Asset Pricing in a Production Economy

Jeffrey Scott Schafer
Ormond Beach, FL

M.A. Economics, University of Virginia, 2012
B.A. Economics and Statistics, University of Florida, 2011

A Dissertation presented to the Graduate Faculty
of the University of Virginia in Candidacy for the Degree of Doctor of Philosophy

Department of Economics

University of Virginia
May, 2018

Abstract

I study the joint dynamics of consumption and asset returns. In the first chapter, I estimate a DSGE model of a production economy to explore what features in this class of models can account for the long-run risk and stochastic volatility I observe in the data on consumption growth and equity returns. In the asset-pricing literature, there has been a proliferation of endowment economy models that exogenously contain a persistent shock to the average growth rate, known as long-run risk, and a persistent shock to the volatility, known as stochastic volatility, of consumption and dividend growth. Once calibrated, these models successfully account for several well-known asset-pricing puzzles. I use Quasi-Maximum Likelihood and the Kalman filter to estimate the long-run risk and stochastic volatility present in consumption growth and equity returns data. I show that the persistence of the long-run risk and stochastic volatility in the data is significantly lower than has previously been assumed by the literature.

Next, I use Indirect Inference to estimate a production economy model. The key features of my model are long-run risk and stochastic volatility in the exogenous process for aggregate productivity growth, as well as rare disaster events that destroy a fraction of the capital stock and occur with a probability that varies persistently over time. I show that the estimated model, which delivers empirically plausible long-run risk and stochastic volatility in consumption growth and equity returns, is not able to resolve well-known asset pricing puzzles. I provide evidence that improving the model's asset pricing implications comes at the expense of delivering empirically plausible consumption dynamics. Finally, including both long-run productivity risk and disaster risk together improves the model's ability to account for asset price and business cycle quantity moments simultaneously.

In the second chapter of my dissertation, I use the estimated DSGE model of a production economy from Chapter 1 to analyze options prices. I price five European call options, each with the previous chapter's equity asset as the underlying asset. The options all have a three month maturity and range in strike price from ten percent below the spot price to ten percent above it. I show that the price of the ITM options are increasing in today's level of capital, while the price of the ATM and OTM options are decreasing in the current capital stock. I show that the inclusion of long-run productivity risk has a significant impact on the slope of these pricing functions, while (at least for call options) the inclusion of disaster risk in the model does not. Next, I simulate this economy and show that it produces an implied volatility smirk. I propose a measure for the average level of implied volatility skew. Finally, I show how eliminating long-run productivity risk and rare disaster risk from the model affect this measure.

Keywords: asset pricing, production economy, long-run risk, stochastic volatility, rare disasters, options pricing, implied volatility

Acknowledgements

I would like to thank my advisors and colleagues including Eric Young, Toshihiko Mukoyama, Zach Bethune, Leland Farmer, Everett Grant, and Nick Embrey. I graciously acknowledge the support of the Bankard pre-doctoral fellowship.

I would also like to thank my parents, Dan and JoAnn Schafer, as well as my wife, Elizabeth Garrett, for their constant support and never-ending encouragement.

All errors are, of course, my own.

Contents

Chapter 1

Long-run productivity risk and rare disasters: Asset pricing in a production economy

1.1 Introduction

Mehra and Prescott (1985) coined the term "equity premium puzzle". They used a consumption-based asset-pricing model, similar to the one in Lucas (1978), to show that for "reasonable" values of the risk aversion parameter the model could not account for the large equity premium they observed in the U.S. data. Weil (1989) not only found that adding non-expected utility recursive preferences to this Lucas-style environment did not resolve the equity premium puzzle; he added that the model's inability to account for the low risk-free rate observed in the data was also puzzling. Shiller (1981) examined an efficient markets model, in which stock prices are the discounted value of expected future dividends, and concluded that it could not account for the high volatility of stock price indices. These findings opened the dam for a flood of papers that identified other puzzles regarding the canonical consumption-based asset-pricing model's inability to account for certain features of asset price data.

More recently, the literature has adopted several new types of risk as potential resolution to these asset pricing puzzles. In an influential paper, Bansal and Yaron (2004) explored an endowment economy model where consumption and dividends evolve exogenously. In their model they assume a persistent shock to the average growth rates, known as long-run risk, of both consumption and dividends. They also assumed that there are persistent shocks to the variance, known as stochastic volatility, of these growth rate shocks; as well as to the variance of the more traditional shocks to the levels of consumption and dividends. They were able to calibrate the parameters of this exogenous process for consumption and dividend growth, as well as their representative agent's preference parameters, to account for many of the puzzling aspects of the asset pricing data.

