Risk Adjustment and the Temporal Resolution of Uncertainty: Evidence from Options Markets

Darien Huang

Ivan Shaliastovich *

January 2013

Abstract

Risk-neutral probabilities, observable from options data, contain information on physical probabilities and risk adjustments. Under further assumptions on the preference structure, such as state-independent expected utility, physical probabilities and risk adjustments can be separately recovered from risk-neutral probabilities alone. We extend the market-based recovery approach to the recursive utility structure, which allows for a preference for the timing of the resolution of uncertainty. We implement a market-based recovery using S&P 500 options and find that the data strongly supports a specification of early resolution of uncertainty. Failure to account for the magnitude of the preference for early resolution of uncertainty can significantly overstate the implied probability of bad events, understate risk adjustments, and as a consequence under-estimate average market returns.

*Darien Huang (darienh@wharton.upenn.edu) and Ivan Shaliastovich (ishal@wharton.upenn.edu) are from The Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104. We thank Franklin Allen, Ravi Bansal, John Campbell, João Gomes, Philipp Illeditsch, Travis Johnson, Richard Kihlstrom, Krishna Ramaswamy, Nikolai Roussanov, Amir Yaron, and seminar participants at Wharton for their comments and suggestions.

Introduction

By the Fundamental Theorem of Asset Pricing, investors price financial assets by taking the present value of the assets' payoffs under the risk-neutral probabilities (Dybvig and Ross, 1987). The risk-neutral probabilities are different from the objective physical probabilities of the economic states in the data because risk-neutral probabilities incorporate risk adjustments which investors require for reaching these states. These risk adjustments reflect investors' attitudes toward the underlying economic risk, and are determined by the investors' preferences in equilibrium asset-pricing models. While risk-neutral probabilities can be extracted from the data using the observed cross-section of options prices (see Breeden and Litzenberger, 1978), separately identifying both the physical probabilities and risk adjustments is quite challenging; traditionally, it is achieved under fully parametric specifications of both the physical and risk-neutral measures, and using additional asset-price or macroeconomic data.

Recently, Ross (2011) shows that physical probabilities and risk-adjustments can be recovered separately from the risk-neutral probabilities alone when agents' preferences belong to the class of state-independent expected utility. Motivated by recent evidence on the importance of the timing of the resolution of uncertainty (see e.g. Bansal and Yaron, 2004), we extend the recovery framework to the Kreps and Porteus (1978) recursive preferences of Epstein and Zin (1989) and Weil (1989) and show how to separately recover physical probabilities and risk adjustments without using macroeconomic variables (e.g., consumption), and instead relying only on market-based variables. We implement our market-based recovery framework using S&P 500 index options and find that the data strongly support a preference for early resolution of uncertainty, with preference parameter configurations similar to the literature. Using the data and model simulations, we document significant biases in estimating physical probabilities and risk adjustments when the preference for early resolution of uncertainty is not sufficiently accounted for.

Specifically, we consider a Markov regime-switching economy, stationary in economic growth rates. When agents' preferences are characterized by state-independent expected utility, the stochastic discount factor only depends on the next-period realized state. Then, as shown in the multinomial Recovery Theorem in Ross (2011), the requirement that implied physical probabilities sum to one provides enough identifying restrictions to separate risk-neutral probabilities into physical probabilities and risk adjustments. However, in the case of recursive utility, agents also care about current economic conditions which affect the future evolution of wealth, so the agent's utility and the stochastic discount factor are now state-dependent. In particular, we show that the dependence of the pricing kernel on the current state of the economy can be summarized in terms of the price-consumption ratio and the preference parameter θ, which depends on the risk aversion and inter-temporal elasticity of substitution (IES) of the representative agent. We focus on the case when the IES parameter is above one, which is the value entertained in the long-run risks literature to explain a wide range of asset pricing facts (see Bansal and Yaron (2004)). In this case, $\theta < 1$ ($\theta > 1$) corresponds to a preference for early (late) resolution of uncertainty, while $\theta = 1$ nests the CRRA expected utility case. We take into account the recursive utility structure and the dependence of the pricing kernel on the current state and develop a framework for the recovery of physical probabilities and risk-adjustments from the risk-neutral probabilities. Hence, given measurements of the price-consumption ratio and the preference parameter θ, we can recover the pricing kernel and physical probabilities separately from risk-neutral probabilities alone, without using macroeconomic variables such as consumption.

Using an economic model, we illustrate the importance of the recursive utility structure to correctly identify the physical probabilities and risk adjustments from the risk-neutral probabilities. In particular, we consider a Mehra and Prescott (1985) economy with recursive preferences, as in Weil (1989). The consumption dynamics is calibrated to a three-state regime-shift model to match the annual U.S. consumption growth from 1929 to 2010, and we let the dividend stream be a levered claim on consumption dynamics as in Abel (1990). We choose a risk aversion coefficient of 25 and IES of 2 to match the observed equity premium and average risk-free rates in the U.S. Notably, this parameter configuration implies a preference for early resolution of uncertainty. We solve our economic model for the equilibrium risk-neutral probabilities, risk-free rates and price-consumption ratios, and use them to recover the physical probabilities and risk-adjustments under alternative preference parameters. We show that unless the econometric framework correctly accounts for a preference for early resolution of uncertainty, we obtain biased estimates of physical probabilities, risk adjustments, and in turn biased estimates of average stock returns. In particular, when θ is above its true value, we estimate larger implied probabilities of bad events and smaller risk adjustments. Intuitively, when the magnitude of the preference for early resolution of uncertainty is not fully accounted for, the recovery methodology trades-off higher probabilities of bad economic states for lower risk compensations in these states. For $\theta = 1$ it is effectively a manifestation of the failure of the power utility model to generate sizeable risk premia given smooth dynamics of the macroeconomy. As a consequence of these excessively large probabilities of bad states, the implied population averages of the economic variables are significantly biased downward relative to their equilibrium values. For example, while the true unconditional probability of the lowest consumption growth state is calibrated to 25%, under the assumption of expected utility its implied physical probability is over 60%. As a consequence, average stock returns are significantly underestimated under the expected utility assumption: they are less than 1% under the expected utility assumption, relative to 9% in the equilibrium model.

Next, we implement our framework using S&P 500 options data. Specifically, we first identify three aggregate economic states based on the past three-month returns on the S&P 500 index, with the worst state corresponding to the bottom 25% of the unconditional distribution of market returns and the best state to the top 50% of market returns. In the data, the worst state reflects adverse aggregate economic conditions, with an average annualized market return of -35.51%, aggregate real consumption growth of 0.29%, and a VIX measure of implied uncertainty of 30.29%, relative to an average annualized market return of 22.63%, real consumption growth of 1.8% and VIX of 20.22% in the best state. We use S&P 500 options prices to estimate the conditional risk-neutral probabilities of future stock prices in each state following the approach of Figlewski (2008); in particular, the left and right tail of the risk-neutral probability measure are fitted using the Generalized Extreme Value (GEV) distribution. Finally, consistent with the economic model, log price-consumption ratios are assumed to be proportional to the log price-dividend ratios, and the coefficient of proportionality is chosen to match the volatility of the implied price-consumption ratio in the model to the volatility of the price-dividend ratio in the data.

Given the estimates of the conditional risk-neutral densities in the data, the risk-free rate, and measurements of the price-consumption ratio, we proceed to recover physical probabilities and risk adjustments for different preference parameters. Our empirical findings are in line with the evidence from the economic model. We find that the implied probability of a bad state is quite large when the assumed preference parameter θ is positive or not sufficiently negative. Indeed, when $\theta = 1$ and preferences collapse to expected utility, the implied probability of the bad state is about 60%, compared to the set value for the bad economic state of 25%. The optimal value of θ under the which minimizes the distance between the recovered and actual conditional probabilities of the states in the data (Kullback-Leibler divergence criterion) is given by $\theta \approx -11$. This suggests a strong preference for early resolution of uncertainty, and results in an implied physical probability of the bad state close to the actual data.

The decomposition of the risk-neutral probabilities into the physical probabilities and risk-adjustments further reveals substantial differences between different preference structures. When the preference for early resolution of uncertainty is not strong enough, the recovered probability of bad events are so large that the implied risk compensation in bad states is actually smaller than in good states: the magnitude of the SDF going to the bad state is more than two times smaller than going to the good state, while the opposite is the case for recursive utility when $\theta \approx -11$. To obtain the economically plausible implication that the bad state requires higher risk compensation in the data than the good state, the utility structure needs to incorporate a sufficient degree of preference for early resolution of uncertainty. In our case, this requires θ to be below -8. The measurements of the physical probabilities have direct implications for the moments of stock returns and macroeconomic variables. Using the implied physical probabilities under expected utility leads to negative estimates of average stock returns of about -15%. This is a direct consequence of assigning a large probability to bad states with low negative returns. Using $\theta \approx -11$ results in a more plausible estimate of average returns of about 3%, which is more consistent with the evidence in the data.

We perform various checks to confirm the robustness of our main results. First, we check the sensitivity of our results to the identification of the bad economic state. We change the cut-off points for the left tail to the 20th and 30th percentiles of the distribution of market returns, and find that, as in the benchmark case, the implied probabilities of the bad state are 40% and 50%, respectively, and average market returns are negative under expected utility. On the other hand, the implied probability of the bad state is about 30% and average market returns are positive and similar to the benchmark case under the recursive utility structure with a preference for early resolution of uncertainty. We further check that the results do not materially change if we consider two aggregate states instead of three. Finally, while the benchmark analysis is based on quarterly contract maturities, we verify that our main findings hold using monthly horizons. Under expected utility, the probability of bad states is recovered to be 60%, and average market returns are negative. Under recursive utility, the optimal $\theta \approx -40$, the implied probability of bad states is 30% with average market returns of 6%, consistent with the data.

