectations of consumption growth induces a negative distortion to the subjective expectations of variance changes. We show that this negative distortion is a linear function of the level of the conditional variance process. Consequently, the unconditional distortion is a linear negative function of the objective steady state of the variance process. Therefore, the subjective steady state of the variance process is lower than the objective steady state. In other words, on average, the agent thinks that the conditional variance of consumption growth should decrease. Since the unconditional level of bond yields and the steady state level of the conditional volatility of consumption growth are inversely related, the agent expects, on average, that yields will increase. Consequently, the unconditional yield curve is upward sloping.

Third, our model can replicate the declining term structure of unconditional volatilities of real yields, and the negative correlation between the level and the spread of the real yield curve. The fact that the robust distortion to the conditional variance process is a linear function of the level of the variance implies that the distorted process retains the mean reversion structure of the objective process. Since shocks to the conditional variance are transitory, the short end of the yield curve is more responsive to volatility shocks relative to the long end. Hence, short maturity yields are more volatile than long maturity yields. Also, our model implies that yields are an affine function of the conditional variance of consumption growth. Therefore, all yields are perfectly positively correlated. Short yields are more responsive to volatility shocks than long yields, but both move in the same direction. So, when yields decrease, the spread between long yields and short yields increases and becomes more positive. As a result, the level and spread of the real yield curve are negatively correlated.

Fourth, the model can reconcile two seemingly contradictory bond market regularities: the strong concavity of the short end of the yield curve and the high degree of serial correlation in bond yields.[3] The intuition for this result is closely linked to the mechanism behind the upward sloping real yield curve. Generally, in a one-factor affine term structure model, the serial correlation of yields is driven by the serial correlation of the state variable implied by the objective DGP. In contrast, in our model the slope of the yield curve is shaped by the degree of mean reversion of the conditional variance process implied by the agent's distorted (i.e., subjective) distribution. The state dependent distortion to the variance process not only changes the perceived steady state of the variance but also its velocity of reversion. With positive correlation between consumption growth and the conditional variance process, we show that the subjective mean reversion is faster than the objective one. Ex ante, the agent expects shocks to the variance process to die out fast, but ex-post these shocks have a longer lasting effect than expected. The slope of the yield curve is a reflection of how fast the agent expects the effect of variance surges to dissipate. The positive slope of the yield curve declines rapidly when the subjective mean reversion is high. The serial correlation of yields is measured ex-post, using realized yields. If the objective persistence of the variance process is high, yields are highly persistent, which is in line with the empirical evidence. We also show that when the agent seeks more robustness, the separation between the ex-ante and ex-post persistence is stronger.

The remainder of the paper is organized as follows. In section 2 we present our continuous time model and discuss its implications for the equity and bond markets. We derive analytical affine term structure pricing rules and discuss the distinction between the market price of

[3] In a standard one-factor model, it is difficult to separate these two properties, since both observations are directly tied to the persistence of the underlying univariate shock process.

risk and uncertainty. In sections 3 and 4 we describe our estimation procedure, present the empirical results and explain why our model can match a number of asset pricing stylized facts. We also present independent evidence that supports our modeling assumptions in section 5. We investigate the implied level of uncertainty aversion exhibited by the representative agent in section 6. Finally, we offer our concluding remarks and discuss potential avenues for future research in section 7 .

2 Robustness in a Continuous Time Model with Stochastic Volatility

In this section we present an infinite horizon, continuous time, general equilibrium model in which a robust representative agent derives optimal policies about consumption and investment.[4] For simplicity we assume a Lucas tree type economy with a conditionally heteroskedastic growth rate of output. Our ultimate goal is to analyze the implied equilibrium yield curve in this economy and, in particular, identify the implications of robustness for the term structure of interest rates.[5]

2.1 Reference and Distorted Models

The representative agent in our economy uses a *reference* or *approximating* model. However, since he fears that this model is potentially misspecified, he chooses to diverge from it when making his decisions.[6] In the context of this paper, the reference model is assumed to generate

[4] In the working paper version available at http://neumann.hec.ca/pages/nicolas.vincent/ we present the robust control idea in a simple two-period asset pricing model.

[5] Our work is closely related to Gagliardini et al. (2009). There are however significant differences between the two approaches: in their case, they use a Multiple Prior implementation of robust control in the context of a two-factor model related to Longstaff and Schwartz (1992);since a closed-form solution cannot be derived, they rely on an approximation around the equilibrium without ambiguity aversion; they use nominal instead of real data; and finally, their empirical exercise is aimed at fitting exclusively moments for the term structure of interest rates, and not the equity market.

[6] Another possibility is to claim that for some reason the agent dislikes extreme negative events and wants to take special precautionary measures against these events. If we choose this behavioral interpretation, we can

the observed data. In contrast with the rational expectations paradigm, the agent entertains alternative DGPs. The size of the set of possible models is implicitly defined by a penalty function (relative entropy) incorporated into the agent's utility function. So the agent chooses an optimal distorted distribution for the exogenous processes. In other words, the agent optimally chooses his set of beliefs simultaneously with the usual consumption and investment decisions. The robust agent distorts the approximating model in a way that allows him to incorporate fear of model misspecification. We will refer to the optimally chosen model as the *distorted* model.[7]

2.2 The Economy

There is a single consumption good which serves as the numeraire. Let D be an exogenous output process that follows a geometric Brownian motion and solves the following stochastic differential equation (SDE),

$$dD_t = D_t \mu dt + D_t \sqrt{v_t} dB_t. \tag{2.1}$$

The probability measure $\mathbb{P}$ on the Brownian process represents the reference or approximating model. However, the agent entertains a *set* of possible probability measures which size is determined by a penalty function (relative entropy) that is incorporated into the agent's utility function. We denote the distorted measure which the agent chooses by $\mathbb{Q}$.[8] The conditional

then assume the agent knows the true DGP, but that his marginal utility function is very high in bad states of the world. Low consumption is so costly that the agent requires policies that are robust to these states. Even though there is complete observational equivalence between the two approaches, they are utterly different from a behavioral perspective.

[7] An alternative is to allow for the possibility that a different, unspecified model, is actually the DGP. In this scenario, it is likely that neither the distorted nor the reference model generate the data. The agent must in this case infer which model is more likely to generate the data. See Hansen and Sargent (2010) for an example.

[8] Formally, we fix a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ supporting a univariate Brownian motion $B = \{B_t : t \geq 0\}$. The diffusion of information is described by the filtration $\{\mathcal{F}_t\}$ on $(\Omega, \mathcal{F})$. All stochastic processes are assumed to be progressively measurable relative to the augmented filtration generated by B. The *set* of possible probability measures on $(\Omega, \mathcal{F})$ entertained by the agent is denoted by $\mathcal{P}$. Every element in $\mathcal{P}$ defines the same null events as $\mathbb{P}$. Note that the assumption that the penalty function is the relative entropy imposes a lot of structure on the possible distorted measures. By Girsanov's theorem we require the distorted measure to be absolutely continuous with respect to the reference measure.

expectation operators under $\mathbb{P}$ and $\mathbb{Q}$ are denoted respectively by $\mathbb{E}_t(\cdot)$ and $\mathbb{E}_t^{\mathbb{Q}}(\cdot)$.

We can obviously think of D as a general dividend process of the economy. We allow the trading of ownership shares of the output tree. The parameters μ and v are the local expectations (drift) and the local variance of the output growth rate, respectively. We assume that v follows a mean-reverting square-root process:

$$\begin{aligned} dv_t &= (a_0 + a_1 v_t)\, dt + \sqrt{v_t}\sigma_v dB_t, \\ a_0 &> 0, \quad a_1 < 0, \quad \sigma_v \in \mathbb{R}, \quad 2a_0 \geq \sigma_v^2. \end{aligned} \tag{2.2}$$

Note that the same shock (Wiener increments) drives both the output growth and the output growth volatility processes.[9] We impose this assumption to retain parsimony. The requirement $a_1 < 0$ guarantees that v converges back to its steady state level $-\frac{a_0}{a_1} (= \bar{v})$ at a velocity $-a_1$. The long run level of volatility is positive since $a_0 > 0$. The Feller condition $2a_0 \geq \sigma_v^2$ guarantees that the drift is sufficiently strong to ensure that $v > 0$ a.e. once $v_0 > 0$. The parameter σ_v is constant over time and will play an important role in our model.

When v is constant over time, the market price of risk is state independent, and the expectations hypothesis of the term structure of interest rates holds. This result stands in sharp contrast to the empirical evidence (e.g., Fama and Bliss (1987), Campbell and Shiller (1991), Backus et al. (1998), Cochrane and Piazzesi (2002)). We discuss in the next section how stochastic volatility interacts with robustness considerations to affect the predictions of our model.

Let dR_t be the instantaneous return process on the ownership of the output process and S_t be the price of ownership at time t. Then, we can write

$$\begin{aligned} dR_t &\equiv \frac{dS_t + D_t dt}{S_t} \\ &= \mu_{R,t} dt + \sigma_{R,t} dB_t, \end{aligned} \tag{2.3}$$

[9] We could also make the expected instantaneous output growth rate, μ, stochastic. By assuming, for example, an affine relation between μ_t and v_t, the model remains tractable and can be solved analytically. For the purpose of this paper, however, we maintain the assumption of a constant drift in the output process.

where μ_R and σ_R are determined in equilibrium. We also let r be the short rate process, which is determined in equilibrium.

2.3 The Dynamic Program of the Robust Representative Agent

The robust representative agent consumes continuously and invests both in a risk-free and a risky asset. The risky asset corresponds to the ownership of a share of the output process (the tree). The risk free asset is in zero net supply in equilibrium. The agent chooses optimally a distortion to the underlying model in a way that makes his decisions robust to statistically small model misspecification. Formally, the agent has the following objective function

$$\sup_{C,\alpha} \inf_{\mathbb{Q}} \left\{ \mathbb{E}_t^{\mathbb{Q}} \left[\int_t^\infty e^{-\rho(s-t)} u\left(C_s\right) ds \right] + \theta \mathcal{R}_t\left(\mathbb{Q}\right) \right\}, \tag{2.4}$$

subject to his dynamic budget constraint

$$dW_t = \left[r_t W_t + \alpha_t W_t \left(\mu_{R,t} - r_t \right) - C_t \right] dt + \alpha_t W_t \sigma_{R,t} dB_t, \tag{2.5}$$

where $\mathbb{Q}$ is the agent's subjective distribution, W is the agent's wealth, ρ is the subjective discount factor, C is the consumption flow process, α is the portfolio share invested in the risky asset, and θ is the multiplier on the relative entropy penalty $\mathcal{R}$. The level of θ can be interpreted as the magnitude of the desire to be robust. When θ is set to infinity, (2.4) converges to the expected time-additive utility case. A lower value of θ means that the agent is more fearful of model misspecification and thus chooses $\mathbb{Q}$ further away from $\mathbb{P}$ in the relative entropy sense. In other words, the set of alternative DGPs is larger the smaller θ is.

Under some regularity conditions and by Girsanov's theorem, we can define a Brownian motion under $\mathbb{Q}$ as

$$B_t^{\mathbb{Q}} = B_t - \int_0^t h_s ds, \quad t \geq 0. \tag{2.6}$$

With this setup at hand, the relative entropy process $\mathcal{R}\left(\mathbb{Q}\right)$ for some $\mathbb{Q}$ can be expressed conveniently as[10]

$$\mathcal{R}_t(\mathbb{Q}) = \frac{1}{2}\mathbb{E}_t^{\mathbb{Q}}\left[\int_t^\infty e^{-\rho(s-t)} h_s^2 ds\right], \quad t \geq 0. \tag{2.7}$$

and (2.4) becomes

$$\sup_{C,\alpha} \inf_h \left\{ \mathbb{E}_t^{\mathbb{Q}} \int_t^\infty e^{-\rho(s-t)} \left[u(C_s) + \frac{\theta}{2} h_s^2 \right] ds \right\}. \tag{2.8}$$

The change of measure (2.6) also allows us to write (2.1), (2.2) and (2.5) under the distorted measure $\mathbb{Q}$. This introduces a drift distortion, which in the context of the market return has an obvious interpretation: it is the uncertainty premium the agent requires for bearing the risk of potential model misspecification

$$dR_t = \left[\mu_{R,t} - \underbrace{(-h_t \sigma_{R,t})}_{\text{Uncertainty premium}} \right] dt + \sigma_{R,t} dB_t^{\mathbb{Q}}. \tag{2.9}$$

The process h is the (negative of the) process for the *market price of model uncertainty.* The diffusion part $\sigma_{R,t}$ on the return process is, as usual, the *risk exposure* of the asset. The product $-h_t \sigma_{R,t}$ is the equilibrium *uncertainty premium.*[11]

2.4 Optimal Policies with Robust Control

We focus on the case of log utility. This allows us to derive closed-form solutions for the model and eventually claim that even with very conservative degrees of risk aversion and intertemporal substitution, the model is able to match empirical regularities from the bond and equity markets.