This elegant result directly motivates the questions that I answer in this paper. First, is there long-run risk or stochastic volatility in consumption growth and equity returns, and is it consistent with the calibrated process from Bansal and Yaron (2004)? Second, what features in a DSGE model with production can account for the long-run risk and stochastic volatility we observe in consumption growth and equity returns? Finally, can the model that accounts for the observed consumption behavior resolve the aforementioned asset pricing puzzles?

To answer the first question, I estimate a joint process for consumption growth and equity returns. This joint process takes the same form as the exogenous process from the model in Bansal and Yaron (2004); allowing for both long-run risk and stochastic volatility in consumption growth and equity returns. Bansal, Kiku, and Yaron (2012) examined vector autoregressions (VARs) containing consumption growth, price-dividend ratio, and the real risk-free rate. They find evidence for time-varying expected consumption growth as well as time-varying consumption volatility. Croce (2014) used similar methodology and finds evidence for predictable fluctuations in the growth rate of aggregate productivity. There are also many papers that provide evidence for the existence of stochastic volatility in a number of asset returns.[1] By directly estimating the joint process for consumption growth and equity returns I not only provide further empirical evidence for long-run risk and stochastic volatility in both time series, I also provide estimates for the persistence of both of these shocks from the data.

Estimation of this joint process is not trivial, however, since the joint process is non-Gaussian and contains unobserved variables. I use a novel approach to estimate the joint process in two steps. In the first step I ignore the stochastic volatility and assume that the variance of all remaining shocks are time invariant. What remains after assuming away the stochastic volatility equation can easily be written in state space form (SSF). This SSF consists of two measurement equations, one for consumption growth and one for equity returns, and one transition equation for the unobserved long-run risk component. Using the Kalman filter, I follow the prediction error decomposition method from Harvey (1989) to obtain maximum likelihood (ML) estimates for the parameters in this first SSF. In the second step I use the Kalman-smoothed estimate for the time series of the unobserved long-run risk component from step one to obtain residuals for step one's measurement equations. Taking the natural loga-

[1] A few examples include Barndorff-Nielsen and Shephard (2002), Bollerslev and Zhou (2002), and Andersen et al. (2003).

rithm of the squares of these residuals I create two new linear measurement equations, and combine these with the transition equation for the unobserved stochastic volatility component to form another SSF. However, the measurement equations in this SSF are now non-Gaussian. Following Ruiz (1993), I estimate this second SSF using Quasi-Maximum Likelihood (QML). I find evidence of long-run risk and stochastic volatility in consumption growth and equity returns, as the variance of both shocks is statistically different from zero. However, both the long-run risk and the stochastic volatility in these time series are significantly less persistent than has commonly been assumed by the previous literature.

Next, I examine a DSGE model with production, in which consumption is now determined endogenously. Several of the features of the model have become almost ubiquitous in production-based asset-pricing models. Epstein and Zin (1989) preferences disentangle the relationship between risk aversion and the intertemporal elasticity of substitution (IES). And convex capital adjustment costs, in the style of Jermann (1998), allow for fluctuations in the price of capital. In addition to these, the model examined here combines two features that have shown particular promise in the literature for resolving production-based asset-pricing puzzles. First, the model in this paper contains both long-run risk and stochastic volatility in the exogenously specified process for aggregate productivity growth. Croce (2014) first added long-run risk to the exogenous process for aggregate productivity growth; and provided empirical evidence for long-run risk in the data on aggregate productivity. The model in Croce (2014) was able to deliver a large equity premium. Second, the model I present here includes rare disaster events that destroy a random fraction of the capital stock. The probability of such a disaster event occurring varies over time and is also persistent. Almost as soon as Mehra and Prescott (1985) discovered the equity premium puzzle, Reitz (1988) put forward the idea that rare disaster events might be the solution. Gourio (2012) included rare disasters that diminish the economy's aggregate productivity, as well as destroy a fraction of the economy's capital stock. He showed that it was important that these disasters occur with a time-varying probability; otherwise they were shown to be observationally equivalent to a preference shock. I examine a model with both long-run productivity risk and rare disasters, and assess this model's ability to account for the long-run risk and stochastic volatility in the consumption growth and equity returns data. Combining these two features allows me to explore the relative importance of each in accounting for both asset-pricing data and business cycle data.