The key goal of our paper is to develop and implement a market-based approach to estimate physical probabilities and risk adjustments separately from the risk-neutral measure, under the recursive utility structure which allows for a preference for the timing of the resolution of uncertainty. Traditionally, the main approach in the literature to identify the economic transitions and the risk adjustments is to specify parametric models for the physical and risk-neutral dynamics (equivalently, for the physical measure and the risk adjustments); see e.g. Pan (2002), Eraker, Johannes, and Polson (2003), and Eraker (2004). Andersen, Fusari, and Todorov (2012) develop a parametric estimation approach for state recovery based on the panel of options prices. Alternatively, Jackwerth (2000), Ait-Sahalia and Lo (2000), and Bliss and Panigirtzoglou (2004) use non-parametric methods to estimate the physical and risk-neutral distributions from the data to derive implications for equilibrium risk adjustments and preference parameters of the representative agent. Jurek (2009) uses a similar method involving transformations of the payoff function adapted from Bakshi, Kapadia, and Madan (2003) to obtain model-free estimates of the moments of the risk-neutral distribution from put and call option prices. In the context of equilibrium models, Shaliastovich (2010) and Eraker and Shaliastovich (2008) consider implications of the recursive utility structure and jumps in economic uncertainty for the equilibrium pricing of options, while Liu, Pan, and Wang (2005) and Gabaix (2007) consider the role of rare jumps for out-of-the money option prices. Bollerslev, Zhou, and Tauchen (2009) and Drechsler and Yaron (2011) discuss predictability implications of the variance risk premia in equilibrium models. Backus, Chernov, and Martin (2011) consider the implications of options markets for the distribution of implied consumption growth disasters. In a related literature, Carr and Wu (2008) consider model-free measurements of the risk premia constructed from options markets.

Our paper proceeds as follows. In Section 1 we describe the theoretical framework first under expected utility and then extend it to recursive preferences. Following that, in Section 2 we set up our model economy, and use model calibrations to highlight the biases for the physical probabilities and the stochastic discount factors that arise without properly considering the preference for early resolution of uncertainty. Section 3 is devoted to our empirical analysis. We estimate the risk-neutral distribution from options prices, and provide evidence that supports the recursive utility structure and preference for early resolution of uncertainty. Section 4 concludes our paper.

1 Theoretical Framework

1.1 Setup of the Economy

We consider a Markovian regime-switching economy with N discrete states. As is typical in the literature, we specify our economy to be stationary in growth rates, in line with classic models such as Mehra and Prescott (1985) and Weil (1989) where consumption growth rates and dividend growth rates are stationary and follow a regime-switching process.

Let $m_{i,j}$ denote the stochastic discount factor (SDF), which in standard structural models corresponds to the intertemporal marginal rate of substitution of the representative agent between state i today and state j next period. According to no-arbitrage pricing conditions, the value of an asset is given by the expectation of its payoff times the SDF. In particular, let $q_{i,j}$ denote the price of an Arrow-Debreu security, which is a state-contingent claim in state i that pays one unit of consumption in state j next period. From the pricing equation, the Arrow-Debreu price is given by the product of the SDF and the physical transition probability $p_{i,j}$:

$$q_{i,j} = m_{i,j} \, p_{i,j}. \tag{1.1}$$

By construction, the sum of the Arrow-Debreu prices given the current state is equal to the price of a one-period risk-free bond in this state, and the Arrow-Debreu price scaled by the price of a risk-free bond, $q_{i,j}/\sum_j q_{i,j}$, corresponds to the risk-neutral probability of going from state i today to state j next period. Following Breeden and Litzenberger (1978), we can estimate the risk-neutral probabilities in a model-free manner from the cross-section of options prices (see Section 3.2). Hence, the Arrow-Debreu prices can be directly identified using market data on option and bond prices alone.

Arrow-Debreu prices incorporate information on both physical probabilities and risk adjustments, as shown in (1.1). To identify the physical probabilities and risk adjustments separately, a standard approach in the literature is to specify parametric models for the physical and risk-neutral dynamics (equivalently, physical dynamics and the SDF), which are then estimated in the data. Alternatively, one can use non-parametric estimates of the risk-neutral and the physical dynamics from the data to identify the implied SDF. As shown in Ross (2011), under restrictions on the class of underlying preferences in the economy, the physical transition probabilities and the SDF can be recovered directly from the Arrow-Debreu prices. Notably, under state-independent expected utility, the SDF and physical probabilities can be recovered using market data alone - without knowledge of the underlying endowment dynamics. We highlight the main steps of this approach below, and then extend our market-based recovery framework to the case of recursive utility.

1.2 Expected Utility Structure

Consider a specification where the representative agent has state- and time-independent expected utility. In this case, the marginal utility between consumption in states i and j depends only on the state the economy is transitioning to and not on the current state. For example, in the classic model of Mehra and Prescott (1985), the agent has expected power utility over future consumption. In this case, the SDF is independent of the current state, and is given by:

$$m_{i,j} = m_j = \delta \lambda_j^{-\gamma}, \tag{1.2}$$

where δ is the agent's time preference parameter, λ_j is the consumption growth rate in state j, and γ is the coefficient of relative risk aversion. With power utility, the SDF is directly related to the endowment dynamics and the preference parameters of the agent. Our subsequent analysis is more general and is not limited to power utility preferences. In fact, it is valid under any specification of time- and state-independent expected utility, and only relies on the identification assumption that the SDF does not depend on the current state.

Let Q denote the matrix of Arrow-Debreu prices, and P the matrix of physical transition probabilities. Let M be the diagonal matrix with the stochastic discount factors m_j on the diagonal. Then, we can rewrite (1.1) in matrix form:

$$Q = PM. \tag{1.3}$$

or more explicitly,

$$\begin{bmatrix} q_{1,1} & \dots & q_{1,N} \\ \vdots & \ddots & \dots \\ q_{N,1} & \dots & q_{N,N} \end{bmatrix} = \begin{bmatrix} m_1\, p_{1,1} & \dots & m_N\, p_{1,N} \\ \vdots & \ddots & \dots \\ m_1\, p_{N,1} & \dots & m_N\, p_{N,N} \end{bmatrix}. \tag{1.4}$$

With N discrete states, this is a system of N^2 equations in $N^2 + N$ unknowns. As pointed out in Ross (2011), to identify the physical probabilities and the SDF separately one can use the further restriction that conditional physical probabilities must sum to 1:

$$P\,\mathbf{1} = \mathbf{1} \tag{1.5}$$

This gives us N additional restrictions, which allows us to identify uniquely both the physical transition probabilities P and the marginal utilities M. Indeed, using the restrictions for the physical probability matrix in (1.3) and (1.5), we obtain:

$$P\,\mathbf{1} = QM^{-1}\mathbf{1} = \mathbf{1} \implies M^{-1}\mathbf{1} = Q^{-1}\mathbf{1}. \tag{1.6}$$

The last condition completely characterizes M since it is a diagonal matrix:

$$M = \left[diag\left(Q^{-1}\mathbf{1}\right)\right]^{-1}. \tag{1.7}$$

We can further recover the implied physical probability transition matrix:

$$P = Q\ diag\left(Q^{-1}\mathbf{1}\right). \tag{1.8}$$

By construction, each row of P sums to one, so subject to the requirement that the recovered P has all positive elements, it represents a valid matrix of conditional probabilities. The implied physical probabilities and the SDF satisfy the pricing equation (1.3), and further, by construction, this decomposition is unique.

Notably, in the analysis above, the physical probabilities and risk adjustments are identified using market data from Arrow-Debreu prices alone, and do not rely on the exact functional form of the SDF, the specification of the endowment process or the agent's preference parameters. Hence, the recovery framework can be implemented empirically using only market data on options and bonds, which are arguably better measured relative to macroeconomic variables. In the next section we consider a recursive utility setup which no longer admits a state-independent SDF. We show that we can maintain a convenient market-based approach for extracting the physical probabilities and the SDF separately from the data, given measurements of the price-consumption ratio and the preference parameter capturing the preference for the timing of resolution of uncertainty.

1.3 Recursive Preferences Structure

For the Kreps and Porteus (1978) recursive utility of Epstein and Zin (1989) and Weil (1989), the life-time utility of the agent V_t satisfies,

$$V_t = \left[(1-\delta) C_t^{1-\frac{1}{\psi}} + \delta \left(\mathbb{E}_t \left[V_{t+1}^{1-\gamma} \right] \right)^{\frac{1-\frac{1}{\psi}}{1-\gamma}} \right]^{\frac{1}{1-\frac{1}{\psi}}} \tag{1.9}$$

where where δ is the time discount factor, $\gamma \geq 0$ is the risk aversion parameter, and $\psi \geq 0$ is the intertemporal elasticity of substitution (IES). For ease of notation, the parameter θ is defined as $\theta \equiv \frac{1-\gamma}{1-\frac{1}{\psi}}$. Note that when $\theta = 1$, that is, $\gamma = 1/\psi$, the above recursive preferences collapse to the standard case of expected utility. As is pointed out in Epstein and Zin (1989), in this case the agent is indifferent to the timing of the resolution of uncertainty of the consumption path. When risk aversion exceeds the reciprocal of IES, the agent prefers early resolution of uncertainty of consumption path, otherwise, the agent has a preference for late resolution of uncertainty. Note that when IES parameter $\psi > 1$, preference for early (late) resolution of uncertainty corresponds to $\theta < 1$ ($\theta > 1$). As shown in Bansal and Yaron (2004), and as we further demonstrate in subsequent sections, the restriction $\psi > 1$ plays an important role to explain a wide range of asset-market features. Hence, we assume that $\psi > 1$ and interpret the parameter θ as capturing the preference of the agent for temporal resolution of uncertainty. An alternative interpretation of θ arises in the robust control and model uncertainty literature (see e.g. Hansen and Sargent (2006)), where the parameter θ captures the agent's aversion to model mis-specification.