Let $J(W_t, v_t)$ denote the agent's value function at time t where W_t and v_t correspond to current wealth and the conditional variance level respectively.[12] One can show that optimal

[10] See, for example, Hansen et al. (2006) and section 3, and especially proposition 4, in Skiadas (2003).

[11] By rewriting the return process under the risk neutral measure, one can also show that there is a perfect correlation between risk and uncertainty premia in our model.

[12] See the working paper version at http://neumann.hec.ca/pages/nicolas.vincent/ for a more detailed derivation of the policies and the value function. Anderson et al. (2003) and Maenhout (2004) also use similar formulations.

distortion h is given by

$$h_t = -\frac{1}{\theta}\left(J_{W,t}\sigma_{W,t} + J_{v,t}\sigma_v\sqrt{v_t}\right). \tag{2.10}$$

A few observations are in order at this point. First, since volatility is stochastic, the robustness correction h is state dependent: the robust agent derives the distorted conditional distribution in such a way that the reference conditional distribution first order stochastically dominates the chosen distorted conditional distribution. If it was not the case then there would be states of the world in which the robust agent would be considered optimistic. Also, the agent wants to maintain the optimal relative entropy penalty constant since θ is constant. In order to achieve this when conditional volatility is stochastic, the distortion has to be stochastic and increase with volatility. Second, the size of the distortion is inversely proportional to the penalty parameter θ: the distortion vanishes as $\theta \to \infty$. Third, whenever the marginal indirect utility and volatility of wealth (J_W and σ_W) are high, the agent becomes more sensitive to uncertainty and distorts the objective distribution more. Low levels of wealth imply large marginal indirect utility of wealth. These are states in which the agent seeks robustness more strongly. The second term in the parentheses corresponds to the effect of the state v on the distortion h. Since $J_v < 0$ for all reasonable parametrizations, the sign of σ_v dictates the optimal response of the agent. Consider the benchmark case when σ_v is positive. Following a positive shock, marginal utility falls as consumption rises, and volatility v increases. Therefore, the investment opportunity set deteriorates exactly when the agent cares less about it. Since the evolution of v serves as a natural hedge for the agent, he reduces the distortion h. The opposite occurs when $\sigma_v < 0$.

Given the choice of a distortion level, the optimal portfolio holding of the risky asset at time t, α_t, can be expressed in two equivalent forms, each emphasizing a different aspect of the intuition:

$$\begin{aligned} \alpha_t &= \frac{\mu^{\mathbb{Q}}_{R,t} - r_t}{\sigma^2_{R,t}} \\ &= \frac{\mu_{R,t} - r_t}{\sigma^2_{R,t}} + \frac{h_t}{\sigma_{R,t}}. \end{aligned} \tag{2.11}$$

The first line of equation (2.11) states that the demand for the risky asset is myopic: the agent only cares about the current slope of the mean-variance frontier. However, this slope is constructed using his subjective beliefs. From an objective point of view, the agent deviates from the observed mean-variance frontier portfolio due to his (negative) distortion to the mean h: he optimally believes the slope is lower and thus decreases his demand for the risky asset.

We posit the guess that the value function is concave (log) in the agent's wealth and affine in the conditional variance, which allows us to rewrite (2.10) as

$$h_t = -\frac{1}{\theta}\left(\frac{1}{\rho} + \delta_1 \sigma_v\right)\sqrt{v_t}. \tag{2.12}$$

Here, we see that the distortion, or the (negative of the) market price of model uncertainty is linear in the conditional volatility of the output growth rate. In equilibrium $\sqrt{v}$ is the conditional volatility of the consumption growth rate.

We can also rewrite (2.11) as

$$\alpha_t = \underbrace{\frac{1}{1+\frac{1}{\rho\theta}} \frac{\mu_{R,t} - r_t}{\sigma^2_{R,t}}}_{\text{Myopic demand}} - \underbrace{\frac{\frac{1}{\theta}}{1+\frac{1}{\rho\theta}} J_v \sigma_v}_{\text{Hedging demand}}.$$

The first element on the RHS corresponds to a variant of the usual trade-off in a log-utility setup between excess return compensation and units of conditional variance.[13] The second

[13] Note that the coefficient is not unitary, as in the usual log problem. The reason is that when introducing robustness, we effectively increase risk aversion, but maintain the unitary EIS. This pushes down the demand schedule for the risky asset.

element is the hedging-type component arising from uncertainty aversion. It is positive since $J_v \sigma_v > 0$ (recall the discussion of (2.10)), and larger in absolute terms the larger J_v or σ_v, ceteris paribus.

The consumption policy is unchanged when the agent seeks robust policies: $C = \rho W$. Here, robustness entails that the agent perceives the local expectations on the risky asset to be lower than the objective drift on the same asset. The substitution effect implies that the agent should invest less since the asset is expected to yield low return in the future. In contrast, the wealth effect predicts that he should consume less today and save instead. In the case of log utility, these two effects cancel each other. Consequently, the effect of robustness on the consumption policy is eliminated. Changing a log-agent's desire to be robust will only affect the risk free rate and the return on the risk free asset.

2.5 Robust Equilibrium

In this section we solve for the equilibrium prices of assets and discuss the pricing of the term structure of interest rates in our model. First, in our setup a robust equilibrium is defined as:

Definition 1 *A robust equilibrium is a set of consumption and investment policies/processes* (C, α) *and a set of prices/processes* (S, r) *that support the continuous clearing of both the market for the consumption good and the equity market* $(C = D, \alpha = 1)$ *and (2.8) is solved subject to (2.5), (2.2) and (2.6).*[14]

In equilibrium, since the agent consumes the output $(C = D)$ the local consumption growth rate and the local output growth rate are the same $(\mu_C = \mu)$. Also, the agent's equilibrium path of wealth is identical to the evolution of the price of the 'tree' since $\alpha = 1$. Therefore, $W = S$. Hence, $D = C = \rho W = \rho S$. As is usually the case with a log representative agent, not only the consumption wealth ratio is constant but so is the dividend-price ratio $\left(\frac{C}{W} = \frac{D}{S} = \rho\right)$.

[14] The same definition also appears in Maenhout (2004). Without stochastic volatility considerations, he also derives the equilibrium risk free rate and equity premium.

Consequently, robustness considerations do not affect the consumption policy and the pricing of the 'tree'. Instead, the implications of uncertainty aversion show up in the risk free rate and the way expectations are formed about growth rates or the return on the risky asset. The equilibrium risk free rate can be derived from (2.11)

$$
\begin{aligned}
r_t &= \rho + \mu_{C,t} + \sqrt{v_t} h_t - v_t \\
&= \rho + \mu - \phi v_t. \qquad (2.13)
\end{aligned}
$$

where

$$
\phi \equiv 1 + \frac{1}{\theta}\left(\frac{1}{\rho} + \delta_1 \sigma_v\right).
$$

As in a standard framework, a larger subjective discount rate ρ or higher future expected consumption growth both make the agent want to save less today and lead to a higher equilibrium real short rate. Also, higher consumption volatility activates a precautionary savings motive, so that the real rate must be lower to prevent the agent from saving. However, here the presence of θ implies a role for uncertainty aversion. Intuitively, robustness amplifies the effect of the precautionary savings motive ($h < 0$ when $\theta < \infty$), and thus lowers the equilibrium level of the short rate. In other words, the robust agent wants to save more than an expected utility agent and therefore needs a stronger equilibrium disincentive to save in the form of lower risk free rate.

The equilibrium local expected return on the risky asset can immediately be derived from (2.3) and the fact that $S = D/\rho$

$$
\begin{aligned}
dR_t &= \left(\mu_{D,t} + \rho\right) dt + \sigma_{D,t} dB_t \\
&= \left(\mu_{D,t} + \rho + h_t \sigma_{D,t}\right) dt + \sigma_{D,t} dB_t^{\mathbb{Q}}.
\end{aligned}
$$

The observed equity premium is[15]

$$\begin{aligned}
\mu_{R,t} - r_t &= \phi v_t = \underbrace{v_t}_{\text{Risk Premium}} + \underbrace{(\phi - 1)\, v_t}_{\text{Uncertainty Premium}} \\
&= cov_t\left(\frac{dC_t}{C_t}, dR_t\right) + \frac{1}{\theta}\left(\frac{1}{\rho} + \delta_1 \sigma_v\right) v_t.
\end{aligned}$$

The equity premium has both a risk premium and an uncertainty premium components. The former is given by the usual relation between the agent's marginal utility and the return on the risky asset. If the correlation between the agent's marginal utility and the asset return is negative, the asset commands a positive risk premium $\left[cov_t\left(\frac{dC_t}{C_t}, dR_t\right) > 0\right]$ and vice versa. The higher the degree of robustness (i.e., the smaller the parameter θ), the larger the uncertainty premium and the market price of model uncertainty. While a decrease in θ increases the equity premium, it also decreases the risk free rate through the precautionary savings motive. The EIS is independent of θ.[16]

We see that robustness can potentially account for both a high observed equity premium and low level of the risk free rate. What about the volatility of the risk free rate? Since we do not change the substitution motive, the only magnification is through the precautionary savings motive. Empirically v is extremely smooth and contributes very little to the volatility of r.[17]

[15] We use the qualifier 'observed' to emphasize again that what the agent treats as merely a reference model is actually the DGP. Therefore, anything under the reference measure is what the econometrician observe when he has long time series of data.

[16] Previous studies (e.g., Anderson et al. (2003), Skiadas (2003), Maenhout (2004)) have shown that when eliminating wealth effects from robustness considerations, a robust control economy is observationally equivalent to a recursive utility economy in the discrete time case (Epstein and Zin (1989), Weil (1990)) or to a stochastic differential utility (SDU) in the continuous time economy as in Duffie and Epstein (1992a) and Duffie and Epstein (1992b). Thus, our combined market price of risk and uncertainty can be viewed as an effective market price of risk in the SDU economy. The difficulty with such approach is that it requires implausibly high degrees of risk aversion. Another difficulty arises in the context of the Ellsberg paradox. Our approach assumes that agents do not necessarily know the physical distribution and want to protect themselves against this uncertainty.

[17] If we allow for a stochastic μ with positive correlation with v, fluctuations in v will be countered by movements in μ since they affect the risk free rate with opposite sign. In other words, if we allow the substitution effect and the precautionary motive to vary positively over time, the risk free rate can be very stable.

Next, we need to price the term structure of interest rates, the main object of this study.[18] Denote the intertemporal marginal rate of substitution (IMRS) process by Λ where $\Lambda_t \equiv e^{-\rho t}/C_t$. Ito's lemma allows us to characterize the dynamics of Λ as

$$\frac{d\Lambda_t}{\Lambda_t} = -r_t dt - \sqrt{v_t} dB_t^{\mathbb{Q}}, \tag{2.14}$$

where the drift is the (negative of) the short rate and the diffusion part is the market price of risk. It is then straightforward to price default-free bonds using (2.14). The excess expected return on a bond over the short rate is given by

$$\mathbb{E}_t^{\mathbb{Q}}\left(\frac{dp_t}{p_t}\right) - r_t dt = -\frac{d\Lambda_t}{\Lambda_t}\frac{dp_t}{p_t} \tag{2.15}$$

$$= \beta_1(\tau)\,\sigma_v v_t. \tag{2.16}$$

where the second line follows from our guess of an affine yield structure, and β_1 is positive and determines the cross section restrictions amongst different maturity bonds. The excess expected return is determined by the conditional covariance of the return on the bond and marginal utility, or alternatively, by the product of the market price of risk and the risk exposure of the bond. The sign of the risk premium is determined by the correlation of the output growth rate and the conditional variance, σ_v. In times of high volatility, the agent's decision to shift his portfolio away from the equity market and towards bonds leads to a rise in bond prices. When $\sigma_v > 0$, this implies that bonds pay well in good times and the risk premium is positive.

Moreover, the *observed* excess return that long term bonds earn over the short rate is not completely accounted for by the risk premium component. Under the objective measure we

[18] Our paper belongs to the vast literature on affine term structure models. The term structure literature is too large to summarize here but studies can be categorized into two strands - equilibrium and arbitrage free models. Our paper belongs to the former strand. The advantage of the equilibrium term structure models is mainly the ability to give meaningful macroeconomic labels to factors that affect asset prices. Dai and Singleton (2003) and Piazzesi (2003), for example, review in depth the term structure literature.

have,

$$\frac{dp(\tau; v_t)}{p(\tau; v_t)} = \left[r_t + \underbrace{\beta_1(\tau)\,\sigma_v v_t}_{\text{Risk Premium}} + \underbrace{\beta_1(\tau)\,\sigma_v v_t\,(\phi - 1)}_{\text{Uncertainty Premium}} \right] dt + \beta_1(\tau)\,\sigma_v \sqrt{v_t} dB_t. \tag{2.17}$$

In the presence of uncertainty aversion, there is an uncertainty premium that drives a wedge between the return on a τ-maturity bond and the short rate. The more robust the agent, the larger the market price of uncertainty is in absolute terms (i.e., ϕ is larger so $-h = (\phi - 1)\sqrt{v}$ is larger). Also, higher conditional variance increases the uncertainty premium since the agent distorts the mean of the objective model more. In other words, higher σ_v also increases the uncertainty exposure of the asset. In section 4 we discuss the intuition behind the predictions of the model, and especially the role robustness plays in our context.