I use Indirect Inference to estimate many of the key parameters in my model. Smith (1993) and Gourieroux, Monfort, and Renault (1993) concurrently defined the Indirect Inference estimator. This approach has a lot in common with the better-known method of moments estimator. However, the hallmark of Indirect Inference is the use of a so-called "auxiliary model". Rather than specifying a set of moments in the data to use as targets, as in the method of moments approach, in Indirect Inference estimation one specifies an auxiliary model whose parameters define the moments in the data that will be targeted. The auxiliary model can be anything, so long as it can be estimated using both real data and data obtained from simulating the economic model being estimated. The joint process for consumption growth and equity returns, whose empirical estimation I have already described above, is a natural choice for the auxiliary model in my Indirect Inference estimation procedure. This estimation technique allows me to directly target the long-run risk and stochastic volatility in the data on consumption growth and equity returns. This auxiliary model is estimated first using real data, and then using data obtained from simulating the economic model. Indirect Inference estimation stipulates that I choose the set of parameters for my economic model such that the simulated data obtained from this solution will yield estimates of the auxiliary model's parameters that are as close to those obtained using the real data as possible.

I show that this economic model is able to account for the long-run risk and stochastic volatility that we observe in the data on consumption growth and equity returns. Perhaps unsurprisingly, most of the model's ability to accomplish this relies on the long-run productivity risk in the model. When the disaster risk in the model is shut down it still performs this task reasonably well, whereas when the long-run productivity risk is shut down the model performs rather poorly in this dimension. However, I show that including both forms of risk in the model improves its ability to account for a larger set of business cycle quantity moments, without hurting its asset pricing implications.

Next, I explore the asset pricing implications of the estimated model. I find that this model, with empirically plausible long-run risk and stochastic volatility in consumption, is not able to resolve asset-pricing puzzles. I also find, by changing the estimation procedure, that improvement in the model's ability to account for the asset-pricing data comes at the expense of its ability to replicate the estimated consumption dynamics. This result suggests that previous work, which claimed to have resolved these puzzles, may have relied on consumption behavior that is not supported by the data.

Finally, I explore the model's inability to account for one feature in the data captured by my auxiliary model. The model is not able to account for the fact that equity returns are roughly 20 times more volatile, on average, than consumption growth. This lack of volatility is a common shortcoming of production based asset pricing models, present in both the model in Croce (2014) and the model in Gourio (2012). I show that improvement in the model's ability to account for this dimension of the data mitigates the model's ability to match the persistence of the long-run risk and stochastic volatility components. I also explore the returns on a levered equity claim. By re-estimating my model using the levered equity claim, I can dramatically improve the model's ability to account for the volatility of excess returns.

In the next section of the paper, I define the full model. In section 3, I explain my discrete approximation to the model's continuous process for aggregate productivity and provide evidence for the accuracy of this method. Section 4 describes my Indirect Inference estimation technique. I analyze the results of this estimation procedure in section 5. I conclude in section 6.

1.2 Model

In this section of the paper, I describe a DSGE model of a production economy. This representative agent model builds upon the standard real business cycle (RBC) framework. I begin by describing the representative consumer's preferences.

1.2.1 Consumer Preferences

Following Epstein and Zin (1989), consumers have recursive preferences

$$U_t(C_t, \ell_t, \psi_t) = \{(1-\beta)\tilde{C}_t^{\gamma} + \beta\psi_t^{\gamma}\}^{1/\gamma},$$

where $\tilde{C}_t$ is the per period utility from consumption, C_t, and leisure, ℓ_t,

$$\tilde{C}_t = C_t^{\theta}(Z_{t-1}\ell_t)^{1-\theta}$$

and ψ_t is the certainty equivalent of future uncertain utility

$$\psi_t = (E[U_{t+1}(C_{t+1}, \ell_{t+1}, \psi_{t+1})^{1-\delta}])^{1/1-\delta}.$$

I assume leisure is multiplied by the previous period's level of aggregate productivity, Z_{t-1}, in order to ensure balanced growth.

In a model where the representative consumer has expected utility preferences the intertemporal elasticity of substitution (IES) and the coefficient of relative risk aversion are inverses of one another. Recursive preferences include one additional parameter in order to break this relationship apart. In this model, IES is $(1-\gamma)^{-1}$ and the Arrow-Pratt coefficient of risk aversion is δ. Using a model with recursive preferences Tallarini (2000) was able to show that, while holding IES constant at one, increasing the level of risk aversion improved the model's ability to account for asset-pricing data while having almost no impact on the model's implications for business cycle quantities. This demonstrated the value of separating risk aversion from IES.