As shown in Epstein and Zin (1989), the real stochastic discount factor implied by these preferences is given by

$$m_{i,j} = \delta^{\theta} \lambda_j^{-\frac{\theta}{\psi}} R_{c,i,j}^{\theta - 1} \tag{1.10}$$

where λ_j is the growth rate of aggregate consumption and $R_{c,i,j}$ is the return on the asset which delivers aggregate consumption as its dividends each time period (the wealth portfolio). This return is different from the observed return on the market portfolio as the levels of market dividends and consumption are not equal: aggregate consumption is much larger than aggregate dividends. Let us decompose the consumption return into its cash flow growth rate, λ_j, and the change in price-consumption ratio PC between states i and j :

$$R_{c,i,j} = \lambda_j \frac{PC_j + 1}{PC_i}. \tag{1.11}$$

Substitute the return decomposition above into the recursive utility SDF in (1.10) to obtain:

$$m_{i,j} = \left[\delta^{\theta} \lambda_j^{-\gamma} (PC_j + 1)^{\theta - 1}\right] PC_i^{1-\theta}. \tag{1.12}$$

Therefore, in the case of the recursive preferences, the SDF depends on both current and future economic states. Unlike the power utility case where agents care just about the next-period consumption shock, with recursive preferences they are concerned about the endogenous dynamics of their wealth, so that the economic variables which affect their wealth now determine the marginal rates of substitution between the periods. When $\theta = 1$, the preferences collapse to standard expected utility in (1.2), and the SDF depends only on the next-period consumption growth rate λ_j.

Note that the Epstein-Zin SDF can be decomposed multiplicatively into a component which depends only on the next-period state, $\tilde{m}_j$, and the component which depends only on the current state through $PC_i^{1-\theta}$:

$$m_{i,j} = \tilde{m}_j PC_i^{1-\theta}. \tag{1.13}$$

The SDF component which involves the next-period state, $\tilde{m}_j$, in general depends on the endowment dynamics and all the preference parameters. However, if we can characterize the current-state component ${PC_i}^{1-\theta}$, we can extend the recovery approach and identify the physical probabilities and the SDF directly using market data alone. Indeed, let us modify the matrix equation for the Arrow-Debreu prices (1.3) in the following way:

$$Q = PC(\theta)P\tilde{M}, \tag{1.14}$$

where $PC(\theta)$ and $\tilde{M}$ are the diagonal matrices of $PC_i^{1-\theta}$ and $\tilde{m}_i$, respectively. Then, from the condition that physical probabilities sum to one, we obtain:

$$P\mathbf{1} = PC(\theta)QM^{-1}\mathbf{1} = \mathbf{1} \implies M^{-1}\mathbf{1} = Q^{-1}PC(\theta)\mathbf{1}, \tag{1.15}$$

which uniquely characterizes the SDF and the physical probabilities in terms of the Arrow-Debreu prices and the price-consumption ratios:

$$\begin{aligned} \tilde{M} &= \left[diag\left(Q^{-1}PC(\theta)\mathbf{1}\right)\right]^{-1}, \\ P &= PC(\theta)^{-1}Q\; diag\left(Q^{-1}PC(\theta)\mathbf{1}\right). \end{aligned} \tag{1.16}$$

Hence, given measurements of the price-consumption ratio and the preference parameter θ, we can still recover the SDF and physical probabilities separately from Arrow-Debreu prices alone, without using macroeconomic data[1].

An important parameter which affects the measurements of the physical probabilities and the implied SDF is the preference parameter θ. When $\theta = 1$, the model reduces to power utility; otherwise, the measured price-consumption ratios affect the recovery and identification of the probabilities and risk adjustments. In the next section, we use a calibrated model to illustrate the importance of this parameter to identify the physical probabilities and risk adjustments. Then, we implement this framework in the data and argue that the market data supports a specification with preference for early resolution of uncertainty ($\psi > 1, \theta < 1$).

[1] Of course, one can alternatively specify an exogenous consumption process to characterize the equilibrium risk-neutral dynamics and its relation to the physical measure, as shown in the multinomial setting of Ross (2011), and in long-run risks type models in Eraker and Shaliastovich (2008) and Shaliastovich (2010)

2 Economic Model

2.1 Economy Dynamics

Consider an economic environment which is similar to Weil (1989). The representative agent has recursive utility, described in (1.9), which allows for a preference for the timing of the resolution of uncertainty. In our economy, as in Mehra and Prescott (1985), consumption growth rate λ_j is stationary and follows time-homogeneous Markovian dynamics with a probability matrix P describing the transition across states. To evaluate the model implications for the stock market, following Abel (1990) we model equity as a leveraged claim on aggregate consumption and specify the dividend dynamics in the following way:

$$\lambda_j^d = 1 + \mu_d + \phi(\lambda_j - \mu - 1), \tag{2.1}$$

where μ_d is the mean dividend growth and $\phi > 1$ is the dividend leverage parameter. For simplicity, we abstract from equity-specific shocks, which can be easily introduced as in Bansal and Yaron (2004).

Given the specification of the endowment dynamics and the agent's preferences, we can use a standard Euler equation to compute the equilibrium prices of the consumption asset, equity, the risk-free bond, and the set of the Arrow-Debreu prices. Details of the numerical computations are provided in the Appendix.

2.2 Model Calibration

The model is calibrated on a quarterly frequency to match the key features of U.S. real consumption growth and financial asset markets for a long historical sample from 1929 to 2010. Specifically, we first consider AR(1) dynamics of consumption growth on a quarterly frequency,

$$\lambda_{t+1} = 1 + \mu + \rho(\lambda_t - \mu - 1) + \sigma\epsilon_{t+1}, \tag{2.2}$$

and calibrate the mean μ, persistence ρ and volatility σ parameters to match the annual consumption growth in the data. To calibrate the dividend process, we set the average dividend growth to be the same as the mean consumption growth and fix the dividend leverage parameter to $\phi = 3$. We then discretize the AR(1) dynamics of consumption growth into 3 states using the Tauchen and Hussey (1991) quadrature approach. We have verified that our main results are virtually unchanged using different numbers of economic states, or using an annual model frequency.

The calibrated parameter values are presented in Table 1, and the key moments for consumption growth in the model and in the data are shown in Table 2. In the data, our measure of consumption growth rates is standard and corresponds to annual real expenditures on non-durables and services from the BEA tables; in the model, the population moments are computed from a long simulation of quarterly consumption, time-aggregated to an annual horizon. As shown in Table 2, our calibration ensures that the model matches very closely the mean, variance, and persistence of consumption growth in the data.

We calibrate the preference parameters δ, γ and ψ to match the key moments of financial asset markets. The subjective discount factor δ is set at 0.988, annualized. The risk aversion is calibrated to $\gamma = 25$ and the intertemporal elasticity of substitution parameter ψ is equal to 2.[2] As shown in Table 2, our model delivers an average market price-dividend ratio of 60, a market risk premium of 7% and risk-free rate of 1.5%, which is consistent with evidence in the data. The implied volatility of the market return is about 10% in the model, which is lower than typical estimates in the data of $15 - 20\%$. Recall that for simplicity, our model does not entertain equity-specific shocks (see our dividend growth specification in (2.1)); allowing for these shocks will enable the model to match the volatility of dividends and returns in the data.

Notably, as in the long-run risks literature, we focus on the case when the intertemporal elasticity of substitution is above one and the agent prefers early resolution of uncertainty ($\psi > 1/\gamma$), so the preference parameter $\theta = -48$ is below one. As discussed in Bansal and Yaron (2004), this parameter configuration ensures that asset valuations rise in good times. For example, the price-dividend ratio in the best state

[2]Our value of risk aversion is higher than the value entertained by Bansal and Yaron (2004) using a monthly calibration of the model, and is similar to the estimate in Bansal and Shaliastovich (2012) on a quarterly frequency. Note that our model is specified at the quarterly frequency, which leads to time-aggregation issues and an upward bias in risk aversion measurements, as discussed in Bansal, Kiku, and Yaron (2009).

(state 1) is 62.77 and it drops to 57.12 in the worst state (state 3), and similarly for the price-consumption ratio.

2.3 Model Implications for Probability and Risk Adjustment

We use our model to generate equilibrium Arrow-Debreu prices Q, which can be further decomposed into the physical probability P and the risk adjustment through the stochastic discount factor SDF:

$$\underbrace{\begin{bmatrix} 0.47 & 0.50 & 0.03 \\ 0.07 & 0.61 & 0.33 \\ 0.01 & 0.17 & 0.82 \end{bmatrix}}_{Q} = \underbrace{\begin{bmatrix} 0.68 & 0.31 & 0.01 \\ 0.17 & 0.67 & 0.17 \\ 0.01 & 0.31 & 0.68 \end{bmatrix}}_{P} \times \underbrace{\begin{bmatrix} 0.69 & 1.60 & 3.44 \\ 0.39 & 0.91 & 1.96 \\ 0.24 & 0.56 & 1.21 \end{bmatrix}}_{SDF}, \tag{2.3}$$

where $\times$ indicates element-by-element multiplication. Note that the SDF depends on both the current and future states. The SDF value is highest going from the best state 1 which has the highest consumption growth rate to the worst state 3 in which consumption growth rate is lowest, and smallest values of SDF obtain when we transition to the best economic state 1. The Arrow-Debreu prices incorporate both the risk adjustment and the physical transition probabilities. In this case, because going from the best to the worst state is very unlikely, the Arrow-Debreu prices between states 1 and 3 are inexpensive. On the other hand, consumption in the worst state is very valuable given low current consumption growth, which can be attributed both to a relatively high risk compensation for remaining in state 3 and the persistence of the Markov chain.

The Arrow-Debreu price decomposition above is based on the equilibrium solution of the model given the full calibration of the endowment dynamics and preference parameters. Let us now consider the case when the researcher only has access to the model-generated data on Arrow-Debreu prices and the price-consumption ratios and tries to identify the physical probabilities of economic states and the implied risk compensation. Following our discussion in Section 1.3, the recovered values are based on the preference parameter θ, so we consider a range of possible values for θ and use the conditions in (1.16) to evaluate the magnitudes of the mis-specification of the physical probabilities and the SDF.