Finally, the yield on a given bond is simply an affine function of the conditional variance

$$\mathcal{Y}(\tau; v_t) = -\frac{1}{\tau} \ln p(\tau; v_t).$$

The two extreme ends of the yield curve are $\lim_{\tau \to 0} \mathcal{Y}(\tau; v_t) = r_t$ and $\lim_{\tau \to \infty} \mathcal{Y}(\tau; v_t) = \rho + \mu - a_0 \bar{\beta}_1$. Thus the spread is

$$\lim_{\tau \to \infty} \mathcal{Y}(\tau; v_t) - \lim_{\tau \to 0} \mathcal{Y}(\tau; v_t) = -a_0 \bar{\beta}_1 + \phi v_t,$$

where $\bar{\beta}_1$ is a function of the model parameters.

3 Model Estimation

Since the model permits closed-form expressions for first and second moments we use the generalized method of moments (GMM) to estimate our model parameters (Hansen (1982)). Our procedure is similar to the one used by, for example, Chan et al. (1992). Our approach is to exploit mainly the time series restrictions to estimate the structural parameters. We do not focus on the cross sectional restrictions of the model as in Longstaff and Schwartz (1992) and Gibbons and Ramaswamy (1993). Since we have a single factor model, yields are perfectly correlated. Therefore, including cross sectional restrictions may reduce the power of the overidentifying restrictions in small samples. We use our point estimates to generate the model's implied real yield curve and compare it to its empirical counterpart. In that sense, our approach is more ambitious. It is important to note that since our model only makes statements about the real economy, we are careful to use data denominated in real terms. The description of the data is relegated to Appendix A.[19]

We need to estimate 6 parameters $\{a_0, a_1, \mu, \rho, \theta, \sigma_v\}$. We form orthogonality conditions implied by the model using the following notation

$$\begin{aligned} Y_{t+1} &\equiv \left[\Delta \mathcal{Y}(1; v_{t+1}), R_{t+1}, \frac{\Delta C_{t+1}}{C_t}, \mathcal{Y}(4; v_t) - \mathcal{Y}(1; v_t)\right], \\ X_t &\equiv \mathcal{Y}(1; v_t), \\ Z_t &\equiv \left[1, \mathcal{Y}(1; v_t), R_t, \frac{\Delta C_t}{C_{t-1}}\right], \end{aligned}$$

where Y_{t+1} is observed at time $t+1$ and contains the change in the one-quarter real yield $(\Delta \mathcal{Y}(1; v_{t+1}))$, the realized real aggregate market return (R_{t+1}), the realized real aggregate consumption growth rate $\left(\frac{\Delta C_{t+1}}{C_t}\right)$ and the real spread between the 1-year and 3-months real yields $(\mathcal{Y}(4; v_t) - \mathcal{Y}(1; v_t))$. X_t is the explanatory factor. Even though conditional variance is not directly observable in the data it is theoretically an affine function of the short rate (or any other real yield with arbitrary maturity). Therefore, we use the short rate (3-months real yield) as an observable that completely characterizes the behavior of the conditional variance. Last, we use lagged 3-months real market return and realized consumption growth rate as instruments

[19] A few studies, for example Brown and Schaefer (1994) and Gibbons and Ramaswamy (1993), also use real data to estimate a term structure model. However, they do not draw restrictions from the equity market and consumption data and their preferences assumption is standard which implies that the equity premium and risk free rate puzzles are still present in the models they estimate.

Also, some authors have used nominal data to estimate real models (e.g., Brown and Dybvig (1986)).

in the vector Z_t.

The stacked orthogonality conditions are given in m

$$\begin{aligned} u_{1,t+1} &\equiv Y_{t+1} - \mu_{Y,t}|X_t, \\ u_{2,t+1} &\equiv diag\left(u_{1,t+1}u'_{1,t+1} - \sigma_{Y,t}\sigma'_{Y,t}|X_t\right), \\ m_{t+1} &\equiv \begin{bmatrix} u_{1,t+1} & u_{2,t+1} \end{bmatrix} \otimes Z_t. \end{aligned}$$

We draw first and second moment restrictions. $\mu_{Y,t}|X_t$ and $\sigma_{Y,t}|X_t$ have the parametric forms implied by the model and are affine in X_t.

4 Empirical Results

In this section we first present the parameter estimates over various samples before analyzing the empirical fit of our model. We pay particular attention to discussing how the interaction of robust control and stochastic volatility allows us to address a number of asset pricing facts.

4.1 Point Estimates of the Structural Parameters

Table 4.1 presents the point estimates over different time periods. Aside from the robustness parameter θ, all coefficients are immediately interpretable. All parameters are statistically different from zero. Also, the model is not being rejected according to the J-test. We will explain the DEP's column later.

Note that both μ (equal to the average real aggregate quarterly consumption growth rate) and ρ are stable over different samples.[20] What is obvious from Table 4.1 is that the estimation

[20] In results available from the working paper version, we also estimated the model without imposing the volatility of consumption growth. What we find is that the parameters μ, ρ, a_0 and a_1 are mostly unaffected. The penalty parameter θ, however, is higher: without imposing the smooth consumption process , the implied pricing kernel (SDF) is much more volatile and less robustness is needed to justify the observed asset prices. Similarly, the implied evolution of v is much more volatile leading to a higher value for σ_v.

Table 4.1: Model estimation with consumption volatility over different time intervals. The data is in quarterly frequency and in quarterly values. 6 parameters are estimated using an iterated GMM. There are 8 moments and 4 instruments that produce 32 orthogonality conditions. T is the number of observation in each estimation. Robust t-statistics are indicated below each point estimate. The standard error are corrected using the Newey-West procedure with 4 lags. p-val is the p-value for the J-test statistic distributed χ^2 with 26 degrees of freedom. The DEP column reports the detection error probabilities.

Period	T	a_0	a_1	μ	ρ	θ	σ_v	J-stat/P-val	DEP
Q2.52-Q4.06	218	0.0005	-0.1951	0.0052	0.0145	9.5730	0.0189	33.5147	4.41%
		6.7904	-10.1242	29.0936	7.6824	3.7579	7.7160	14.77%	
Q1.62-Q4.06	180	0.0004	-0.1594	0.0051	0.0125	14.9578	0.0175	28.0792	11.01%
		6.9588	-10.7284	28.2098	6.3202	2.7028	5.5972	35.46%	
Q1.72-Q4.06	140	0.0004	-0.1611	0.0048	0.0148	11.8660	0.0149	22.3470	10.86%
		6.3636	-10.3456	28.4321	6.1634	2.5565	5.1858	66.96%	
Q1.82-Q4.06	100	0.0004	-0.1314	0.0054	0.0208	6.6591	0.0082	17.3000	7.31%
		6.2872	-6.6014	34.0290	8.7880	3.7862	7.4882	89.97%	
Q1.90-Q4.06	68	0.0003	-0.0997	0.0050	0.0172	9.0779	0.0064	12.2261	14.49%
		8.1423	-9.7387	36.4052	10.5310	5.7235	10.3818	98.98%	
Q2.52-Q4.81	118	0.0007	-0.2528	0.0050	0.0115	13.4790	0.0431	20.2067	17.11%
		7.1360	-13.8242	20.8878	5.6148	2.7839	6.0434	78.17%	
Q2.52-Q4.89	149	0.0006	-0.2232	0.0053	0.0140	10.6126	0.0290	24.2714	10.10%
		6.2481	-9.5348	24.3649	6.3682	3.2241	6.9336	56.04%	

procedure detects mostly high frequency movements and not the slow moving component in consumption growth volatility identified by Bansal and Yaron (2004). Hence, it appears that the high-frequency component from the market data dominates in the full-model estimation. We will discuss further our parameter estimates in the context of the fit of the model.

4.2 Theoretical and Empirical Moments

Table 4.2 presents a comparison of model-implied and empirical moments over different time spans for the main focus of our study, the bond market. In general, the model fares very well: it can generate average bond yields very similar to those in the data. In addition, it is able to fit the autocorrelation of yields as well as the return from a strategy of holding a one-year bond for three quarters.

Table 4.3 shows some key moments for the equity and consumption data. While the model-implied moments are comparable to their empirical counterparts, in fact we may say that our

Table 4.2: Empirical and theoretical bond market moments (with consumption volatility restriction). The period column represents the time interval of the data that is used to estimate the model. The data is in quarterly frequency with quarterly values. T is the number of quarterly observations used to estimate the model. Columns with the number (1) present the empirical moments. Empirical moments computed with the data and theoretical moments are implied by the estimated model. Columns with the number (2) present the theoretical moments. The theoretical moments were generated using $1,000$ replications of the economy that was calibrated using the estimated parameters over the corresponding period. Robust standard errors are given below each moment. The standard errors were corrected using the Newey-West procedure with 4 lags. The standard errors for the theoretical moments were computed over the $1,000$ replications. All moments, aside from the autocorrelations, are given in % values. $\mathcal{Y}_{3m}$, $\mathcal{Y}_{1y}$, and $\rho(\mathcal{Y}_{3m})$ are the real 3 month yield, real 1 year yield and the first order autocorrelation coefficient of the real 3 month yield, respectively. The last column reports real holding period return for buying a one year to maturity bond and selling it after three quarters.

Period	T	$\mathcal{Y}_{3m}$		$\mathcal{Y}_{1y}$		$\rho(\mathcal{Y}_{3m})$		$\ln\left[\frac{p(1;v_{t+3})}{p(4;v_t)}\right]$	
		(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
$Q2.52 - Q4.06$	218	1.531	1.838	2.250	2.570	0.843	0.863	2.465	2.802
		0.263	0.083	0.234	0.055	0.060	0.004	0.274	0.047
$Q1.62 - Q4.06$	180	1.749	2.191	2.241	2.614	0.835	0.878	2.394	2.750
		0.292	0.011	0.279	0.008	0.071	0.000	0.321	0.007
$Q1.72 - Q4.06$	140	1.686	2.216	2.214	2.680	0.830	0.875	2.359	2.847
		0.362	0.080	0.351	0.058	0.075	0.004	0.406	0.047
$Q1.82 - Q4.06$	100	2.190	2.584	2.742	3.071	0.862	0.891	2.932	3.232
		0.366	0.109	0.378	0.084	0.062	0.003	0.443	0.073
$Q1.90 - Q4.06$	68	1.666	1.914	2.015	2.233	0.897	0.890	2.118	2.346
		0.422	0.135	0.371	0.111	0.058	0.006	0.420	0.110
$Q2.52 - Q4.81$	118	0.965	1.572	1.865	2.471	0.804	0.784	1.976	2.764
		0.330	0.210	0.247	0.118	0.096	0.010	0.275	0.090
$Q2.52 - Q4.89$	149	1.469	2.207	2.358	2.996	0.823	0.825	2.586	3.257
		0.328	0.143	0.292	0.088	0.077	0.007	0.348	0.070

model is performing *too* well: it generates an equity premium *larger* than in the data, while still matching the low consumption growth rate as well as the bond market facts. The same conclusion seems to hold over different time horizons. Note, however, that since we are imposing consumption growth volatility, the model compromises on the implied market return volatility being somewhere between the empirical consumption growth rate volatility and the empirical market return volatility.[21]

Table 4.3: Empirical and theoretical equity and goods market moments (with consumption volatility restriction). The period column represents the time interval of the data that is used to estimate the model. The data is in quarterly frequency with quarterly values. T is the number of quarterly observations used to estimate the model. Columns with the number (1) present the empirical moments. Empirical moments computed with the data and theoretical moments are implied by the estimated model. Columns with the number (2) present the theoretical moments. The theoretical moments were generated using $1,000$ replications of the economy that was calibrated using the estimated parameters over the corresponding period. Robust standard errors are given below each moment. The standard errors were corrected using the Newey-West procedure with 4 lags. The standard errors for the theoretical moments were computed over the $1,000$ replications. All moments are given in % values. μ_R, μ_C, σ_R, and $\mathcal{Y}_{3m}$ are the real return on the market (including dividends), real growth rate of consumption, volatility of real aggregate market return and real 3 month yield, respectively.