Empirical estimates of IES vary widely. However, many papers, including Bansal and Yaron (2004), have highlighted mechanisms that rely on IES to be greater than one in order to generate a large equity risk premium. In this paper I estimate the parameters that control IES and risk aversion. I do not restrict risk aversion to be large, nor the IES to be greater than one.

1.2.2 Production

In this model, the representative firm produces output, Y_t, according to a Cobb-Douglas production function

$$Y_t = K_t^{\alpha}(Z_t n_t)^{1-\alpha},$$

where K_t is capital, n_t is labor input, and Z_t is aggregate productivity. The capital stock evolves according to the law of motion

$$K_{t+1} \leq (1-\delta_K)K_t + I_t - G_t K_t,$$

where capital depreciates at the rate δ_K, and I_t is investment. Following Jermann (1998) this law of motion for capital contains a convex capital adjustment cost, given by G_t. I assume a quadratic form for the cost

$$G_t = \frac{\phi}{2}\left(\frac{K_{t+1}}{K_t} - 1\right)^2.$$

In a model without capital adjustment costs, the supply of capital is perfectly elastic. Capital adjustment costs are of crucial importance as they create fluctuations in the price of capital. I estimate ϕ, which controls the size of the cost.

The economy in the model is subject to two aggregate resource constraints

$$C_t + I_t \leq Y_t,$$

and

$$l_t + n_t \leq 1.$$

1.2.3 Productivity Process

The process for aggregate productivity growth includes both long-run risk and stochastic volatility, as in Croce (2014). Let Z_t be aggregate productivity and z_t be the natural log of its value. Then

$$\Delta z_{t+1} = log(Z_{t+1}) – log(Z_t)$$

approximates the growth rate of aggregate productivity. Aggregate productivity growth evolves according to

$$\Delta z_{t+1} = \mu + x_t + \sigma_t \epsilon_{t+1}, \tag{1.1}$$

where μ is the average growth rate of aggregate productivity and ϵ_{t+1} is an *iid* shock to the level of aggregate productivity growth, with standard deviation σ_t at time t. x_t is a shock to the average growth rate of aggregate productivity. This persistent shock to the average growth rate, which evolves according to the AR(1) process

$$x_{t+1} = \rho_x x_t + \varphi_x \sigma_t \eta_{t+1}, \tag{1.2}$$

is referred to as long-run risk. The long-run component, x_t is mean zero, its persistence is captured by ρ_x, and the relative volatility of this shock is given by φ_x. The modelling of these two separate shocks allows aggregate productivity growth to contain both an *iid* short-run component and a slow moving long-run component. Most of the long-run risk literature assume that the long-run component is small but extremely persistent. This means that the effects of a large realization of the shock to this long-run component would be extremely long-lasting. In this paper I estimate the persistence and relative volatility of this long-run component.

The natural log of the variance, $log(\sigma_t^2)$, of both shocks is time-varying. This property is referred to as stochastic volatility. The stochastic volatility evolves according to another AR(1) process

$$log(\sigma_{t+1}^2) = (1-\rho_\sigma)\mu_\sigma + \rho_\sigma log(\sigma_t^2) + \sigma_\nu \nu_{t+1}. \tag{1.3}$$

Here, μ_σ is the unconditional mean of the process and ρ_σ is its persistence. σ_ν is the standard deviation of the shock to the volatility process. Because of this stochastic volatility, the processes for the evolution of aggregate productivity growth and its long-run component are non-Gaussian. Through Monte Carlo simulation experiments it can be shown that an increase in σ_ν increases the kurtosis in the distributions of both Δz_{t+1} and x_t. The stochastic volatility can be interpreted as fluctuation in the overall level of uncertainty surrounding the fundamental variables of the economy. I also estimate the persistence and volatility of this stochastic volatility process in this paper.