In the top two panels of Figure 1 we show the implied unconditional probability of being in the bad state, and the implied value for the SDF which corresponds to transitioning from the good to the bad state relative to remaining in the good state (i.e. $\frac{m_{1,3}}{m_{1,1}}$). When θ is equal to its calibrated value of -48, the unconditional probability of remaining in the bad state and the relative value of the SDF in the bad state are equal to their equilibrium values of 25% and 5.02, respectively. When the candidate value of θ is above its calibrated value, recovery is based on the understated magnitude of the preference for early resolution of uncertainty, leading to biased estimates of physical probabilities and risk adjustments. Specifically, the probabilities of bad events are biased upwards, while the risk adjustments of the bad events are biased downwards. Consider, for example, the decomposition of the Arrow-Debreu prices into the implied physical probabilities and implied SDF under the assumption of expected utility:

$$\underbrace{\begin{bmatrix} 0.47 & 0.50 & 0.03 \\ 0.07 & 0.61 & 0.33 \\ 0.01 & 0.17 & 0.82 \end{bmatrix}}_{Q} = \underbrace{\begin{bmatrix} 0.47 & 0.50 & 0.03 \\ 0.07 & 0.61 & 0.33 \\ 0.01 & 0.17 & 0.82 \end{bmatrix}}_{P(\theta=1)} \times \underbrace{\begin{bmatrix} 0.99 & 1.00 & 1.00 \\ 0.99 & 1.00 & 1.00 \\ 0.99 & 1.00 & 1.00 \end{bmatrix}}_{SDF(\theta=1)}. \tag{2.4}$$

With expected utility, the SDF does not depend on the current state. Further, as is well-recognized in the literature, time- and state-independent expected utility models do not generate enough volatility of the SDF to account for asset market returns. Consistent with this intuition, the decomposition above shows that under the expected utility assumption there is very little action coming from the risk compensation, and virtually all the difference in the Arrow-Debreu prices is attributed to the difference in the implied physical probabilities. Essentially, this is just another manifestation of the risk-free rate and the equity premium puzzles, which states that standard expected utility models cannot simultaneously explain the levels of the risk-free rate and equity risk premium given the actual dynamics of consumption growth in the data (see e.g., Mehra and Prescott (1985)). To account for these asset market features, the recovery framework needs to twist the physical dynamics of the endowment. In particular, it needs to put more weight on the likelihood of bad events to generate expensive Arrow-Debreu prices of going into bad states. Because of that, under expected utility the unconditional probability of bad states are significantly biased upwards and equal to 62% versus its true value of 25%, as shown in Figure 1.

The mis-specification of the economic dynamics has important implications for the implied moments of macroeconomic variables and financial market variables, as shown in the last column of Table 2 where we document the moments of consumption growth and stock returns computed using the recovered physical probabilities under the expected utility assumption. As evident in the Table, attributing more likelihood to bad events leads to a significant downward bias of the measured average consumption growth rate and the mean returns. For example, while the mean consumption growth rate is calibrated to 1.9%, using the recovered physical probabilities under the expected utility framework, the implied mean is -0.8%. Similarly, as shown in the bottom panel of Figure 1, the average return on the market implied by the physical probabilities under expected utility is less than 1%, relative to its calibrated value of 8.9%.

To formally evaluate the mis-specification of the economic dynamics, we compute the Kullback-Leibler (KL) divergence between the implied and calibrated physical probabilities. This is the standard measure of the fit of distributions, and is calculated as the distance between the true distribution and a candidate distribution; smaller values of KL divergence imply a better fit. It is given by the following equation:

$$KL(P||P_\theta) = \sum_i \pi_i \left[\sum_j p_{i,j} \log \left(\frac{p_{i,j}}{{p^\theta}_{i,j}} \right) \right], \tag{2.5}$$

where π_i are the unconditional probabilities of being in each state, $p_{i,j}$ are the calibrated physical transition probabilities, and ${p^\theta}_{i,j}$ are the recovered transition probabilities for a particular candidate value of θ. We plot the value of KL divergence for candidate values of θ on the bottom right panel of Figure 1. The criterion function is minimized at zero (no deviation from the true distribution) at the true value of θ, and significantly rises for alternative values of θ.

Overall, our findings suggest that while analysis based on the expected utility framework suggests a high probability of bad events (e.g., disasters), these results have to be interpreted with caution and might just indicate a mis-specification of the underlying preference structure of the agent. Indeed, when the preference structure allows for a preference for the timing of the resolution of uncertainty, the burden of explaining the cross-section of Arrow-Debreu prices falls less on the physical probabilities, and the differences in Arrow-Debreu prices are attributed more to variations in risk compensation across states.

3 Empirical Analysis

3.1 Data

We use the OptionMetrics database to obtain daily closing prices for exchange-traded S&P 500 index options on the CBOE from 1996 to 2011. On each trading day, there are an average 840 put and call options contracts written on the S&P 500 index and differing with respect to the expiration date and the strike price; however, a significant number of the contracts is subject to liquidity concerns, such as zero trading volume and large bid-ask spreads. To mitigate possible microstructure issues, we follow Figlewski (2008) to apply standard data filters and exclude contracts with zero trading volume, quotes with best bid below \$0.50, and very deep-in-the-money options. Further, our benchmark analysis is conducted on a quarterly frequency, where we use options with 3 months to maturity and track their prices on the expiration dates of the contract. We focus on the quarterly frequency for several reasons. First, the main liquidity in the options markets lies in the primary quarterly expiration cycle: the main hedging instruments for the options are the S&P 500 futures which feature quarterly expirations, so the majority of S&P 500 options also trade on the primary quarterly cycle with expirations in March, June, September, and December of each year. Second, in our empirical analysis, aggregate economic states are identified from the distribution of market returns, and focusing on a relatively lower quarterly frequency helps to reduce non-systematic noise in prices. Finally, using quarterly frequency in the data allows us to directly relate our findings to the economic model in Section 2. We have verified our findings are robust to using a monthly data horizon, and we report the results in Section 4. In addition to options prices, we use data on interest rates which correspond to the 3 month U.S. Treasury rate, and the returns and price-dividend ratio on the S&P 500 index.

As the options data and the implied risk-neutral probabilities are based on the S&P 500 index, we use the distribution of the capital gains on the index, $r^m_{t+1} = \log \frac{S_{t+1}}{S_t}$, to identify the aggregates state of the economy. As evident from the histogram of the capital gains in Figure 2, the return distribution is fat-tailed and negatively skewed. Indeed, the skewness of capital gains over the 1996-2011 period is -0.73, and its kurtosis is 4.62 on a quarterly frequency, which is higher than for a normal distribution. Large negative moves in quarterly returns are likely to contain important information about the aggregate economy and are in general more important from the perspective of a risk-averse investor relative to large positive shocks in returns. Motivated by such considerations, we identify 3 economic states, good, medium, and bad, where the good state corresponds to the upper 50% percentile of the return distribution, the bad state represents the lowest 25% of the returns, and the medium state is in between. The median return in each bin identifies the level of return in each of the economic states, and is given by:

$$\begin{bmatrix} r_1^m & r_2^m & r_3^m \end{bmatrix} = \begin{bmatrix} 22.63\% & 0.19\% & -35.51\% \end{bmatrix}, \tag{3.1}$$

annualized. The estimated transition matrix for the states in the data is equal to,

$$P = \begin{bmatrix} 0.53 & 0.34 & 0.13 \\ 0.47 & 0.20 & 0.33 \\ 0.44 & 0.13 & 0.44 \end{bmatrix}. \tag{3.2}$$

Because the good state corresponds to the upper 50% of the return distribution, there is a considerable probability of remaining in the good state (53%) or transitioning to the medium one (34%). The overall persistence of the aggregate state implied by the transition matrix is low and matches the persistence of returns in the data. The persistence of returns in the data is 0.13, while the persistence of the estimated Markov chain above is 0.23.

We further verify that the states identified by the returns on the index contain meaningful information about the aggregate economy, and we report the average values of the key economic variables, such as real consumption growth, VIX and asset prices, in Table 3. As shown in the Table, there is a significant difference between the average economic variables in the two extreme states. The median real consumption growth is 1.79%, annualized, in the good state, whereas it is a much lower 0.29% in the bad state. The VIX index, which measures uncertainty about the market, is 20.22 in the good state, relative to a much higher value of 30.29 in the bad state, and the price-dividend ratio for the index increases from about 50 to 55 going from the bad state to the good state. An increase in price-dividend ratio in good states is consistent with our assumption that the value of inter-temporal elasticity of substitution is above 1. Looking at Figure 3, we see that the implied volatility curves for each state generated from options data are increasing across the range of moneyness as the aggregate state worsens. In particular, at-the-money implied volatility increases from 18% in the good state to 26% in the bad state. Overall, our economic states meaningfully capture the real growth and uncertainty prospects in the aggregate economy.

3.2 Estimation of the Risk-Neutral Distribution

Theoretically, the entire risk-neutral probability distribution can be extracted directly using a continuum of options contracts, as shown in Breeden and Litzenberger (1978). Let $\tilde{P}(x) = \int_{-\infty}^{x} \tilde{p}(z)\, dz$ denote the risk neutral cumulative distribution function, K be the strike price, and r the risk-free interest rate. Then, given the current value S of the underlying, the price of a European call option expiring at time T is given by:

$$\begin{aligned} C(K; S, r) &= PV\{\mathbb{E}^{Q}[\max(S_T - K, 0)]\} \\ &= \int_{K}^{\infty} e^{-rT}(S_T - K)\, \tilde{p}(S_T)\, dS_T. \end{aligned}$$

Differentiating the call price with respect to strike price allows us to relate the risk-neutral distribution to the prices of call options:

$$\frac{\partial C}{\partial K} = e^{-rT}\left[-(K-K)\,\tilde{p}(K) + \int_{K}^{\infty} -\tilde{p}(S_T)\, dS_T\right] = -e^{-rT}\left[1 - \tilde{P}(K)\right], \tag{3.3}$$

so that the risk-neutral probability is determined by the second derivative of the price of call options:

$$\tilde{p}(K) = e^{rT}\frac{\partial^2 C}{\partial K^2}. \tag{3.4}$$

In practice, we do not observe the entire continuum of options prices, and we do not observe very deep in- and out-of-the money contracts to capture the tails of the distribution. To address the first issue, we interpolate the data to fill in the quotes between listed strikes. Following Shimko (1993), we first transform option prices into Black and Scholes (1973) implied volatilities and interpolate the implied volatility surface, and then transform the interpolated curve back to find a theoretical profile of call option prices by strike. As we are interested in the conditional probabilities of being in the good, medium and bad aggregate states, we use quarterly data to calculate the average volatility surface in each of the three states and then compute the implied risk-neutral distributions conditional on each state. To deal with the unavailability of deep in- and out-of-the money contracts, we follow the steps in

Figlewski (2008) and approximate the tails of the risk-neutral density by a Generalized Extreme Value (GEV) distribution. The tail parameters of the distribution are chosen to match the curvature of the risk-neutral probability density function at two extreme points, along with the requirement that the tail probabilities in both the observed risk-neutral distribution and the GEV tail distributions must equal. A similar approach is also pursued by Vilkov and Xiao (2012). All of the details for the estimation of the risk-neutral distribution are provided in the Appendix.