Period	T	μ_R		$\mu_R - \mathcal{Y}_{3m}$		σ_R		μ_C	
		(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
$Q2.52 - Q4.06$	218	8.690	12.346	7.069	10.354	16.445	9.861	2.096	2.099
		1.120	0.078	1.133	0.118	1.112	0.081	0.086	0.075
$Q1.62 - Q4.06$	180	7.440	10.589	5.606	8.245	17.092	9.708	2.074	1.526
		1.233	0.111	1.239	0.160	1.254	0.172	0.093	0.110
$Q1.72 - Q4.06$	140	7.973	11.023	6.198	8.645	17.418	10.224	1.927	1.524
		1.456	0.093	1.457	0.132	1.440	0.111	0.101	0.092
$Q1.82 - Q4.06$	100	11.159	12.294	8.807	9.501	16.579	10.507	2.199	1.588
		1.544	0.151	1.532	0.204	1.478	0.168	0.088	0.152
$Q1.90 - Q4.06$	68	9.110	10.033	7.342	7.991	15.984	10.068	2.028	1.349
		1.920	0.178	1.894	0.241	1.841	0.136	0.100	0.176
$Q2.52 - Q4.81$	118	6.787	15.933	5.777	14.187	16.223	9.901	1.988	1.802
		1.541	0.152	1.592	0.256	1.624	0.300	0.138	0.141
$Q2.52 - Q4.89$	149	8.498	14.082	6.944	11.665	16.653	9.897	2.128	1.826
		1.388	0.117	1.418	0.188	1.379	0.199	0.116	0.114

Next, we provide intuition about the role of robustness in matching bond market moments.

[21] We also tried to reestimate the model without imposing consumption growth volatility. The results were qualitatively similar. Not surprisingly, this version was better able to match the aggregate market return volatility, at the expense of higher consumption growth rates. Results for the yield curve were however relatively unchanged.

4.2.1 Bond Returns and Upward Sloping Yield Curve

Results for the 3-months and 1-year real yields reveal that the model is doing a good job at reproducing the shape of the short end of the yield curve. To get a better idea of its ability to replicate the term structure of interest rates, the top panel in Figure 4.1 presents estimation results over the years $1997-2006$. During this period TIPS bonds were traded in the U.S. and thus provide a good proxy to real yields. The solid line is the average level of the yield curve over this period with 95% confidence bands. The dot-dashed line is the model-implied average yield curve. Note that we only impose two bond market restrictions in the estimation procedure and yet the model can imitate reasonably well the shape of the entire yield curve.

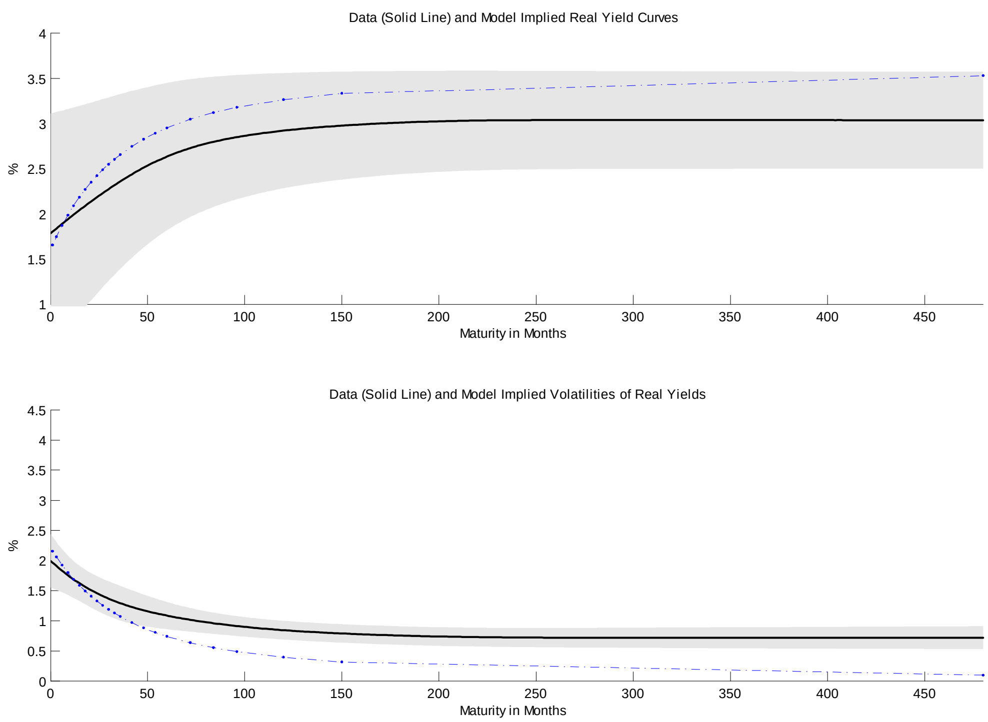

Figure 4.1: Top panel: average real yield curve extracted from the TIPS data from $M1.97-M12.06$ (solid line) with 95% confidence bands with Newey-West (12 lags) correction. Model implied average yield curve (dot-dashed line). The model is estimated over the same period as the empirical yield curve. Bottom panel: empirical term structure of unconditional volatilities of the TIPS data (solid line). with 95% confidence bands with Newey-West (12 lags) correction. The model is estimated over the same period as the empirical yield curve.

Two elements allow our model to match the term structure of interest rates. First, recall from expression (2.17) that robustness introduces an uncertainty premium in addition to the usual risk premium through the precautionary savings motive. Both premia are positive as long as $\sigma_v > 0$, (i.e. shocks to consumption growth and volatility are positively correlated). The magnifying role of robustness means that one can match the excess return on long term bonds relative to short term bonds with a moderate amount of stochastic volatility.

Second, with robustness the (perceived) evolution of v under the distorted measure $\mathbb{Q}$ is different from the evolution of v under the objective measure $\mathbb{P}$ in two respects. In Figure 4.2 we plot the objective and perceived impulse response functions for the conditional variance v following a shock. Note that, unlike a rational expectations agent, the robust agent is on average wrong about the future evolution of v: uncertainty aversion leads him to distort his beliefs such that, on average, he expects the conditional variance to decrease over time.[22] In Section 6 we argue that this bias is not statistically unreasonable. This observation provides an additional channel for an upward-sloping yield curve: through the precautionary savings motive, lower future volatility results in higher expected yields.

Another manifestation of the forces just described is apparent in the last column of Table 4.3. It captures the return of a strategy in which the agent buys a 1-year bond and sells it after 3 quarters. Backus et al. (1989) point to the difficulty of representative agent models to account for both the sign and magnitude of holding period returns in the bond market. Note that we did not impose any holding period returns conditions in the estimation procedure and

[22] To see this algebraically, write (2.2) under both measures

$$\begin{aligned} dv_t &= -\kappa_v \left(v_t - \bar{v}\right) dt + \sigma_v \sqrt{v_t} dB_t \\ &= -\kappa_v^{\mathbb{Q}} \left(v_t - \bar{v}^{\mathbb{Q}}\right) dt + \sigma_v \sqrt{v_t} dB_t^{\mathbb{Q}}. \end{aligned} \tag{4.1}$$

Here, κ_v is the velocity of reversion and $\bar{v}$ is the steady state of v, both under the reference measure. However, the subjective velocity of reversion is

$$\kappa_v^{\mathbb{Q}} = \kappa_v - \sigma_v \left(1 - \phi\right) > \kappa_v \tag{4.2}$$

and the subjective steady state is

$$\bar{v}^{\mathbb{Q}} = \frac{\kappa_v}{\kappa_v^{\mathbb{Q}}} \bar{v} < \bar{v}. \tag{4.3}$$

where the inequality follows from our finding that $\kappa_v < \kappa_v^{\mathbb{Q}}$ across all samples.

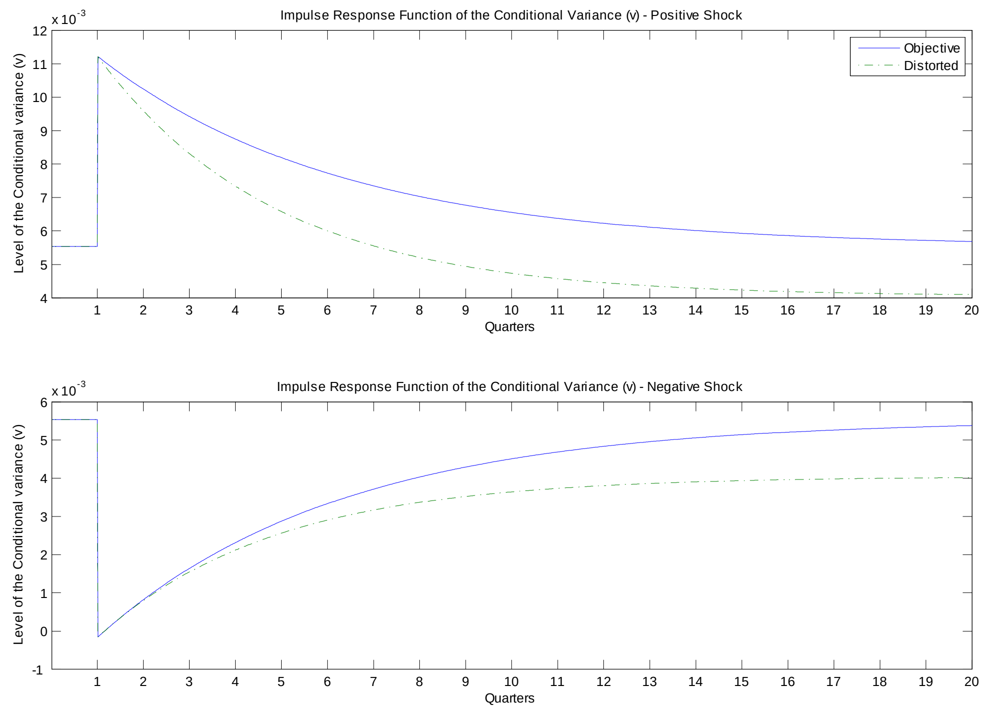

Figure 4.2: Biased expectations. Using the parameters estimated over the entire period $Q2.52 - Q4.06$, the figure shows the impulse response function of the conditional variance to a positive and negative shocks. The solid line represents the objective evolution of v and the dashed line represents what the robust agent believes the evolution of v is going to be.

yet the model captures the returns dynamics well.

The bottom panel of Figure 4.1 depicts the term structure of the volatilities of yields. Clearly, the model can replicate the downward slope due to the mean reversion in the estimated conditional variance process, as discussed earlier. The impression is that the procedure anchors the implied first and second moments of the 1-year yield to its empirical counterpart, but it is still doing a good job in approximating the entire curve.

4.2.2 Slope of the Yield Curve and Persistence of Yields

Traditionally, one-factor models encounter an inherent difficulty in trying to account simultaneously for the rapidly declining slope of the yield curve (i.e., strong convexity of the slope of the yield curve) and the high persistence of yields. Time-series evidence implies that interest-rate shocks die out much more slowly than what is implied from the rapidly declining slope of the average yield curve (Gibbons and Ramaswamy (1993)).

Even though we present a one-factor model, we can account for these two facts with a single parametrization. Figure 4.1 shows that the agent prices the yield curve as if shocks to v die out fast. However, Table 4.2 confirms that the model can still match the persistence of the short rate $(\rho(\mathcal{Y}_{3m}))$.[23]

The key for understanding the success of our model on that dimension lies in expression (4.2). The agent believes that the conditional variance reverts to its steady state *faster* than under the objective measure $\left(\kappa_v^{\mathbb{Q}} > \kappa_v\right)$. Since yields are affine functions of the conditional variance of consumption growth, they inherit the velocity of reversion of v under the objective model. In other words, the persistence of yields is measured ex-post and is solely determined by the objective evolution of v without any regard to what the agent actually believes.

At the same time, the slope of the yield curve (or the pricing of bonds) is completely determined by what the agent believes the evolution of v will be, as discussed earlier. If $\kappa_v^{\mathbb{Q}}$ is substantially larger than κ_v, the slope of the yield curve can flatten at relatively short horizons, reflecting the beliefs of the agent that v will quickly revert to its steady state level. Since the agent persistently thinks that $\kappa_v^{\mathbb{Q}} > \kappa_v$ the slope can be on average rapidly declining.