With some simple algebraic manipulation of the terms in equation 1 we can arrive at

$$log(Z_{t+1}) = \mu + x_t + log(Z_t) + \sigma_t \epsilon_{t+1},$$

which makes it obvious that the evolution of aggregate productivity contains a unit root. As a result, the process is non-stationary and the effects of this shock are permanent. In order for me to be able to solve the model I follow Linde (2009) and scale each variable by the level of aggregate productivity today. In what follows I use a hat symbol to denote the transformed quantities. $\hat{C}_t = C_t/Z_t$, $\hat{Y}_t = Y_t/Z_t$, and so on. Since K_{t+1} is a date t choice variable $\hat{K}_{t+1} = K_{t+1}/Z_t$, and thus $\hat{K}_t = K_t/Z_{t-1}$. Hours worked, n_t, need not be transformed since the functional form for per period utility is specified to ensure that it is stationary. Along the balanced growth path in this economy all of the other variables will grow at the common rate μ from equation 1.

1.2.4 Disasters

Each period there is some probability, π_t, that a disaster event will occur. These disasters are assumed to be rare, meaning that the average probability of a disaster is quite small, but very severe. This stands in stark contrast to the type of risk, typically to aggregate productivity, modelled by the standard real business cycle framework. The more traditionally modelled risks are smaller in magnitude, but negative realizations of these shocks occur much more frequently. As far back as Reitz (1988), rare disaster risk has been thought to explain asset-pricing puzzles. In this paper, similar to the model in Gourio (2012), I assume that a disaster event destroys a fraction of the capital stock in the economy. The size of this disaster event is random with

$$\zeta_t \sim NID(\mu_\zeta, \sigma_\zeta^2),$$

where ζ_t determines the fraction of the capital stock destroyed. To incorporate these disasters, I can re-write the law of motion for capital in the economy as

$$K_{t+1} \leq \Big((1-\delta_K)K_t + I_t - G_t K_t\Big) e^{D_{t+1}\zeta_{t+1}},$$

where D_{t+1} is an indicator function that takes on a value of one if a disaster event occurs. I assume that a disaster is realized only after the representative agent has made all decision for the current period t. The agent makes all decisions for the period fully aware of the current probability of a disaster π_t. However, only after all decisions have been made does D_{t+1} take on a value of zero or one; and if $D_{t+1} = 1$, only then is the size of the disaster event, ζ_{t+1} , realized.

The probability of a disaster event, π_t, changes over time according to

$$\log(\pi_{t+1}) = (1 - \rho_\pi)\log(\bar{\pi}) + \rho_\pi \log(\pi_t) + \sigma_\upsilon \upsilon_{t+1}, \tag{1.4}$$

where $\bar{\pi}$ is the average probability of a disaster. σ_υ is the standard deviation of the shock to the probability of a disaster, and ρ_π governs the persistence of this shock. Gourio (2012) showed that if the probability of a disaster, where disasters were modelled similarly, was constant then these disaster events were observationally equivalent to a preference shock. The results from that paper suggested that time-variation in the probability of disaster events was crucial to resolving asset-pricing puzzles. In this paper I will estimate both ρ_π and σ_υ.

1.2.5 Solution Method

It is straightforward to show that the solution to the planner's problem is equivalent to the competitive equilibrium in this economy. I now define the planner's dynamic programming problem. The aggregate state of the economy is described by five state variables: aggregate capital stock K_t, aggregate productivity Z_t, long-run component x_t, volatility σ_t, and disaster probability π_t. Let $\boldsymbol{\kappa}$ be a vector containing these state variables. The planner optimizes over four choice variables: consumption C_t, hours worked n_t, investment I_t , and tomorrow's capital stock K_{t+1}. Let $\boldsymbol{\chi}$ be a vector of the choice variables. The planner's value function is therefore

$$V(\boldsymbol{\kappa}) = \max_{\boldsymbol{\chi}} \left((1-\beta) \left(C_t^{\theta} (Z_{t-1}(1-n_t))^{1-\theta} \right)^{\gamma} + \beta \left(E[V(\boldsymbol{\kappa'})^{1-\delta}] \right)^{\frac{\gamma}{1-\delta}} \right)^{\frac{1}{\gamma}},$$

where $1-n_t$ has been substituted in for ℓ_t since the time constraint always binds. The planner maximizes this value function subject to two non-linear inequality constraints

$$C_t + I_t \leq K_t^{\alpha} (Z_t n_t)^{1-\alpha}$$

and

$$K_{t+1} \leq \left((1-\delta_K) K_t + I_t - \frac{\phi K_t}{2} \left(\frac{K_{t+1}}{K_t} - 1 \right)^2 \right) e^{D_{t+1} \zeta_{t+1}}.$$

The first constraint is obtained by simply substituting the production function into the aggregate resource constraint. Aggregate productivity, Z_t , evolves according to equations 1 through 3. The second constraint is the law of motion for capital, where the probability that $D_{t+1} = 1$ is π_{t+1} from equation 4, and ζ_{t+1} is the previously specified *iid* normal random variable.