Figure 4 shows the estimated risk-neutral distributions for each of the aggregate states, together with the GEV adjustments of the right and left tails of the distribution. The bottom panel of Table 3 summarizes the conditional moments of the risk-neutral distribution. Going from good to bad state, risk-neutral volatility increases from 18.92% to 25.27% ; the risk-neutral 3rd central moment becomes about two times more negative, and the 4th moment of the distribution increases twofold as well. Overall, good states are characterized by relatively lower volatility and a relatively lower left tail, which is consistent with our findings on the behavior of VIX and asset prices in the previous section. Our evidence is also consistent with Chang, Christoffersen, and Jacobs (2012) who find that risk-neutral skewness correlates negatively with market returns.

3.3 Data Implications for Probability and Risk Adjustment

Using our estimates of the risk-neutral distributions, we can compute the conditional risk-neutral probabilities between the three aggregate states implied by the options prices. The risk-neutral probabilities, adjusted by the risk-free rates, allow us to calculate the matrix of Arrow-Debreu prices which is specified below:

$$Q = \begin{bmatrix} 0.47 & 0.14 & 0.38 \\ 0.48 & 0.13 & 0.38 \\ 0.50 & 0.09 & 0.41 \end{bmatrix}. \tag{3.5}$$

The Arrow-Debreu prices appear relatively low for bad states: for example, the Arrow-Debreu price of going from the good state to the bad state is 0.38, relative to 0.50 for going from the bad state to the good state. Ex-ante, it is not clear whether the difference in these prices is attributable to the difference in physical transition probabilities (as suggested by our estimate in (3.2), since going from good to bad is less likely than going from bad to good), or by the difference in the magnitudes of risk compensation between the states. To separate the Arrow-Debreu prices into the implied physical probabilities and the risk-adjustment, we implement our market-based recovery methodology outlined in Section 1, allowing for recursive state-dependent utility and a preference for the timing of the resolution of uncertainty.

The recovery of the probabilities and the SDF relies on measurements of the price-consumption ratio and the preference parameter θ. The price-consumption ratio is not directly observed in the data. Consistent with our economic model, we assume that the log price-consumption ratio is proportional to the log price-dividend ratio, $\log PC \approx \alpha + \beta \log PD$, and set the scale parameter β to match the volatility of the price-consumption ratio in the model relative to the volatility of the price-dividend ratio over the long historical sample in U.S.. Given our economic model, the implied estimate of $\beta \approx 1\%$, which is consistent with empirical findings in Lustig, Nieuwerburgh, and Verdelhan (2012) that the price-consumption ratio is less volatile than the price-dividend ratio. We examine the robustness of our results to the scale parameter β in Section 4.

Given these measurements of the price-consumption ratio, we entertain a range of possible preference parameters θ and identify the implied physical probability distribution and the SDF for each of the values of this parameter. For all the values of θ, the top panels of Figure 5 depict the implied unconditional probability of being in the bad state and the implied value of the SDF for transitions from the good to the bad state, while the bottom panels show the implied average market returns, computed under the implied physical probabilities, and the Kullback-Leibler divergence criterion between the recovered transition probabilities and estimated transition probabilities in the data. The actual estimates for the probability of being in bad states, the value of the SDF and the implied average market returns in the case of early resolution of uncertainty and the expected utility are provided in Table 4.

Our empirical findings based on the options data are consistent with the output of the economic model in Section 2. As shown in the bottom right panel of Figure 5, the implied and actual physical probabilities of the states are best matched when θ is sufficiently negative, and the the Kullback-Leibler divergence criterion is minimized at $\theta = -11.28$.

Measurements of θ have important implications for the recovery of the physical probabilities and the risk adjustments. As in the economic model, the recovered probability of the bad state is significantly higher when the preference parameter θ is positive or not sufficiently negative. Indeed, for $\theta = 1$ the preferences collapse to expected utility and the implied probability of the bad state is about 60%, compared to its set value of 25%. When θ equals its Kullback-Leibler optimal value of -11.28, the representative agent has a strong preference for early resolution of uncertainty, and the recovered physical probability of the bad state is close to the actual value in the data. The economic channel which accounts for the upward bias in the recovered probability of bad events is the one highlighted in the economic model: expected utility features very little risk adjustment across states, so physical probabilities have to bear all the burden of explaining the cross-section of Arrow-Debreu prices in the data.

To further illustrate the importance of the recursive utility structure, consider the decomposition of the Arrow-Debreu prices into implied physical probabilities and the SDF in the expected utility framework and one that features a preference for early resolution of uncertainty,

$$\begin{aligned}
\underbrace{\begin{bmatrix} 0.47 & 0.14 & 0.38 \\ 0.48 & 0.13 & 0.38 \\ 0.50 & 0.09 & 0.41 \end{bmatrix}}_{Q} &= \underbrace{\begin{bmatrix} 0.55 & 0.18 & 0.28 \\ 0.55 & 0.17 & 0.28 \\ 0.58 & 0.12 & 0.30 \end{bmatrix}}_{P(\theta=-11.28)} \times \underbrace{\begin{bmatrix} 0.87 & 0.78 & 1.36 \\ 0.87 & 0.79 & 1.37 \\ 0.86 & 0.77 & 1.35 \end{bmatrix}}_{SDF(\theta=-11.28)} \\
&= \underbrace{\begin{bmatrix} 0.26 & 0.17 & 0.57 \\ 0.27 & 0.16 & 0.57 \\ 0.28 & 0.11 & 0.61 \end{bmatrix}}_{P(\theta=1)} \times \underbrace{\begin{bmatrix} 1.79 & 0.82 & 0.67 \\ 1.79 & 0.82 & 0.67 \\ 1.79 & 0.82 & 0.67 \end{bmatrix}}_{SDF(\theta=1)}.
\end{aligned} \tag{3.6}$$

where, again, $\times$ denotes element-by-element multiplication.

In the benchmark case featuring early resolution of uncertainty, the implied physical probabilities are close to their estimates in the data, and the SDF correctly identifies bad states as the ones with the highest risk compensation. This is consistent with the intuition from our economic model, as shown in (2.3). Notably, there is less variation in the SDF across states, which can be explained by a lower persistence of the aggregate states in the data relative to the economic model. In the case of expected utility, implied physical probabilities are quite different from their estimates in the data. The recovered probability of bad events are so large that the implied risk compensation in bad states is actually considerably smaller than in good states: the magnitude of the SDF going to the bad state is more than two times smaller than going to the good state, while the opposite is the case for recursive utility. To obtain the economically plausible implication that the bad state requires higher risk compensation in the data than the good state, the utility structure should incorporate a sufficient degree of preference for early resolution of uncertainty. In our case, this requires θ to be below -8.

The measurements of the physical probabilities have direct implications for the moments of stock returns and macroeconomic variables. As we show in the bottom left panel of Figure 5, using the implied physical probabilities under expected utility leads to negative estimates of average returns of about -15%. This is a direct consequence of assigning a large probability to bad states with low negative returns. Using negative values of $\theta = -11.28$ results in a more plausible estimate of average returns of about 3%, which is more consistent with the evidence in the data. We further evaluate the implications for the measurements of physical probabilities on higher order moments of returns. The measurements of physical probabilities do not have a significant effect on the implied physical volatility of returns: it is 12.92% under expected utility, relative to 12.61% under early resolution of uncertainty with θ of -11.28. However, both skewness and kurtosis of returns are closer to the actual data under the recursive utility structure: for example, the return skewness is about -0.7 for $\theta = -11.28$ relative to 0.6 under the expected utility.

In sum, consistent with our economic model, our empirical findings suggest that the recovery of physical probabilities and the SDF is significantly affected by the underlying preference for the timing of the resolution of uncertainty. When the magnitude of the preference for early resolution of uncertainty is not fully accounted for, the implied physical probabilities tend to trade-off higher probabilities of bad economic events for lower risk compensation in these events. A specification with enough preference for early resolution of uncertainty matches the actual physical probabilities and the moments of returns in the data quite well.

4 Robustness

4.1 Measurements of Price-Consumption Ratio

In our benchmark case, we estimate the log price-consumption ratio assuming it is proportional to the log price-dividend ratio, $\log PC \approx \alpha + \beta \log PD$, and the coefficient β is identified from the volatility of the price-consumption ratio in the model relative to the volatility of price-dividend ratio in the data. To check the robustness of our findings, we consider alternative values of β, and plot the implied physical probability of bad states for a range of β and θ values on Figure 6. As a shown in the plot, to match a calibrated probability of the bad state of 0.25 for values of $\beta > 0$, which is the economically plausible case, the implied preference parameter θ should be below one. While the actual value of the preference parameter depends on the choice of β, for all configurations of $\beta > 0$ the implied preference structure suggests a preference for early resolution of uncertainty.

4.2 Measurements of Aggregate States

Our benchmark specification features three economic states which are identified using the 25th and 50th percentile cut-offs for the return distribution in the data. We consider various robustness checks with respect to the location and the number of bins, and we report the results for the recovery of physical probabilities and the SDF under alternative identification of the aggregate states in Table 4.