4.2.3 Spread and Level of Yields

In quarterly data over the sample $52.Q2 - 06.Q4$ the correlation between the level and slope of the real yield curve is -0.5083 with standard errors of 0.0992 (where the slope is the difference between the 1-year and 3-months yields). This finding is robust over different time intervals

[23] The term structure literature usually identifies 3 factors that account well for most of the variation in the yield curve (Litterman and Scheinkman (1991)): level, slope and curvature. The level slope is very persistent and, thus, accounts for most of the observed persistence of yields. Also, note that we do not impose this restriction in our estimation and yet the model is able to match this moment with high accuracy.

and different frequencies. The model can account for this fact in the following way[24]: recall that a positive shock to conditional volatility lowers yields. Also note that yields are perfectly (positively) correlated since all of them are an affine function of the same factor. However, short yields are more sensitive to conditional volatility shocks. To understand why, it helps to think about the mean reversion of the conditional variance (the ergodicity of its distribution). The effect of any shock is expected to be transitory. The full impact of the shock happens at impact and then the conditional variance starts reverting back to its steady state. Therefore, the effect of, say, a positive shock is expected to dissipate and yields are expected to start to climb back up. This expected effect is incorporated into long term yields immediately. Short yields in the far future are almost unaffected by the current shock since it is expected that the effect of the shock will disappear eventually. Since long term yields are an average of future expected short yields plus expected risk premia, they tend to be smoother than short term yields.

The expected risk premium is also a linear function of the state, and thus inherits its mean reversion. Therefore, the expected risk premium in the far future is also smoother than the risk premium in the short run. This also contributes to the rotation of the yield curve: since the short end of the yield curve is very volatile relative to the long end (recall Figure 4.2), whenever yields decrease, the spread increases (or become less negative, depending on the initial state). The opposite also holds true.

5 Additional Evidence

The success of our model in replicating numerous moments for both the equity and bond markets rests on two ingredients: (1) state-dependent volatility of consumption growth, and (2) a positive correlation between shocks to consumption growth and volatility ($\sigma_v > 0$). In this

[24] We explain the intuition through the time variation of the conditional volatility of consumption growth rate. One can alternatively use the substitution channel and focus on time variation in expected consumption growth rate.

subsection we provide direct empirical evidence about the level and behavior of the conditional variance of real aggregate consumption growth.

5.1 ARMAX-GARCH Real Consumption Growth Rate

We start with a simple univariate time series parametric estimation. The model we are fitting to the consumption growth process is an ARMAX$(2,2,1)$ model and a GARCH$(1,1)$ to the innovations process:

$$\begin{aligned} A(L)\frac{\Delta C_t}{C_{t-1}} &= c + B(L)R_{t-1} + C(L)\eta_{C,t}, \qquad (5.1)\\ \eta_{C,t+1} &= \sigma_{C,t}\varepsilon_{C,t+1}, \quad \varepsilon_{C,t} \sim N(0,1), \\ D(L)\sigma_{C,t} &= \omega + F(L)\eta^2_{C,t}, \end{aligned}$$

where A, B, C, D, F are polynomials of orders $2,1,2,1,1$ respectively, in lag operators. $\frac{\Delta C_t}{C_{t-1}}$, R_t, η_t are, respectively, the realized real consumption growth rate at time t, the real return on the aggregate market index at time $t-1$, and an innovation process with time-varying variance. In Figure 5.1 we plot the GARCH volatility estimates for both real aggregate consumption growth rate and the real return on the aggregate stock market. We also plot a measure of realized volatility for both consumption growth and market return series that we obtain by fitting an ARMA$(2,2)$ to the original data and then use the square innovations to construct the realized variance series. The sample period is $Q2.52 - Q4.06$.

First, there seems to be evidence of what has been dubbed as the 'Great Moderation' (e.g., Stock and Watson (2003)). It is clear that consumption growth volatility has slowly declined over the sample period but the volatility of the market return did not. This pattern is apparent in both measures of conditional volatility.

Second, it appears that there are both high frequency (business cycle) fluctuations and a very low frequency stochastic trend in consumption growth volatility. Panel A of Table 5.1

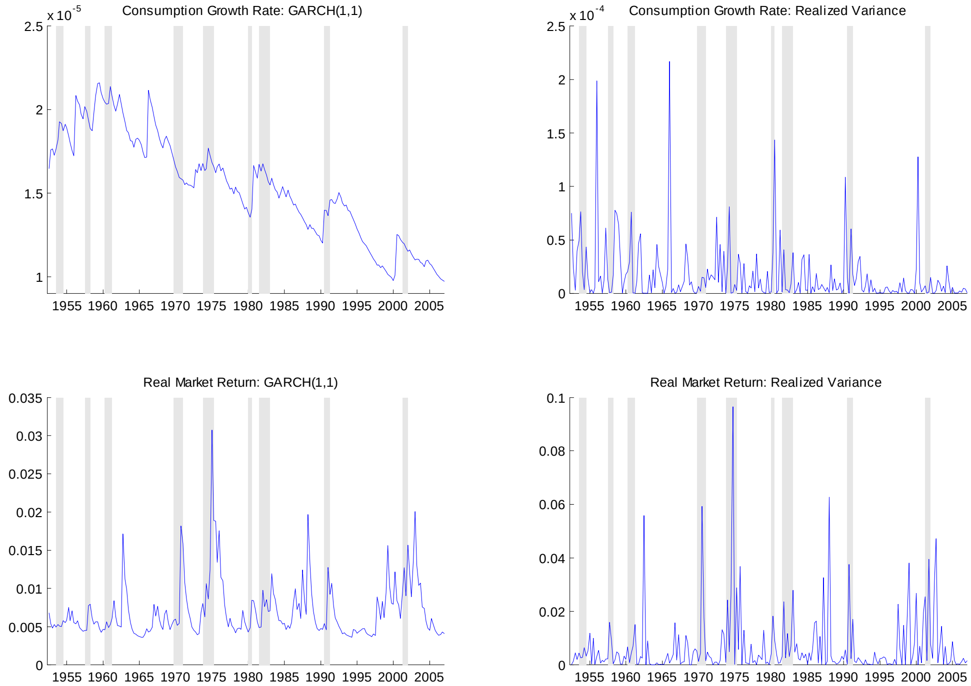

Figure 5.1: ARMAX-GARCH estimation for both real consumption growth rate and real aggregate market return. We fit model (5.1) and present the GARCH estimates for the conditional variance of real consumption growth rate and real aggregate market return in the left panel. The right panel present the square innovations from an ARMX specification to real consumption growth and real aggregate market return. The quarterly data is $Q2.52 - Q4.06$. The gray bars are contraction periods determined by the NBER.

presents the implied reversion coefficient and half life derived using the autoregressive coefficient from the GARCH estimated conditional variance series.[25] For comparison purposes, Panel A of Table 5.1 presents the half life of the volatility shock process implied by our earlier estimation results. We also present in that panel the perceived half life by the robust agent. Expression (4.2) shows that the perceived velocity of mean reversion is faster than the physical speed at which shocks to volatility die out. In general, the point estimates imply that shocks to volatility die out relatively fast. These results confirm that without forcing asset market restrictions on the consumption series, we observe a very slow moving process for conditional variance. At the

[25] Our point estimates correspond to quarterly data. In general, with data sampled at quarterly frequency one can map an autoregressive coefficient to a coefficient governing the speed of reversion as our κ_v. Let $\hat{\alpha}$ denote the autoregressive coefficient. Then, the quarterly speed of reversion coefficient $\kappa_v = -\ln(\hat{\alpha})$ and the half life is $\ln(2)/\kappa_v$.

same time, the conditional variance of the aggregate market return is much less persistent. The general estimation procedure results in panel A are, to some extent, a combination of these two effects.[26]

Table 5.1: Panel A: Point estimates of the velocity of reversion coefficient and the implied half life (in quarters) of the conditional variance process. Objective referes to the physical rate in which the conditional variance gravitates to its steady state. Distorted referes to the rate in which the robust agent believes the conditional variance gravitates to its steady state. These point estimates are from the estimation procedure that imposes the volatility of real aggregate consumption growth rate as a moment condition. Panel B: implied reversion coefficients and half lifes (in quarters) for the conditional volatility of consumption growth rate and aggregate market return derived from the GARCH procedure. The consumption growth rate mean is modeled as an ARMAX(2,2,1) and the aggregate market return is modeled as ARMA(2,2).

Panel A:	$Q2.52 - Q4.06$			
	Estimate		Half Life (Q)	
Consumption	0.010		68.373	
Market	0.4069		1.704	
Panel B:	$Q2.52 - Q4.06$		$Q1.90 - Q4.06$	
	Estimate	Half Life (Q)	Estimate	Half Life (Q)
Objective	0.1951	3.553	0.0997	6.952
Distorted	0.2994	2.315	0.1343	5.161

The recent 'long run risks' literature usually calibrates asset pricing models with a highly persistent conditional variance process.[27] For example, Bansal and Yaron (2004) assume that the autoregressive coefficient (with monthly frequency data) in the conditional variance of the consumption growth process is 0.987.[28] This number implies a half life of 13.24 quarters, which is almost 4 times higher than the number we obtain in our empirical results. As explained earlier, this difference is driven largely by the inclusion of equity and bond markets in our set of moments. What we show in this paper is that robust decision making coupled with state

[26]We conduct this comparison only for the entire period $Q2.52 - Q4.06$ since we want to examine evidence concerning very low frequency components. Even our longest sample is somewhat short to conveniently detect the slow moving component. We believe that shorter samples will make the detection exercise impossible.

[27]Bansal and Yaron (2004) find that introducing a small highly persistent predictable component in consumption growth can attenuate the high risk aversion implications of standard asset pricing models with recursive utility preferences. However, this persistent component is difficult to detect in the data. Croce et al. (2006) present a limited information economy where agents face a signal extraction problem. Their model addresses the identification issues of the long run risk component. Hansen and Sargent (2010) is another example for the difficulty in identifying the long run risk component. However, in addition to a signal extraction problem, their agent seeks robust policies and consequently his estimation procedure is modified.

[28]See table IV in Bansal and Yaron (2004).

dependent volatility requires moderate levels of persistence in the conditional variance of the consumption growth process. Recall that we assume a constant drift in consumption growth. If we assume a stochastic and highly persistent μ, as in Bansal and Yaron (2004), we would need to worry about the volatility of the risk free rate. In other words, if the substitution effect channel is very persistent and the precautionary savings motive is much less persistent, the short rate can potentially be very volatile. If shocks to μ were to die out much slower than shocks to v, the ergodic distribution of the short rate would be very volatile. In that sense, we might be able to reconcile our results with the calibration exercise of Bansal and Yaron (2004) if we assumed an expected consumption growth rate process.

Our hypothesis is that higher frequency fluctuations are channeled through the asset market while there are other aspects which we do not identify that contribute to the low frequency fluctuations. In other words, when we estimate the full model, the effect of the equity and bond market restrictions is reflected in the implied persistency of the conditional variance process. Here, we use the Hodrick-Prescott filter with parameter 1600 to disentangle these two components of consumption growth volatility. Figure 5.2 presents this result and makes clear that the decline in the low frequency component started in the $'60$, before the Great Moderation.[29]

We also use the volatility estimates to explain asset prices (see also Chapman (1997), Bansal and Yaron (2004), Bansal et al. (2005)). In particular, in figure 5.3 we examine the dynamic

[29] In our model it is hard to make 'conditional' statements about the economy, mainly because we modeled a constant drift to the consumption growth rate process. It is obviously interesting to think about the correlation structure of expected consumption growth rate and the conditional variance process. Empirically, there is evidence that suggests that interest rates are procyclical (e.g., Donaldson et al. (1990)) and volatility is either countercyclical or at least slightly leads expected growth rates which are believed to be countercyclical (e.g., Whitelaw (1994)). Our conditional variance process is assumed to correlate positively with realized consumption growth rate. Also, the conditional variance correlate negatively with interest rates. In this sense, variance and real interest rates behave as in the data. If, for example, expected growth rate correlate negatively with realized consumption growth rates, they will correlate negatively with the conditional variance. In that case, a positive shock to consumption growth rate will have a double negative effects on real interest rates. Expected growth rates will be low and thus the substitution effect will make equilibrium real interest rates lower. At the same time, conditional variance will be higher and the precautionary savings motive will push the equilibrium real interest rate even lower. Also, Chapman (1997) documented the strong positive correlation of real yields and consumption growth rate when excluding the monetary experiment period of $1979 - 1985$.