I follow Linde (2009) and transform this dynamic programming problem to ensure stationarity. The full algebraic derivation of the following equations can be found in an appendix to this paper. Recall that variables with a hat have been divided by aggregate productivity ($\hat{C}_t = C_t/Z_t$ for example). Let $g_t = Z_t/Z_{t-1}$, then the planner's value function can be rewritten as

$$\hat{V}(\hat{\boldsymbol{\kappa}}) = \max_{\hat{\boldsymbol{\chi}}} \left((1-\beta) \left(\hat{C}_t^{\theta} (g_t^{-1}(1-n_t))^{1-\theta} \right)^{\gamma} + \beta \left(E[\hat{V}(\hat{\boldsymbol{\kappa'}})^{1-\delta}] \right)^{\frac{\gamma}{1-\delta}} \right)^{\frac{1}{\gamma}}.$$

The two non-linear inequality constraints can be rewritten as

$$\hat{C}_t + \hat{I}_t \leq g_t^{-\alpha} \hat{K}_t^{\alpha} n_t^{1-\alpha}$$

and

$$\hat{K}_{t+1} \leq \left(g_t^{-1}(1-\delta_K)\hat{K}_t + \hat{I}_t - \frac{\phi \hat{K}_t}{2 g_t} \left(g_t \frac{\hat{K}_{t+1}}{\hat{K}_t} - 1 \right)^2 \right) e^{D_{t+1}\zeta_{t+1}}.$$

$\Delta z_t = log(g_t)$ evolves as previously described in equations 1 through 3.

I use value function iteration to solve the planner's stationary dynamic programming problem. I discretize the aggregate state space by constructing grids for each of the aggregate state variables. I use 20 grid points in the grid for the aggregate capital stock, the values for these grid points range from 75% below to 75% above the steady state level. In order to evaluate the mathematical expectation in the planner's value function I use Markov chain approximations to the continuous processes that describe the evolution of the other state variables. I describe my discrete approximation method for the processes in equations 1 through 3 at length in the next subsection of the paper. I use the method from Rouwenhorst (1995) to construct a Markov chain approximation to the continuous process for the probability of a disaster. This method has been shown to approximate autoregressive processes well, even those that are very highly persistent. The grid points and transition matrix are chosen in a manner such that the discrete approximation will match the unconditional mean and variance of the continuous process exactly. I construct a grid for the probability of a disaster containing five grid points that span the range

$$\left[log(\bar{\pi}) - \sqrt{\frac{4}{1-\rho_\pi^2}}\sigma_\upsilon, log(\bar{\pi}) + \sqrt{\frac{4}{1-\rho_\pi^2}}\sigma_\upsilon \right]$$

I use the FFSQP FORTRAN algorithm for numerical optimization to maximize the value function during each iteration in the value function iteration routine. Since this numerical optimization procedure optimizes over the continuous state space, as opposed to my discretized representation, I need to interpolate function evaluations of the value function. I construct piecewise continuous cubic hermite splines to interpolate the value function. Finally, I used Howard's improvement algorithm to decrease the number of iterations necessary for the value function iteration routine to converge to the solution.

Discrete Approximation

In this section of the paper I describe the method used to discretely approximate the continuous process for aggregate productivity growth described by equations 1 through 3. I begin by using the method from Rouwenhorst (1995) to construct a Markov chain approximation for equation 3. This volatility process is simply a Gaussian AR(1) process. As discussed in the previous subsection the Rouwenhorst method has been shown to approximate this type of continuous process well, even when the process in close to unit root. I create a grid for the values of $log(\sigma_t^2)$ containing M grid points, and an MxM transition matrix containing the probability of transitioning to each of the M values for $log(\sigma_{t+1}^2)$ tomorrow given today's value. This transition matrix depends only on the number of grid points M and the persistence of this process ρ_σ. The values for the M grid points are spread evenly across the range

$$\left[\mu_\sigma - \sqrt{\tfrac{M-1}{1-\rho_\sigma^2}}\sigma_\nu, \mu_\sigma + \sqrt{\tfrac{M-1}{1-\rho_\sigma^2}}\sigma_\nu\right],$$

which depends upon σ_ν. In all of the results presented in this paper I use $M = 5$ grid points.