Specifically, we first entertain the case where the cut-off point for the bad state corresponds to the 20th percentile of the return distribution. As shown in Table 4, the recovery under expected utility still over-estimates the probability of the bad state to be 40% which results in a -8.47% estimate for the average market return. The value of $\theta = -11.28$ which corresponds to the preference for early resolution of uncertainty minimizes the KL divergence criterion and leads to positive average market returns and the bad state probability closer to the data. Similarly, when the left tail is set above our benchmark specification to the 30th percentile of the distribution of market returns, the implied probability of the bad state is 50% and average market returns are negative under the expected utility. On the other hand, the implied probability of the bad state is 32% while average market returns are 5.7% under the recursive utility structure with preference for early resolution of uncertainty. Note that in both of these configurations we can still meaningfully identify the aggregate economic states. For example, for the 20th percentile left tail specification, the PD ratio in the best state is 55 relative to 52 in the worst state, consumption growth in the best state is 1.8% versus 0.3% in the worst state, and the average VIX is 20 in the best state compared to 35 in the worst state. Similar results hold for the 30th percentile for the specification of the left tail. Interestingly, as the bad state corresponds more and more to the tail events (20th percentile compared to 25th and 30th percentile), the implied value of θ becomes more negative. This suggests that the recursive utility structure and the preference for early resolution of uncertainty play an increasingly important role to account for the larger tail events in stock markets; see e.g. Drechsler and Yaron (2011), Eraker and Shaliastovich (2008) and Shaliastovich (2010) for the discussion of equilibrium recursive utility models with jumps.

While our benchmark specification is a three state model, similar results hold for a two state model as well, which we report in Table 4. We consider as a robustness check a two-state specification where the left tail (bad state) is defined as, respectively, the 20th, 25th, and 30th percentiles of quarterly returns. For all cases, an expected utility specification over-estimates the probability of bad states, and in all the cases except the 30% cut-off, the implied average market returns are negative. The value of θ that minimizes KL divergence between the recovered conditional distribution and the transition probabilities in the data are all negative and imply a preference for early resolution of uncertainty. Under recursive utility, both the probability of bad states and average market returns are much closer to the data compared to the case of expected utility.

We have also considered the robustness of our results with respect to the alternative specifications for the good state. Generally, if the bin structure permits us to meaningfully identify aggregate economic states, our results remain robust.[3]

[3]For a sufficiently high cut-off point used to define a good state, the relative magnitude of economic variables in good states versus bad states are reversed compared to the benchmark case. This might be due to microstructure and data issues associated with the right tail of the return distribution.

4.3 Monthly Horizon

While our main results are presented on a quarterly frequency, our results are robust to using a monthly horizon as well.

We use options with one month to maturity and track their prices on the same expiration cycle as in our benchmark specification, that is, we use quarterly observations in March, June, September, and December. We construct the risk-neutral distribution implied by the monthly options prices and compute Arrow-Debreu prices following our discussion in Section 1. Economic states are defined using the capital gains to the index over the past month and are binned at the 25th and 50th percentiles, as in our benchmark setup.

We report the results for the recovery of physical probabilities and the SDF for the monthly horizons in Table 4, and show the implied unconditional probability of being in the bad state, the implied value of the SDF for transitions from the good to the bad state, the implied average market returns and the Kullback-Leibler divergence criterion for each value of θ in Figure 7. As seen in the Figure and the Table, the Kullback-Leibler divergence is minimized at $\theta = -39.46$. Under expected utility, the probability of the bad state is significantly biased upwards and is equal to 62%, and the implied average market return is very negative. Under a preference for early resolution of uncertainty, the probability of the bad state is 31% and the average market return is 6%. Overall, our evidence at the monthly horizon is similar to the benchmark specification and the economic model.

5 Conclusions

We show how to separately recover physical probabilities and risk adjustments from the risk-neutral probabilities without using macroeconomic variables and allowing for a preference for timing of resolution of uncertainty, thus extending the recovery framework of Ross (2011) to the Kreps and Porteus (1978) recursive preferences of Epstein and Zin (1989) and Weil (1989). We implement our market-based recovery framework using S&P 500 options and find that the data strongly support a specification of early resolution of uncertainty, with preference parameter values similar to common values in the literature. Using the data and model simulations, we document significant biases in estimating physical probabilities and risk adjustments when the preference for early resolution of uncertainty is not sufficiently accounted for.

To highlight the implications of timing of resolution of uncertainty for the physical probabilities and risk adjustments, we first use in a Mehra and Prescott (1985)-Weil (1989) economic model, which incorporates Epstein-Zin utility. We calibrate our model to match stylized facts of financial market returns such as the equity premium and average risk-free rates. In the model, we see that failing to sufficiently account for a preference for early resolution of uncertainty leads to biased estimates of the physical distribution of returns, because we will attribute too large a proportion of the high state prices of bad states to physical likelihoods rather than risk-adjustment.

We then implement our market-based recovery approach using S&P 500 options data. We extract the risk-neutral distribution from options prices using a standard technique, and identify our economic states based on market returns. We apply the framework to the extracted risk-neutral distribution and recover implied dynamics for the physical return probabilities of the U.S. market. We show that not fully accounting for the preference for early resolution of uncertainty results in an over-estimation of the probabilities of bad states and downward-biased estimates of average returns. In all, the evidence from the S&P 500 index options market suggests that the representative agent for the U.S. economy has a strong preference for early resolution of uncertainty.

A Model Solution

The equilibrium price of the asset is computed using the Euler equation:

$$\mathbb{E}_i \left[M_{i,j} R_j \right] = 1. \tag{A.1}$$

We first use this equation for the consumption asset. Given the expression for the stochastic discount factor in (1.12), we obtain that in equilibrium,

$$E_i \left[\delta^\theta \lambda_j^{1-\gamma} \left(\frac{PC_j + 1}{PC_i} \right)^\theta \right] = 1. \tag{A.2}$$

This provides us the equation for the price-consumption ratio in each state i, and we solve the system of equations numerically using fixed point iteration.

Given equilibrium solutions to the price-consumption ratio, we can characterize the stochastic discount factor and obtain equilibrium prices of the Arrow-Debreu claims, the stock market, and the risk-free rates. Specifically, the Arrow-Debreu prices follow the equation (1.1). For the stock market claim, the Euler equation is given by,

$$E_i \left[M_{i,j} \left(PD_j + 1 \right) \left(\mu_d + \phi(\lambda_j - \mu) \right) \right] = PD_i, \tag{A.3}$$

which leads to a linear matrix system for the price-dividend ratio of the market, PD_i. Risk-free rates satisfies

$$R_{f,i} = E_i \left[M_{i,j} \right]^{-1}, \tag{A.4}$$

which can be solved using the equilibrium dynamics for the stochastic discount factor $M_{i,j}$.

B Estimation of Risk-Neutral Density

The exact steps for the estimation of the risk-neutral density from options prices are provided below:

1. Combine puts and calls, which are out-of-the-money (not too deep out-of-the-money, best bid at least \$0.50), and contracts not more than 20 points in-the-money.

2. Transform mid-prices into implied volatilities using Black and Scholes (1973). In the region of +/- 20 points from at-the-money, take a weighted average of put and call implied volatilities.

3. Fit a 4th order polynomial to the implied volatilities over a dense set of strike prices, and convert back into call option prices using Black-Scholes.

4. Numerically differentiate the call prices using (3.3) to recover the risk-neutral distribution function.

References

Abel, Andrew B., 1990, Asset prices under habit formation and catching up with the joneses, *American Economic Review* 80, 38–42.

Ait-Sahalia, Yacine, and Andrew W. Lo, 2000, Nonparametric risk management and implied risk aversion, *Journal of Econometrics* 94, 9–51.

Andersen, Torben, Nicola Fusari, and Viktor Todorov, 2012, Parametric inference and dynamic state recovery from option panels, working paper.

Backus, David, Mikhail Chernov, and Ian Martin, 2011, Disasters implied by equity index options, *Journal of Finance* 66, 1969–2012.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2003, Stock return characteristics, skew laws and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101–43.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2009, An empirical evaluation of the long-run risks model for asset prices, working paper.

Bansal, Ravi, and Ivan Shaliastovich, 2012, A long-run risks explanation of predictability puzzles in bond and currency markets, forthcoming in Review of Financial Studies.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 54, 1481–1509.

Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.

Bliss, Robert R., and Nikolaos Panigirtzoglou, 2004, Option-implied risk aversion estimates, *Journal of Finance* 59, 407–446.

Bollerslev, Tim, Hao Zhou, and George Tauchen, 2009, Expected stock returns and variance risk premia, *Review of Financial Studies* 22, 4463–92.

Breeden, Douglas, and Robert Litzenberger, 1978, Prices of state contingent claims implicit in options prices, *Journal of Business* 51, 621–651.

Carr, Peter, and Liuren Wu, 2008, Variance risk premiums, *Review of Financial Studies* 22, 1311–41.

Chang, Bo Young, Peter Christoffersen, and Kris Jacobs, 2012, Market skewness risk and the cross-section of stock returns, forthcoming in Journal of Financial Economics.

Drechsler, Itamar, and Amir Yaron, 2011, What's vol got to do with it, *Review of Financial Studies* 24, 1–45.

Dybvig, Philip, and Stephen Ross, 1987, Arbitrage, in M. Milgate J. Eatwell and P. Newman, (eds.) *The New Palgrave Dictionary of Economics* (Palgrave Macmillan).

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

Eraker, Bjorn, 2004, Do stock prices and volatility jump? Reconciling evidence from spot and option prices, *Journal of Finance* 59, 1367–1403.

Eraker, Bjorn, M. Johannes, and N.G. Polson, 2003, The impact of jumps in returns and volatility, *Journal of Finance* 53, 1269–1300.

Eraker, Bjørn, and Ivan Shaliastovich, 2008, An equilibrium guide to designing affine pricing models, *Mathematical Finance* 18, 519–543.

Figlewski, Stephen, 2008, Estimating the implied risk neutral density for the u.s. market portfolio, in Jeffrey R. Russell Timothy Bollerslev and Mark Watson, (eds.) *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle* (Oxford University Press).

Gabaix, Xavier, 2007, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, working paper.

Hansen, Lars, and Thomas Sargent, 2006, Fragile beliefs and the price of model uncertainty, working paper.