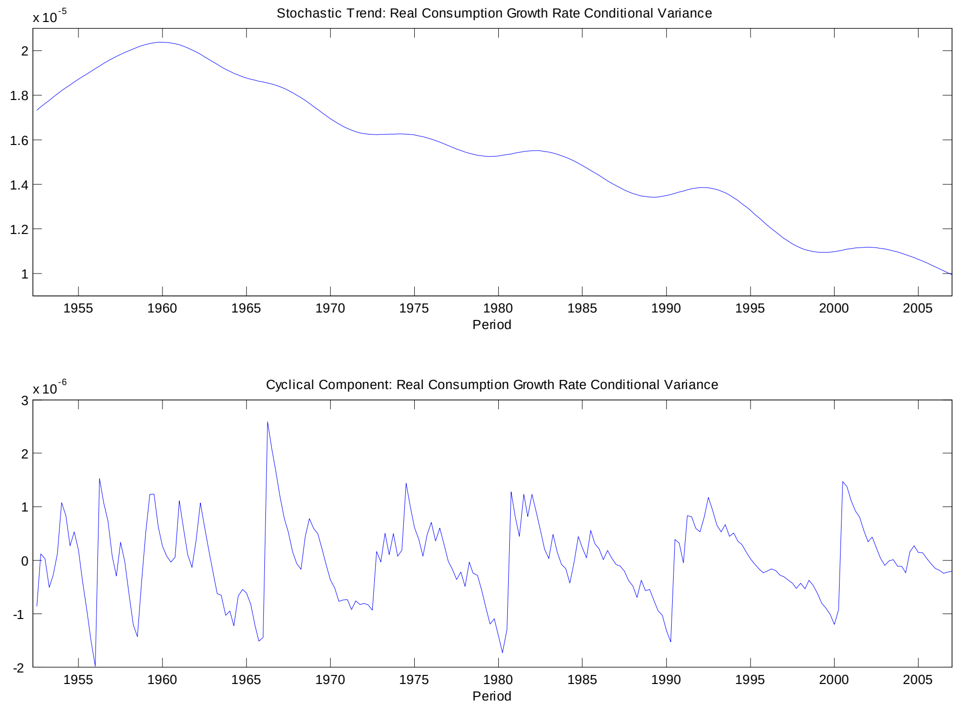

Figure 5.2: HP-flitered conditional variance of real consumption growth rate derived from an ARMAX-GARCH estimation in (5.1). The top panel presents the low frequency trend and bottom panel presents the cyclical component. The HP-filter parameter is 1600. The quarterly data is over the period $Q2.52 - Q4 - 06$.

cross correlation patterns between consumption growth volatility obtained from the GARCH estimation in (5.1) and the spread between the real 1-year real yield and the real 3-months real yield.

These patterns agree with the model's predictions. We know that shorter maturity yields respond more than longer maturity yields to a volatility shock. This result is mainly due to the ergodicity of the state variable that affect yields. If the state is assumed to revert back to a known steady state, we expect the longer yield to have a smaller response to contemporaneous shocks. Note that we do not identify the type of shock in this exercise. We merely observe a shock that happens to affect both consumption growth volatility and the bond market.

The second result is the sign response of the yields to a volatility shock. When conditional volatility increases we see that yields decrease. From the precautionary savings motive effect

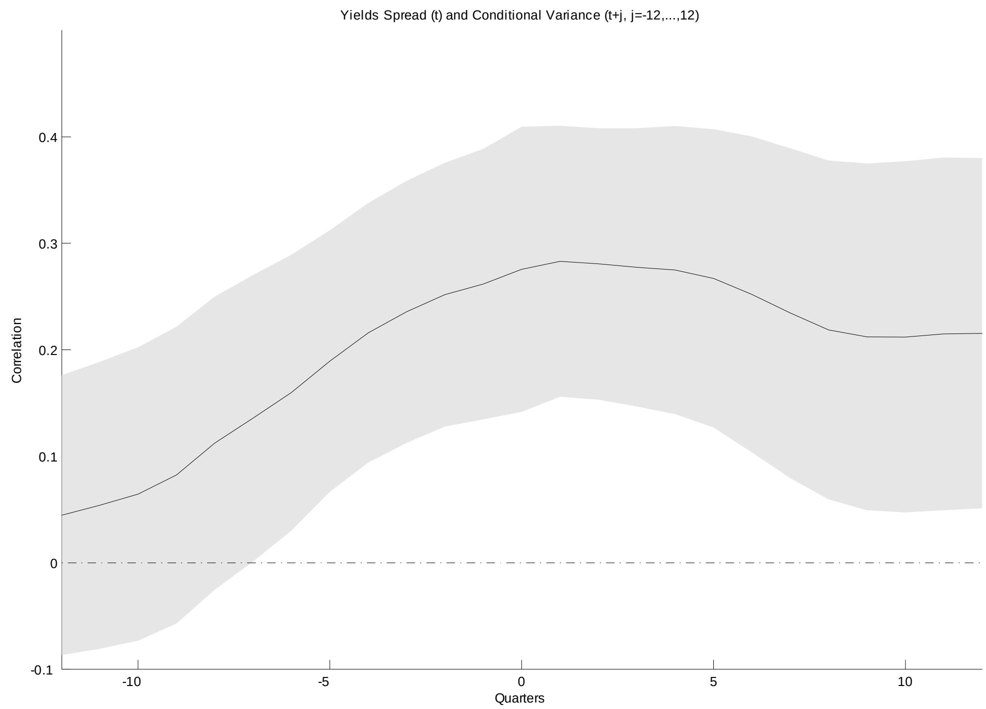

Figure 5.3: Dynamic cross-correlation between real consumption growth rate volatility and the real spread between the 1 year and 3 months yields. The quarterly data covers the period $Q2.52 - Q4.06$.

we do expect such response. Since in our model ambiguity aversion amplifies the precautionary savings motive, we expect this channel to play an important role when linking consumption growth volatility and yields. When combining these two results, we expect the spread to increase with a volatility shock. In other words, on average, the yield curve rotates when a shock to volatility occurs.

There are three caveats to these results. First, the upper left panel in Figure 5.1 depicts the behavior of the conditional variance of real consumption growth. One can argue that the series exhibit a non-stationary behavior. If this is the case, then the GARCH process is potentially misspecified. Given the slow-moving component we identified, it is hard to convincingly argue against such hypothesis. Second, our macro data is sampled at quarterly frequency. Drost and Nijman (1993) have shown that temporal aggregation impedes our ability to detect GARCH effects in the data. Even if our model is not misspecified, the fairly low frequency sampling may suggest it is (see also Bansal and Yaron (2004)). Third, we showed that the (sign of the) correlation between shocks to realized consumption growth and the conditional variance is important in explaining risk and uncertainty premia. The simple GARCH exercise does not help us identify the sign of this correlation. We address this difficulty next.

5.2 Real Dividends Growth Rate: GJR-GARCH

Since we argue that the sign of σ_v plays an important role in understanding risk premia in our model, we also estimate a GJR-GARCH$(1,1)$ (Glosten et al. (1993)). Originally, this model was constructed to capture 'leverage' effects when examining market returns (i.e., a negative shock to returns means lower prices and more leveraged firms, hence higher volatility of future returns). Here we use it with a different interpretation in mind. We use the leverage coefficient to extract information about the sign of the correlation between consumption/dividends growth rate innovations and conditional variance innovations. Since we argue that the sign of σ_v is positive, as indicated by asset prices behavior, we hope to find the reverse of a leverage effect.[30] We fit the following time series model

$$\begin{aligned} \frac{\Delta C_t}{C_{t-1}} &= c + \eta_{C,t}, \qquad (5.2) \\ \eta_{C,t+1} &= \sigma_{C,t}\varepsilon_{C,t+1}, \quad \varepsilon_{C,t} \sim N(0,1), \\ D(L)\,\sigma_{C,t} &= \omega + F(L)\,\eta_{C,t}^2 + G(L)\,I_{\left\{\eta_{C,t}<0\right\}}\eta_{C,t}^2, \end{aligned}$$

where the polynomial G captures the leverage effects and

$$I_{\left\{\eta_{C,t}<0\right\}} = \begin{cases} 1 & \eta_{C,t} < 0 \\ 0 & \text{otherwise} \end{cases}.$$

[30] Even though our interpretation has nothing to do with the leverage effect discussed in Glosten et al. (1993), we still use this term for convenience.

We regress the realized consumption growth rate only on a constant (effectively demeaning the growth rate) since we assume in our model that dividends growth rate drifts on a constant. The more negative η is, the larger is η^2. Thus, we expect the leverage effect coefficient to be negative in order to capture the positive correlation between shocks to growth rates and conditional variance. In most lag specifications we estimated, the leverage coefficients in the G polynomial have a negative sign, which suggests that negative shocks to the dividends growth rate implies a negative shock to the conditional variance. However, and perhaps not surprisingly, with quarterly frequency data it is hard to detect these GARCH effects. Leverage effects are especially hard to detect. In most cases we cannot reject the null that leverage effects are not present. In order to investigate the sign of σ_v further, we use real dividends instead of consumption. To alleviate the problem with the GARCH estimation, we use *monthly* data.[31] Figure 5.4 displays the results of a GJR-GARCH$(1,1)$ estimation where c is the unconditional mean of the real growth rate of aggregate dividends.

This figure shows the presence of volatility clustering. The estimation procedure suggests that σ_v is indeed positive since the leverage coefficient is *always* negative and statistically significant. On average, when a negative shock hits the dividends growth rate, we tend to see a decline in the conditional variance of the same process. Table 5.2 summarizes the estimation results for the leverage coefficient over different time intervals.[32]

It is interesting to note that the earlier post-war data supports more strongly the hypothesis that shocks to dividends are positively correlated with shocks to volatility. This covariation measures the risk exposure of default free bonds to risk and uncertainty. If the market prices of these risks and uncertainty did not move in the opposite direction one should, ceteris paribus, expect to observe higher risk premia in the earlier part of the sample.

In summary, the data seems to confirm two things. First, the existence of a small time-varying component in the volatility of growth rates. Second, the correlation of shocks to

[31] We obtained the real dividends series from Robert Shiller's website. See also Appendix A.

[32] This suggestive evidence is also consistent with different time intervals and with EGARCH estimation (see Nelson (1991)) over the same time intervals. Results are available from the authors upon request.

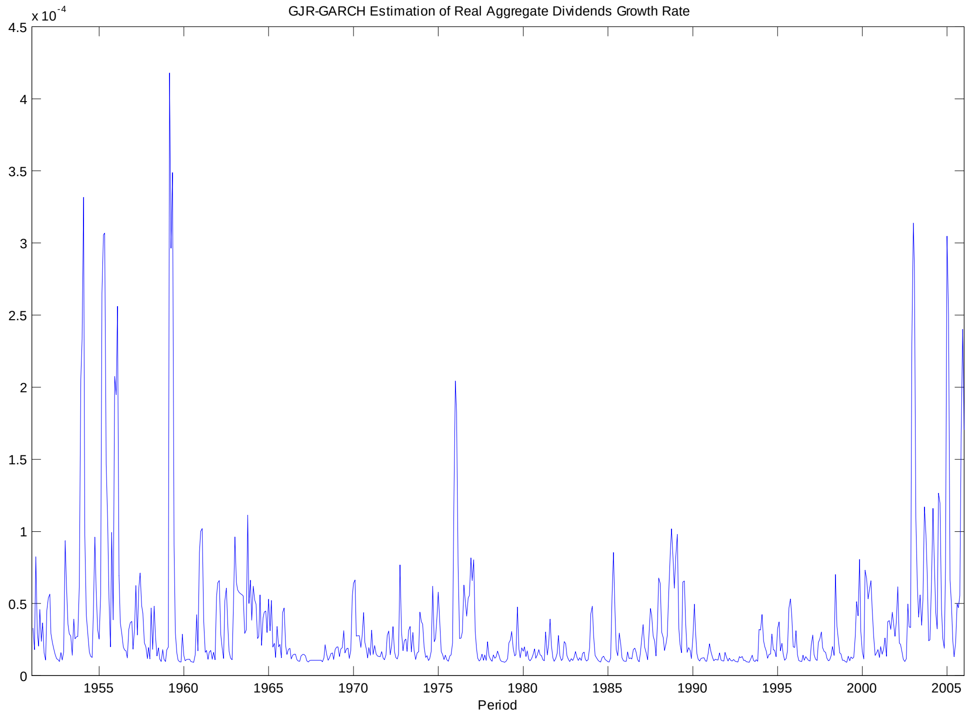

Figure 5.4: GJR-GARCH$(1, 1)$ estimation (model 5.2) of the conditional variance of real aggregate dividends growth rate with monthly observations over the period $M1.52 - M4.06$.

dividends growth rate and shocks to conditional variance is positive.

5.3 Biased Expectations: Pessimism and (the Reverse of) Doubt

Abel (2002) argues that one can potentially account for the equity premium and the risk free rate when modeling pessimism and doubt in an otherwise standard asset pricing (Lucas tree) model. Pessimism is defined as a leftward translation of the objective distribution in a way that the objective distribution first order stochastically dominates the subjective distribution. Doubt is modeled in a way that the subjective distribution is a mean preserving spread of the objective distribution.

There is evidence that people tend to consistently underestimate both market return and

Table 5.2: Estimating the 'leverage' coefficient over different time intervals. The data is monthly real aggregate dividends over $M1.52 - M12.06$ from Robert Shiller's website. A negative point estimate means that a negative shock to realized dividends growth rate is accompanied by a negative shock to the conditional variance of dividends growth rate.