Next, I must approximate the continuous processes in equations 1 and 2. However, these processes are not Gaussian, due to the stochastic volatility. To discretely approximate these continuous processes, I extend the Rouwenhorst method. It is important to note that the specification of the transition matrix, which governs the movement across states over time, does not depend on σ_t. It only depends on the number of grid points to be used, and the persistence parameter of the process to be approximated. As such, I can create a single transition matrix for the aggregate productivity growth process and a single transition matrix for the long-run component process, neither of which depend on the current value of σ_t. Each of the two transition matrices is NxN. I next create M grids, each grid containing N grid points for the value of aggregate productivity growth. The N grid points in each grid are spread evenly across the range

$$\left[\mu - \sqrt{N-1}\sigma_i, \mu + \sqrt{N-1}\sigma_i\right] \qquad i = 1, 2, ..., M,$$

one grid for each of the M discrete values of σ_t. A large value of σ_i means the grid covers a wider range of values, and a smaller value of σ_i means the grid covers a narrower range of values. I use the same method to create M grids, each containing N grid points for the value of the long-run component. The N grid points in each grid are spread evenly across the range

$$\left[-\sqrt{\tfrac{N-1}{1-\rho_x^2}}\varphi_x\sigma_i, \sqrt{\tfrac{N-1}{1-\rho_x^2}}\varphi_x\sigma_i\right] \qquad i = 1, 2, ..., M.$$

In all of the results presented in this paper I use $N = 5$ grid points.

In order to test whether or not this technique reasonably approximates the underlying continuous process I performed a Monte Carlo experiment. I simulated the continuous processes and the discrete Markov approximations each for a large number of periods. Table 1 contains the first four unconditional moments obtained from simulating the continuous processes for Δz_t, x_t, and σ_t, and the same moments obtained from simulating my discrete approximations. The first two columns report these moments for the processes without any stochastic volatility, $\sigma_\nu = 0$. Under these circumstances Rouwenhorst's method guarantees that the mean and variance of the discrete approximation for each of the processes exactly matches its counterpart from the continuous process. We notice also that this method approximates the skewness and kurtosis of these processes well. The results in the third and the fourth columns report the same set of moments, now with a considerable amount of stochastic volatility. The results from Rouwenhorst (1995) continue to guarantee that the discrete approximation will exactly match the mean and variance of the process for $log(\sigma_t^2)$. However, we can no longer rely on that to be the case for the, now non-Gaussian, processes for Δz_t and x_t. The results in Table 1 suggest that the discrete approximation method for these continuous processes performs reasonably well. We also notice that the inclusion of stochastic volatility here dramatically increases the kurtosis in the distribution of aggregate productivity growth.

I also test that my approximation technique is able to match the conditional moments of these distribution as well. The full results from those Monte Carlo experiments can be found in an appendix. The results confirm that this discrete Markov technique approximates the true continuous processes reasonably well. I utilize the technique when solving the model to obtain all of the results in this paper.

Table 1.1: Unconditional Moments

	Continuous	Discrete	Continuous	Discrete
$E[\Delta z_t]$	0.012	0.012	0.012	0.012
$\sigma(\Delta z_t)$	0.010	0.010	0.006	0.004
$skew(\Delta z_t)$	0.000	0.000	-0.002	0.003
$kurt(\Delta z_t)$	2.96	2.90	34.0	30.9
$E[x_t]$	0.000	0.000	0.000	0.000
$\sigma(x_t)$	0.007	0.007	0.004	0.003
$skew(x_t)$	0.000	0.000	0.003	0.003
$kurt(x_t)$	2.95	2.83	35.3	32.8
$E[log(\sigma_t^2)]$	-15.5	-15.5	-15.5	-15.5
$\sigma(log(\sigma_t^2))$	0.000	0.000	3.10	3.10
$skew(log(\sigma_t^2))$	0.000	0.001	0.039	0.000
$kurt(log(\sigma_t^2))$	2.71	2.66	2.68	2.66

1.2.6 Asset Pricing

I use the model in this paper to examine the prices of two assets. One is an equity share in the firm, which entitles its owner to an infinite stream of dividends. The other is a one period risk-free bond that entitles its owner to a single unit of consumption next period. These assets are priced according to the standard Euler equations; one for the price, p_t, of the equity asset

$$p_t = \beta E\big[M_{t+1}(d_{t+1} + p_{t+1})\big],$$

and one for the price, q_t, of the risk-free bond

$$q_t = \beta E\big[M_{t+1}\big].$$

M_{t+1} is the asset-pricing kernel

$$M_{t+1} = \left(\frac{C_{t+1}}{C_t}\right)^{-1} \left(\frac{\tilde{C}_{t+1}}{\tilde{C}_t}\right)^{\gamma} \left(\frac{U_{t+1}}{\mu_t}\right)^{1-\gamma-\delta},$$

which is standard for models whose agents have Epstein-Zin preferences. The last term in this equation depends on the certainty equivalent of uncertain future utility. Persistent or even permanent shocks, whose effects last for many periods, have a large impact on the price of assets through this term.