Jackwerth, Jens Carsten, 2000, Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* 13, 433–451.

Jurek, Jakub, 2009, Crash-neutral currency carry trades, working paper.

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

Liu, Jun, Jun Pan, and Tan Wang, 2005, An equilibrium model of rare-event premia and its implication for option smirks, *Review of Financial Studies* 18, 131–164.

Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan, 2012, The wealth-consumption ratio, working paper.

Mehra, Rajnish, and Edward Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 10, 211–219.

Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.

Ross, Stephen, 2011, The recovery theorem, working paper.

Shaliastovich, Ivan, 2010, Learning, confidence and option prices, working Paper.

Shimko, David, 1993, Bounds of probability, *RISK* 6, 33–37.

Tauchen, George, and Robert Hussey, 1991, Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models, *Econometrica* 59, 317–396.

Vilkov, Grigory, and Yan Xiao, 2012, Option-implied information and predictability of extreme returns, working paper.

Weil, Philippe, 1989, The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* 24, 401–421.

Tables and Figures

Table 1: **Model Calibration**

Preferences	δ	γ	ψ
	0.988^*	25	2
Consumption	μ	ρ	σ
	1.88^*	0.28^*	1.32^*
Dividend	μ_d	ϕ	
	1.88^*	3	

Calibration of model parameters. The model is calibrated on a quarterly frequency. The parameter values with superscript * are annualized, e.g. δ^4 and ρ^4 for the subjective discount factor and consumption growth persistence, 2σ for consumption volatility, and 4μ for the mean. The AR(1) dynamics of consumption growth is discretized into 3 states using the Tauchen and Hussey (1991) quadrature approach. Mean and volatility parameters are in percent.

Table 2: **Model Output**

	Data	Model Under EZ	Model Under EU
$\mathbb{E}[\frac{C_{t+1}}{C_t}]$	1.90	1.90	-0.77
$\sigma\left[\frac{C_{t+1}}{C_t}\right]$	2.20	2.50	1.90
$\rho\left[\frac{C_{t+1}}{C_t}\right]$	0.50	0.50	0.44
$\mathbb{E}[PD_m]$	60.02	59.60	58.16
$\mathbb{E}[R_m]$	7.13	8.89	0.90
$\mathbb{E}[Rf]$	1.19	1.54	0.87

Data and model-implied mean, volatility, and persistence of annual consumption growth (top panel), and average price-dividend ratio, excess returns on the market and the risk-free rate (bottom panel). Data is annual from 1929 to 2010; model statistics are based on a long simulation of quarterly data time-aggregated to an annual horizon. "Under EZ" model output is based on the recursive utility configuration with preference for early resolution of uncertainty, while "Under EU" model output is based on the implied physical probabilities recovered under the expected utility assumption.

Table 3: **Economic Variables in Aggregate States**

	Mean	Std. Dev.	Econ. State Good	Econ. State Medium	Econ. State Bad
Mkt Capital Gains	5.92	17.82	22.63	0.19	-35.51
Mkt PD ratio	57.81	14.00	54.50	57.29	49.54
VIX	22.31	8.08	20.22	18.77	30.29
Real cons. growth	1.31	0.83	1.79	1.74	0.29
RN Distribution:					
Volatility			18.92	19.57	25.27
3rd Moment×1000			-0.74	-0.84	-1.24
4rd Moment×1000			0.37	0.39	0.73

The top panel shows the mean and standard deviation of asset-price and macroeconomic variables, and their median values in good, medium, and bad economic states. Bottom panel shows the volatility, and 3rd and 4th centered moments of the risk-neutral distribution in each state. Economic states correspond to upper 50%, 25%-50%, and lower 25% distribution of capital gains of S&P 500, respectively. Quarterly observations from 1996 to 2011.

Table 4: **Implications for Probabilities and Risk Compensations**

	Under EZ				Under EU			
	θ	Pr. Bad	SDF	Mkt Ret	θ	Pr. Bad	SDF	Mkt Ret
Benchmark	-11.28	0.28	1.57	2.56	1.00	0.59	0.37	-14.85
Three States:								
Left Tail:								
$20th$ pctile	-12.21	0.25	1.23	1.33	1.00	0.39	0.54	-8.47
$30th$ pctile	-7.27	0.32	1.32	5.90	1.00	0.48	0.63	-0.62
Two States:								
Left Tail:								
$20th$ pctile	-52.34	0.21	1.52	1.52	1.00	0.27	1.07	-2.15
$25th$ pctile	-9.43	0.25	1.87	1.94	1.00	0.35	1.19	-2.88
$30th$ pctile	-3.09	0.30	1.47	5.65	1.00	0.34	1.25	4.35
Monthly horizon	-39.46	0.31	1.12	6.07	1.00	0.62	0.22	-26.73

Implied physical probability of a bad aggregate state, the value of the stochastic discount factor from good to bad state relative to staying in a good state, and the implied average annualized market return, recovered under the specifications with recursive preferences ("Under EZ") and under power utility ("Under EU"). Recursive preference model output corresponds to the optimal preference parameter θ which minimizes Kullback-Leibler divergence between the implied and physical conditional probabilities of aggregate states in the data; under expected utility, θ is fixed at 1. The benchmark setup features 3 states and the bad state cut-off at 25% of the return distribution. Robustness specifications include setting the bad state cut-off to 20th and 30th percentiles; using two state, and using monthly data horizons. Quarterly observations from 1996 to 2011.

Figure 1: Implications for Probabilities and Risk Adjustments: Economic Model

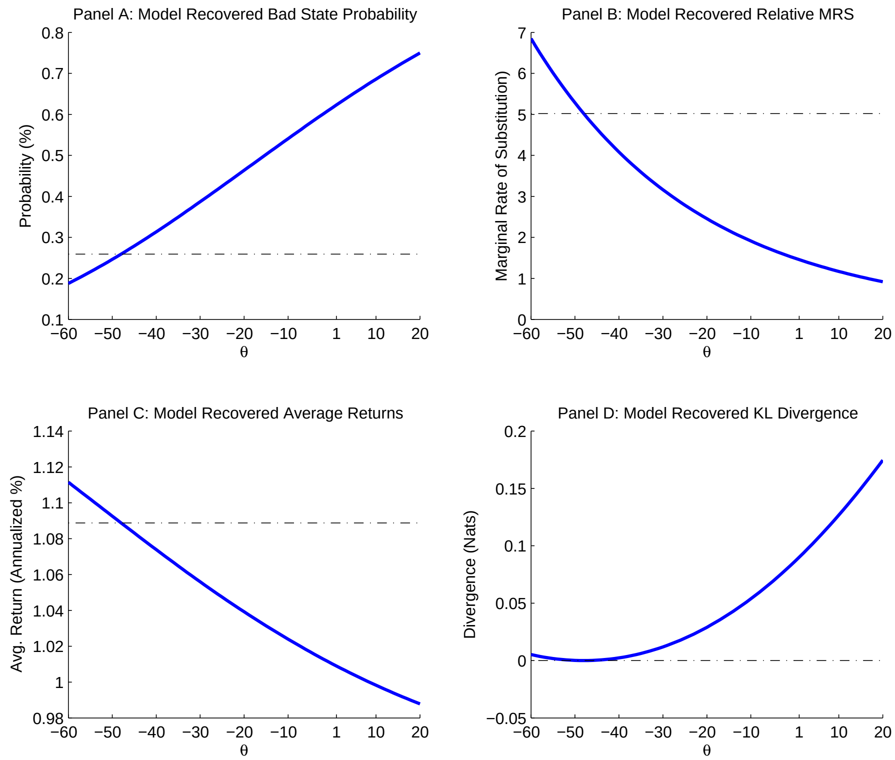

Top panel shows unconditional probability of bad states recovered in the model and the value of the stochastic discount factor from good to bad state relative to staying in a good state as a function of the preference parameter θ. Bottom panel shows the implied average market return and the Kullback-Leibler divergence between the implied and calibrated physical probabilities as a function of the preference parameter θ. The dashed line represents the value given by θ that minimizes the Kullback-Leibler (KL) divergence. The output is based on the economic model.

Figure 2: Empirical Distribution of Market Capital Gains

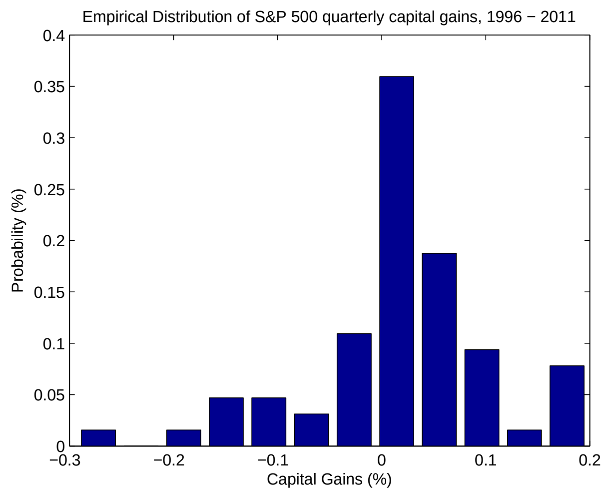

Empirical distribution of log capital gains on S&P 500. Quarterly observations from 1996 to 2011.

Figure 3: Implied Volatility Curves in Economic States

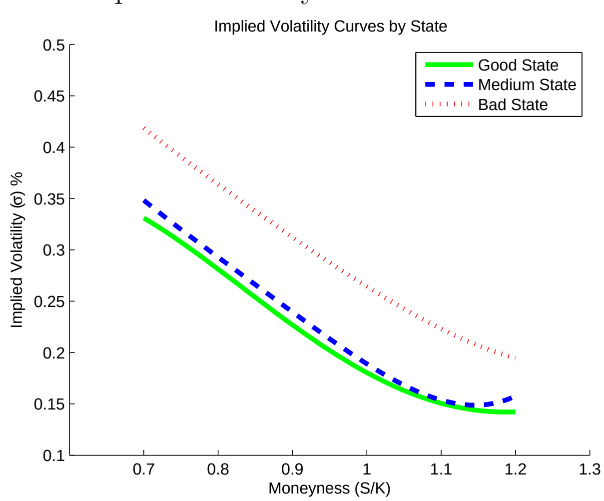

Implied volatility curves for a range of moneyness (spot/strike) in each aggregate economic state. Quarterly observations from 1996 to 2011.