Period	'Leverage' Coefficient	Standard Errors
$M1.52 - M12.06$	-0.390	0.129
$M1.62 - M12.06$	-0.242	0.175
$M1.72 - M12.06$	-0.312	0.215
$M1.82 - M12.06$	-0.263	0.163
$M1.90 - M12.06$	-0.256	0.195
$M1.52 - M12.81$	-0.509	0.167
$M1.52 - M12.89$	-0.442	0.156

the conditional volatility of output growth rate (e.g., Soderlind (2006)). Also, Giordani and Soderlind (2006) confront the Abel (2002) suggestion with survey data and find strong support for the pessimism argument in growth rates of both GDP and consumption. The result is robust over forecasts of different horizon and with both the Livingston survey and the Survey of Professional Forecasters data. However, they also find evidence of overconfidence in the sense that forecasters underestimate uncertainty. Therefore, the evidence suggests the existence of the reverse of doubt.

Our model endogenously predicts both phenomena.[33] First, robustness requirements lead the agent to pessimistic assessments of future economic outcomes (e.g., expression (2.9) in which the agent negatively distorts the expected return on the risky asset). Consequently, the agent persistently underestimates expected growth rates of both the risky asset and consumption. In that sense, robustness endogenizes the pessimism idea of Abel (2002). Our model also predicts biased expectations concerning the dynamics of the conditional variance process v in a way that is consistent with the data. Expressions (4.2) and (4.3) formalize this idea. In the case where $\sigma_v > 0$ (an assumption that we later support empirically), a pessimistic assessment of expected output growth rate leads to what can be interpreted as optimistic beliefs about future output growth volatility. In other words, the model predicts also the reverse of doubt. Note

[33] For a decision-theoretic link between ambiguity averse agent and the setup of Abel (2002), see Ludwig and Zimper (2006).

that here the agent knows exactly the current conditional variance but wrongly estimates its future evolution.

6 'Disciplining Fear': Detection Error probabilities

In this section, we undertake the task of interpreting θ. We showed so far that the model can account for different asset pricing facts and puzzles. Nevertheless, we have yet to tackle an important question - does the model imply too much uncertainty aversion? Even though we showed that coefficients of relative risk aversion and elasticity of intertemporal substitution of unity are sufficient, we still need to gauge the amount of ambiguity aversion implied by the data. Detection error probabilities (DEP's) are the mechanism through which we can interpret θ, and consequently, assess the amount of ambiguity aversion implied by our estimation.

In order to quantify ambiguity aversion, we ask the following: when the agent examines the (finite amount of) data available to him and has to decide whether the reference or the distorted model generated the data, what is the probability of making a model detection mistake? If the probability is very low, this indicates that the two models are far apart statistically, and that the agent should easily be able to distinguish between them. In this case, one might be led to conclude that the degree of robustness implied by our estimation is unreasonably high. If to the contrary, the DEP is high, then it is reasonable to believe that the agent would find it difficult to determine which model is the true representation of the economy.[34]

Technically, DEP's are a mapping from the space of structural parameters to a probability space, which is inherently more easily interpretable than parameter values. Based on our estimate of the parameter θ, we infer the detection error probabilities from the data. It then allows us to interpret whether the degree of ambiguity aversion in our parameterization seems excessive. Appendix B details how to derive the DEP for a given economy using simulations.

[34] For an elaborate discussion of DEP's see, for example, Anderson et al. (2003) and Barillas et al. (2007). For a textbook treatment of robustness and DEP's see chapters 9 and 10 in Hansen and Sargent (2007a).

The last column in Table 4.1 presents the implied DEP's in each economy. First, it is important to point out that DEPs have to be between 0% and 50% (if both models are the same, then there is a 50% chance of making a mistake when assessing which model is the true one). What we find is that our implied DEPs are not unreasonably small, particularly in the context of a framework where the only source of uncertainty is a single shock. This is, once again, an outcome of the interaction between the two main building blocks of our model - robust decision making and state dependent volatility. Together they imply a high enough market price of risk and uncertainty. Consequently, with stochastic volatility our framework does not require a drastic distortion to the reference model. In addition, we should note that imposing consumption growth rate volatility in our procedure makes rejecting the distorted model easier: if we do not impose such restrictions, the DEPs are significantly higher, ranging between 15% and 30%.[35]

Figure 6.1 presents two comparative statics exercises on the implied DEPs. The left panel fixes the benchmark model and varies only θ. The right panel introduces variation only in the number of observations available to the econometrician. We see a clear pattern: Higher θ means less robustness. Thus, it becomes harder to statistically distinguish between the reference and the distorted models as the agent distorts less and less. As $\theta \to \infty$ the DEP reaches 0.5. This is not surprising, since $\theta = \infty$ implies that the distortion to the DGP is zero (recall (2.12)) and both models are therefore indistinguishable. On the other hand, a lower value of θ means more robustness and the models become statistically distant from each other (in the relative entropy sense), reflected in a lower DEP. Similarly, more observations reduce the DEP, in line with our earlier intuition.

[35] See the working paper version for more details.

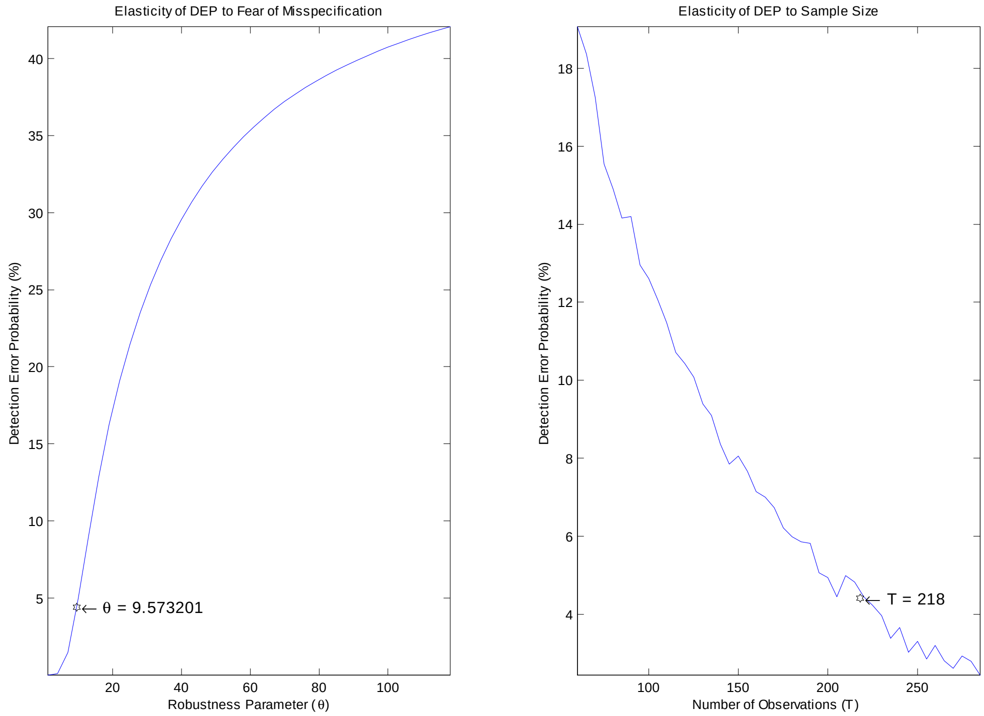

Figure 6.1: Comparative statics on DEP's. Left panel: Fixing all the estimated parameters in the bechmark case with consumption growth volatility as a restriction and over the longest sample $Q2.52 - Q4.06$. Varying robustness in the model (by varying θ on the x-axis) we compute the implied DEP's (y-axis). Right panel: Fixing all the parameters in the benchmark model and varying only the hypothetical number of observations (x-axis) and computing the implied DEP's (y-axis).

7 Conclusion

We presented an equilibrium dynamic asset pricing model that can account for key regularities in the market for default free bonds, while predicting an equity premium, risk free rate and consumption growth as in the data. We estimated the model and showed that it performs well, even though the structural parameters of risk aversion and elasticity of intertemporal substitution are unitary. The results are driven by the interaction of the robust control decision mechanism and state dependent conditional volatility of consumption growth. We interpreted most of what is usually considered risk premium as a premium for Knightian uncertainty. The agent is being compensated in equilibrium for bearing the possibility of model misspecification.

We showed that under the assumption of state dependent conditional volatility of consumption growth, not only the market price of risk is stochastic but also the market price of model uncertainty. As part of our research agenda, we are currently investigating a model with heterogenous robust control agents. Such a model can generate both state dependent risk and uncertainty premia even though the conditional volatility of consumption is constant. The channel through which the model generates stochastic market prices of risk and uncertainty is the trade between the agents and the consequent fluctuations in the agent's relative wealth.[36]

We also suggested that different frequencies in the conditional volatility of consumption growth are potentially important in understanding asset prices. We find it easier to detect high frequency variation in the volatility of consumption growth rate. Also, the full estimation of the model has trouble detecting the lower frequency component. We believe that further investigation of this point is warranted. In addition, an interesting extension would be to consider the link between the evolution of volatility over time and the behavior of asset prices, in the presence of ambiguity aversion. This is directly linked to the recent literature on the Great Moderation in macroeconomics.

We also believe that extending the empirical investigation to a broader asset class can be fruitful. Liu et al. (2005), for example, examine options data in the context of a robust equilibrium with rare events. We believe that one can address different empirical regularities pertaining to the valuation of interest rate sensitive assets with robust considerations. Also, we think that robustness can shed more light on our understanding of exchange rate dynamics, and in particular the failure of uncovered interest rate parity. Finally, our model is a complete characterization of a real economy. One can extend this framework to a nominal one either by assuming an exogenous price level process as in Cox et al. (1985) and Wachter (2001) or by modelling an exogenous money supply process as in Buraschi and Jiltsov (2005) to derive an endogenous price level.

[36] Liu et al. (2005) introduce state dependent market price of uncertainty by modeling rare events. Hansen and Sargent (2010) introduce state dependent market price of uncertainty through the distortion (tilting) of Bayesian model averaging.

References

Abel, A., (2002), "An exploration of the effects of pessimism and doubt on asset returns." *Journal of Economic Dynamics & Control* 26, 1075–1092.

Anderson, E., L. P. Hansen, and T. J. Sargent, (2003), "A quartet of semigroups for model specification, robustness, prices of risk, and model detection." *Journal of the European Economic Association* 1 (1):68–123.

Backus, D., Foresi, S., Mozumdar, A. and Wu, L., (1998), "Predictable changes in yields and forward rates." NBER WP no 6379.

Backus, D., Gregory, A. W. and Zin, S. E., (1989), "Risk premiums in the term structure: Evidence from artificial economies." *Journal of Monetary Economics* 24, 371-99.

Bansal, R., Katchatrian V., and A. Yaron (2005), "Interpretable asset markets?" *European Economic Review* 49, 531-560.

Bansal, R., Kiku, D., and Yaron, A. (2007), "Risks for the long run: estimation and inference." Working paper, Duke and UPenn.

Bansal, R., and A. Yaron. (2004), "Risks for the long run: a potential resolution of asset pricing puzzles." *Journal of Finance* 59(4):1481–509.

Barillas, F., Hansen, L. P., Sargent, T. J., (2007), "Doubts or variability?" Working paper, University of Chicago and NYU.

Bekaert, G., Hodrick, R. and Marshall, D., (1997), "On biases in tests of the expectations hypothesis of the term structure of interest rates." *Journal of Financial Economics* 44, 309-48.

Bliss, R. (1999), "Fitting term structures to bond prices." Working paper, Federal Reserve Bank of Chicago.

Brown, S. J. and Dybvig, P. H. (1986), “The empirical implications of the Cox, Ingersoll, and Ross theory of the term structure of interest rates.” *Journal of Finance* 41, 617-32.

Brown, R. H. and Schaefer, S. (1994), “The term structure of real interest rates and the Cox, Ingersoll, and Ross Model.” *Journal of Financial Economics* 35, 3-42.

Buraschi, A., Jiltsov, A., (2005), "Inflation risk premia and the expectations hypothesis." *Journal of Financial Economics* 75, 429-490.

Cagetti, M., Hansen, L. P., Sargent, T. and Williams, N. (2002), "Robustness and pricing with uncertain growth." *Review of Financial Studies* 15, 363–404.

Campbell, J. Y., (2000), "Asset pricing at the millennium." *Journal of Finance* 55(4): 1515-1567.

Campbell, J. Y. and Shiller, R., (1991), “Yield spreads and interest rates: A bird’s eye view.” *Review of Economic Studies* 58, 495-514.

Chan, K, Karolyi, A., Longstaff, F. and Sanders, A., (1992), “The volatility of the short-term interest rates: An empirical comparison of alternative models of the term structure of interest rates.” *Journal of Finance* 68: 1209-27.

Chapman, D. (1997), “The cyclical properties of consumption growth and the real term structure.” *Journal of Monetary Economics* 39, 145-72.

Cochrane, J. H. and Piazzesi, M., (2002), “Bond risk premia.” *American Economic Review* 95, 138-160.

Cox, J. C., Ingersoll, J. E. and Ross, S. A. (1985), “A theory of the term structure of interest rates.” *Econometrica* 53, 385-407.

Croce, M. M., Lettau, M., Ludvigson, S. C. (2006), "Investor information, long-run risk, and the duration of risky cash flows." Working paper, NYU.

Dai, Q. (2003), "Term structure dynamics in a model with stochastic internal habit." Working paper, UNC.

Dai, Q. and Singleton, K. (2000), "Specification analysis of affine term structure models." *Journal of Finance* 55, 1943-78.

Dai, Q. and Singleton, K. (2002), "Expectation puzzles, time-varying risk premia, and affine models of the term structure." *Journal of Financial Economics* 63, 415-41.

Dai, Q. and Singleton, K. (2003), "Term structure modeling in theory and reality." *Review of Financial Studies* 16(3), 631-678.

Donaldson, J. B., Johnsen, T. H. and Mehra, R. (1990), "On the term structure of interest rates." *Journal of Economic Dynamics and Control* 14, 571–596.

Drost, F.C., Nijman, T.E. (1993), "Temporal Aggregation of Garch Processes." Econometrica 61, 909-927.

Duffie, D., Epstein, L., (1992a), "Asset pricing with stochastic differential utility." *Review of Financial Studies* 5, 411–436.

Duffie, D., and Epstein, L., (1992b), "Stochastic differential utility." *Econometrica* LX, 353–394.

Duffie, D. and Kahn, R. (1996), "A yield-factor model of interest rates." *Mathematical Finance* 6, 379-406.

Duffie, D., and Singleton, K., (1993), "Simulated methods estimation of markov models of asset prices." *Econometrica* 61, 929-952.

Ellsberg, D. (1961), "Risk, ambiguity and the savage axioms." *Quarterly Journal of Economics* 25, 643–69.

Epstein, L. G. and Wang, T., (1994), "Intertemporal asset pricing under knightian uncertainty." *Econometrica* 62(2):283–322.

Epstein,L. G. and Schneider, M., (2003), "Recursive multiple priors." *Journal of Economic-Theory* 113:32–50.

Epstein, L. and Zin, S., (1989), "Substitution, risk aversion and the temporal behavior of consumption and asset returns: a theoretical framework." *Econometrica* 57, 937–969.

Erbaş, S. N. and Mirakhor, A. (2007), "The equity premium puzzle, ambiguity aversion, and institutional quality." IMF WP/07/230.

Fama, E. F. and Bliss, R.R. (1987), "The information in long-maturity forward rates." *American Economic Review* 77, 680-92.

Gagliardini, P., Porchia, P. and Trojani, F., (2009), "Ambiguity aversion and the term structure of interest rates." *Review of Financial Studies* 22(10), 4157-4188.

Gibbons, M. R. and Ramaswamy, K., (1993), "A test of the Cox, Ingersoll and Ross model of the term structure." *Review of Financial Studies* 6, 619-58.

Gilboa, I. and Schmeidler, D., (1989), "Maxmin expected utility with non-unique prior." *Journal of Mathematical Economics* 18:141–153.

Giordani, P., Söderlind, P., (2006), "Is there evidence of pessimism and doubt in subjective distributions? Implications for the equity premium puzzle." *Journal of Economic Dynamics & Control* 30, 1027-1043.

Glosten, L., Jagannathan, R., Runkle, D. (1993), "On the relation between the expected value and the volatility of the nominal excess return on stocks." *Journal of Finance* 48, 1779-1801.

Hansen, L. P., (1982), "Large sample properties of generalized method of moments estimators." *Econometrica* 50, 1029-54.

Hansen, L. P. and Jagannathan, R., (1991), "Implications of security market data for models of dynamic economies." *Journal of Political Economy* 99:225–262.

Hansen, L. P. and Sargent, T. J., (2007a), "Robustness." Princeton, New Jersey: Princeton University Press.

Hansen, L.P. and Sargent, T. J., (2010), "Fragile beliefs and the price of model uncertainty." *Quantitative Economics*, Vol. 1, Issue 1, 129-162.

Hansen, L. P., Sargent, T. J., Turmuhambetova, G. A. and Williams, N. (2006), "Robust control, min-max expected utility, and model misspecification." *Journal of Economic Theory* 128, 45–90.

Hansen, L. P., Sargent, T. and Tallarini, T. (1999), "Robust permanent income and pricing." *Review of Economic Studies* 66, 873–907.

Litterman, R. and Scheinkman, J. (1991), "Common factors affecting bond returns." *Journal of Fixed Income* 1, 54-61.

Liu, J., Pan, J. and Wang, T. (2005), "An equilibrium model of rare-event premia and its implication for option smirks." *Review of Financial Studies* 18(1), 131–64.

Longstaff, F. and Schwartz, E., (1992), "Interest rate volatility and the term structure: A two-factor general equilibrium model." *Journal of Finance* 47, 1259-82.

Lucas, Jr., R. E., (1978). "Asset prices in an exchange economy." *Econometrica* 46, 1426–14.

Ludwig, A. and Zimper, A. (2006), "Rational expectations and ambiguity: A comment on Abel (2002)." *Economics Bulletin*, Vol. 4, No. 2 pp. 1-15.

Maenhout, P. (2004), "Robust portfolio rules and asset pricing." *Review of Financial Studies* 17, 951–983.

Maccheroni, F., Marrinacchi, M., Rustichini, A, (2006) "Ambiguity aversion, robustness, and the variational representation of preferences." *Econometrica* 74, 1447-1498.

McCulloch, J. H. and Kwon, H. C. (1993), "U.S. term structure data, 1947-1991." Working paper, Ohio State University.

Mehra, R. and Prescott, E., (1985), "The equity premium: a puzzle." *Journal of Monetary Economics* 15, 145–61.

Miao, J., (2004), "A note on consumption and savings under knightian uncertainty." *Annals of Economics and Finance* 5, 299-301.

Nelson, D., (1991), "Conditional heteroskedasticity in asset returns: A new approach." *Econometrica* 59, 347-370.

Strzalecki, T. (2007), "Axiomatic foundations of multiplier preferences." Working paper, Northwestern University.

Piazzesi, M. (2004), "Affine term structure models", Y. Aït-Sahalia and L. P. Hansen (eds.), *Handbook of Financial Econometrics*, Amsterdam: North Holland.

Routledge, B. R., and Zin, S. E. (2001), "Model uncertainty and liquidity." Working paper, NBER.

Skiadas, C. (2003), "Robust control and recursive utility." *Finance and Stochastics* 7, 475–489.

Söderlind, P., (2006), "C-CAPM without ex-post data.". WP 344, *NCCR FINRISK*

Stock, J. H. and Watson, M. W., (2003), "Has the business cycles changed? evidence and explanations." in "Monetary policy and uncertainty: adapting to a changing economy." pp. 9-56. Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming.

Uppal, R., Wang, T., (2003), "Model misspecification and underdiversification." *Journal of Finance* 58, 2465–2486.

Wachter, J., (2006), “A consumption-based model of the term structure of interest rates.” *Journal of Financial Economics* 79, 365–399.

Weil, P., (1989), "The equity premium puzzle and the risk free rate puzzle." *Journal of Monetary Economics* 24, 401–421.

Weil, P., (1990), "Nonexpected Utility in Macroeconomics." *Quarterly Journal of Economics* 105 (1):29–42.

Whitelaw, R., (1994), "Time variation and covariations in the expectation and volatility of stock market returns." *Journal of Finance* 49, 515–541.

A Data

Unless otherwise stated, all data are quarterly from $Q2.1952 - Q4.2006$.

- McCulloch-Kwon-Bliss data set: nominal prices and yields of zero coupon bonds - see McCulloch and Kwon (1993) and Bliss (1999). In the estimation exercises we use only the 3 month and 1 year nominal yields at the quarterly frequency to create the real counterparts. The data we use spans the period $Q2.52 - Q4.96$

- Treasury Inflation-Protected Securities (TIPS) data from McCulloch: real yields from $M1.97 - M12.06$. Although the data is available at higher frequencies, we use only observations at quarterly frequency

- Quarterly market index (NYSE/AMEX/NASDAQ) including distributions from CRSP

- Quarterly CPI (all items), SA, from the BLS (see FREDII data source maintained by the federal reserve bank of St. Louis for full description)

- Semi-annual inflation expectations from the Livingston survey (maintained by the federal reserve bank of Philadelphia) - period $H1.52 - H1.81$. From $Q3.81$ quarterly inflation expectations data from the Survey of Professional Forecasters (SPF) becomes availabe

- Quarterly inflation expectations from the SPF maintained by the federal reserve bank of Philadelphia. The sample period covers $Q3.81 - Q4.06$

- Quarterly real Personal Consumption Expenditures (PCE): *services* and *nondurables* from the BEA, SA

- Quarterly real Personal Consumption Expenditures PCE: *imputed services of durables* from the Federal Board of Governors

- Civilian Noninstitutional Population series from the BLS

- Monthly real dividends obtained from Robert Shiller's website over the period $M1.52-M12.06$ (http://www.econ.yale.edu/~shiller/data.htm). This data set was used in the GARCH-GJR exercise

Since we use only real data in the estimation, we convert nominal prices to real ones using the price level data. For the short rate (3 months) we use a 3 year moving average of realized inflation to construct a 3 month ahead expected inflation measure. For the 1 year yield we use both the Livingston and SPF survey data to construct a quarterly series of expected inflation. The SPF is sampled at quarterly frequency but it is available only in the latter part of the sample. We interpolate the semi-annual Livingston data to construct quarterly data using piecewise cubic Hermite interpolation.

B Computing Detection Error Probabilities

In this appendix we shortly discuss how we compute DEP's. The discussion is based on chapter 9 in Hansen and Sargent (2007a). The econometrician observes $\left\{\frac{\Delta C_{t+1}}{C_t}\right\}_{t=1}^{T}$ and construct the log-likelihood ratio of the distorted model relative to the objective model

$$\ell^T = \sum_{t=1}^{T} \log \frac{f\left(\frac{\Delta C_{t+1}}{C_t}|\theta<\infty\right)}{f\left(\frac{\Delta C_{t+1}}{C_t}|\theta=\infty\right)}.$$

The distorted model is denoted with $f(\cdot|\theta<\infty)$ and the reference model is denoted with $f(\cdot|\theta=\infty)$. The distorted model is selected when $\ell^T>0$ and the objective model is selected otherwise.

There are two types of detection errors:

1. Choosing the distorted model when actually the reference model generated the data

$$\mathbb{P}\left(\ell^T>0|\theta=\infty\right) = \mathbb{E}\left(\mathbf{1}_{\{\ell^T>0\}}|\theta=\infty\right).$$

2. Choosing the reference model when actually the distorted model generated the data

$$\begin{aligned} \mathbb{P}\left(\ell^T < 0|\theta < \infty\right) &= \mathbb{E}\left(\mathbf{1}_{\{\ell^T<0\}}|\theta < \infty\right) \\ &= \mathbb{E}\left(\exp\left(\ell^T\right)\mathbf{1}_{\{\ell^T<0\}}|\theta = \infty\right). \end{aligned}$$

Therefore, the *average error (denoted $\wp$) with a prior of equiprobable models* is

$$\begin{aligned} \wp &= \frac{1}{2}\left[\mathbb{P}\left(\ell^T > 0|\theta = \infty\right) + \mathbb{P}\left(\ell^T < 0|\theta < \infty\right)\right] \\ &= \frac{1}{2}\mathbb{E}\left\{\min\left[\exp\left(\ell^T\right), 1\right]|\theta = \infty\right\}. \end{aligned} \tag{B.1}$$

We can write an (approximate) transition likelihood ratio as

$$\begin{aligned} \frac{f\left(\frac{\Delta C_{t+1}}{C_t}|\theta < \infty\right)}{f\left(\frac{\Delta C_{t+1}}{C_t}|\theta = \infty\right)} &= \exp\left[-\frac{1}{2}\times\frac{\left(\frac{\Delta C_{t+1}}{C_t} - \mu - h_t\sqrt{v_t}\right)^2 - \left(\frac{\Delta C_{t+1}}{C_t} - \mu\right)^2}{v_t}\right] \\ &= \exp\left[-\frac{1}{2}\times\frac{-2\left(\frac{\Delta C_{t+1}}{C_t} - \mu\right)(1-\phi)v_t + (1-\phi)^2 v_t^2}{v_t}\right] \\ &= \exp\left[\left(\frac{\Delta C_{t+1}}{C_t} - \mu\right)(1-\phi) - \frac{(1-\phi)^2 v_t}{2}\right]. \end{aligned}$$

We simulate the economy 5,000 times using the point estimates of the parameters and construct a likelihood ratio for each economy. Using (B.1) we can immediately derive $\wp$.