Figure 4: Recovered Risk-Neutral Densities in Economic States

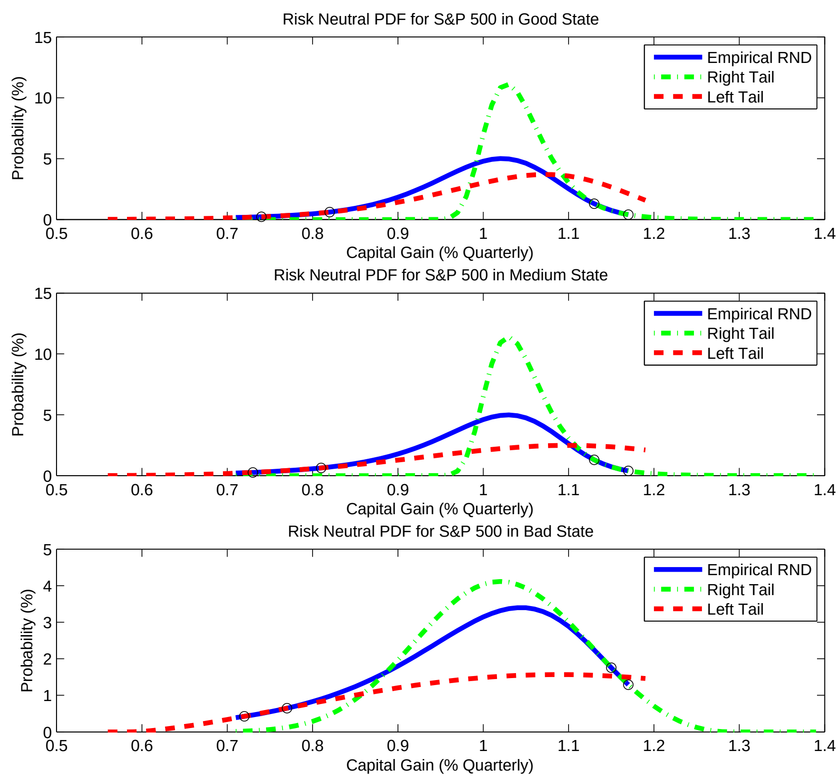

Empirical risk-neutral densities for the market capital gains in good, medium and bad economic states. The blue solid line represents the portion constructed from option data alone; red dashed and greed dashed-dotted lines represent left and right tails, respectively, constructed using the GEV approximation to the underlying data density. Quarterly observations from 1996 to 2011.

Figure 5: Implications for Probabilities and Risk Adjustments: Data

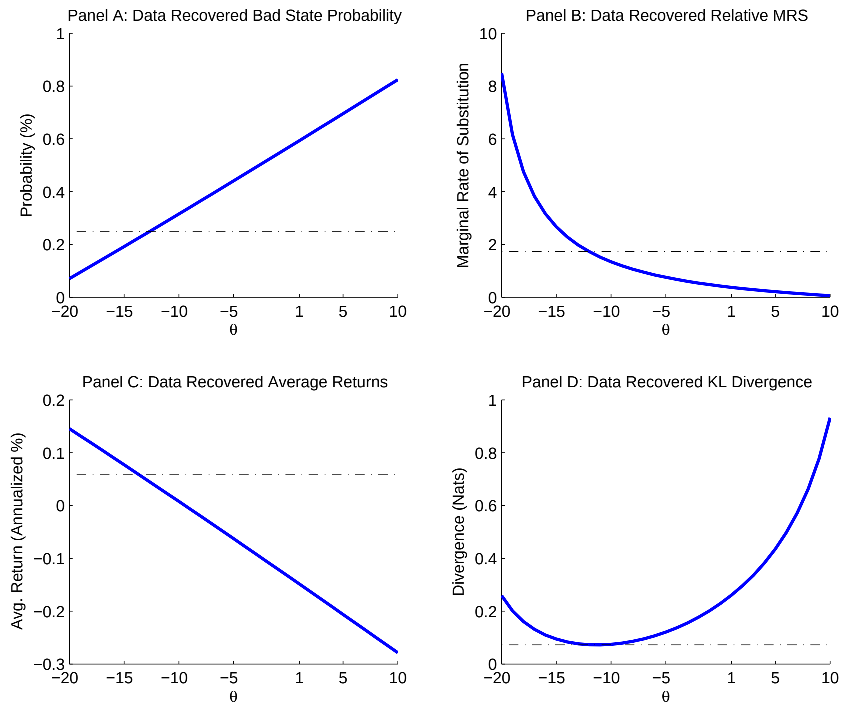

Top panel shows unconditional probability of bad states recovered in the model and the value of the stochastic discount factor from good to bad state relative to staying in a good state as a function of the preference parameter θ. Bottom panel shows the implied average market return and the Kullback-Leibler divergence between the implied and calibrated physical probabilities as a function of the preference parameter θ. The dashed line represents the value given by θ that minimizes the Kullback-Leibler (KL) divergence. The output is based on quarterly observations from 1996 to 2011.

Figure 6: Implied Bad State Probability

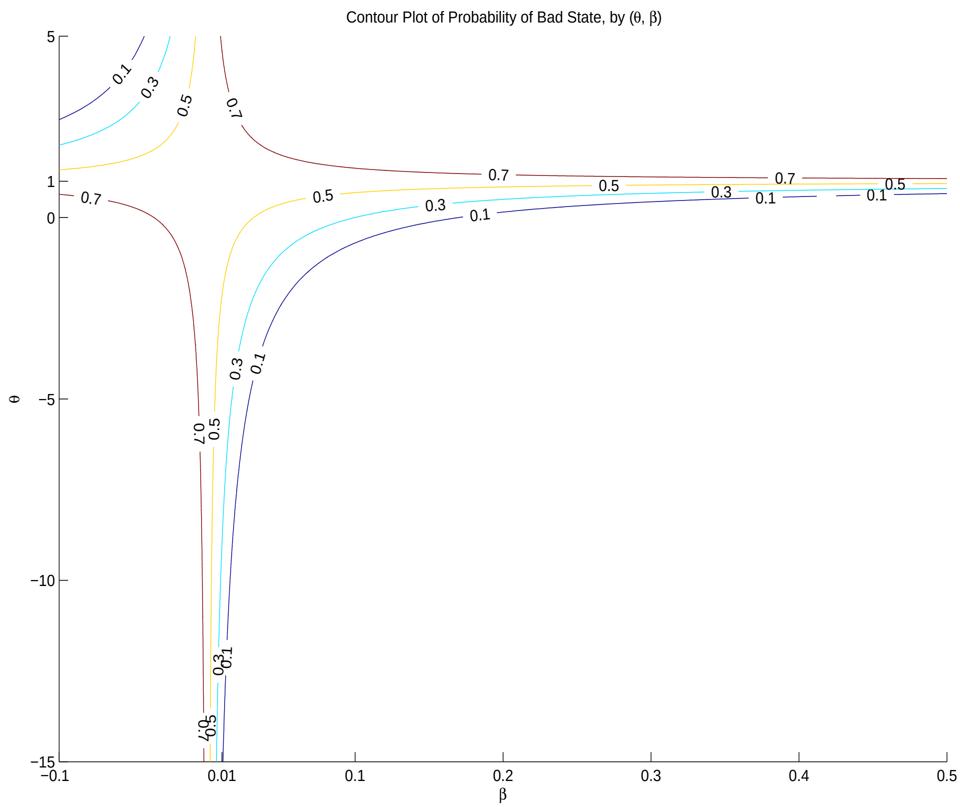

Contour plot of recovered bad state probability in the data, for different combinations of preference parameters θ (y-axis) and price-consumption scale factor β (x-axis). Quarterly observations from 1996 to 2011.

Figure 7: Implications for Probabilities and Risk Adjustments: Monthly Data

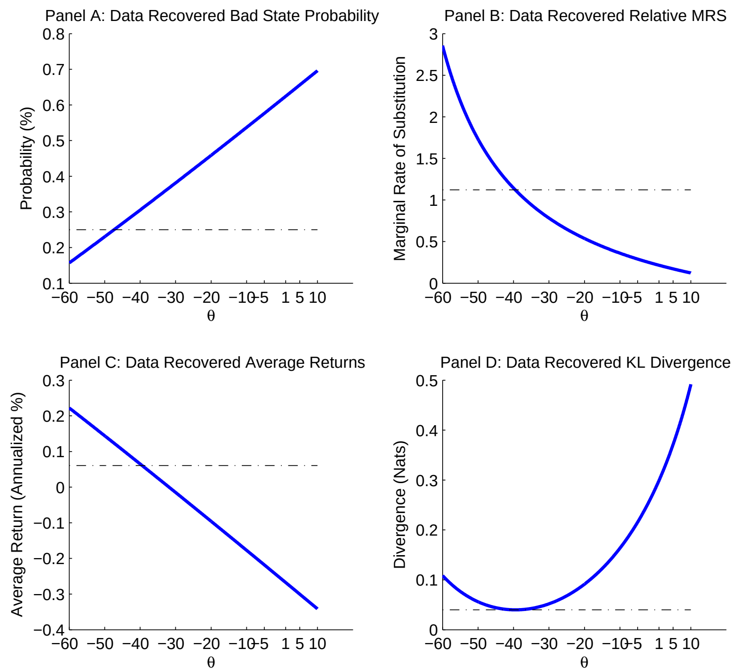

Top panel shows unconditional probability of bad states recovered in the model and the value of the stochastic discount factor from good to bad state relative to staying in a good state as a function of the preference parameter θ. Bottom panel shows the implied average market return and the Kullback-Leibler divergence between the implied and calibrated physical probabilities as a function of the preference parameter θ. The dashed line represents the value given by θ that minimizes the Kullback-Leibler (KL) divergence. The output is based on monthly horizons at primary cycle expirations (March - June - September - December) from 1996 to 2011.