The firm pays out what is left over after labor and investment costs as dividends to their shareholders

$$d_t = Y_t - w_t n_t - I_t.$$

Wages are simply the marginal product of labor

$$w_t = (1-\alpha) K_t^{\alpha} (Z_t n_t)^{-\alpha}$$

as determined by the firm's profit maximization condition.

I perform a transformation, as described in a previous subsection of the paper, to ensure stationarity of the price of the equity asset. This yields the Euler equation

$$\hat{p}_t = \beta E\left[M_{t+1}(\hat{d}_{t+1} + \hat{p}_{t+1})\right],$$

where

$$\hat{d}_t = g_t^{-1} \alpha \hat{K}_t^{\alpha} (g_t n_t)^{1-\alpha} - \hat{I}_t.$$

In this paper I examine the net return on both assets as defined by

$$r_{t+1}^{f} = \frac{1}{q_t} - 1,$$

and

$$r_{t+1}^{e} = \frac{p_{t+1} + d_{t+1}}{p_t} - 1.$$

1.3 Estimation

Smith (1993) and Gourieroux, Monfort, and Renault (1993) first developed the Indirect Inference estimator. At its core this estimation procedure is similar in strategy to the more common Method of Moments estimation procedure. Method of Moments requires one to choose a set of moments in the data to target. Indirect Inference, instead, requires one to specify an auxiliary model. The parameters of this auxiliary model, when estimated using data, make up the set of moments that will be targeted. To avoid any confusion with the auxiliary model, I will refer to the model presented in Section 2 as the economic model for the remainder of this paper. Indirect Inference places very few restrictions on what can serve as the auxiliary model for estimation. The auxiliary model need not be linear, or Gaussian, or have any analytical solution. The only requirement is that one can solve for the auxiliary model's parameters using actual data or using data obtained by simulating the economic model. The auxiliary model parameters are first estimated using the real data. Then the auxiliary model parameters are estimated many more times using simulated data, where each time the simulated data has been obtained by simulating the economic model given a different set of economic model parameters. Indirect Inference selects the set of economic model parameters such that the auxiliary model parameters estimated using data obtained by simulating the economic model, given this set of economic model parameters, are as close as possible to the auxiliary model parameters estimated using the real data.

1.3.1 Auxiliary Model

For my auxiliary model in this paper I use a stochastic process similar to the one used to specify consumption and dividend growth in Bansal and Yaron (2004). My auxiliary model is a stochastic process that describes the joint behavior of the growth rate of consumption, g_t^c, and equity returns, r_t^e. The auxiliary model is described by the four equation system

$$g_{t+1}^c = \mu_c + x_t + \sigma_t \epsilon_{t+1} \tag{1.5}$$

$$r_{t+1}^e = \mu_e + \phi_e x_t + \varphi_e \sigma_t \eta_{t+1} \tag{1.6}$$

$$x_{t+1} = \rho_x x_t + \varphi_x \sigma_t \nu_{t+1} \tag{1.7}$$

$$log(\sigma_{t+1}^2) = (1 - \rho_\sigma)\mu_\sigma + \rho_\sigma log(\sigma_t^2) + \sigma_\omega \omega_{t+1}, \tag{1.8}$$

where x_t is a persistent shock to average consumption growth and equity returns and σ_t is a shock to the volatility at time t. ϵ_{t+1}, η_{t+1}, ν_{t+1}, and ω_{t+1} are all assumed to be independent and identically distributed $N(0,1)$. The four equations contain a total of nine auxiliary model parameters. Let $\boldsymbol{\beta} = (\mu_c, \mu_e, \mu_\sigma, \phi_e, \varphi_e, \rho_x, \varphi_x, \rho_\sigma, \sigma_\omega)$. These nine auxiliary model parameters, estimated using actual data, are the targets of my Indirect Inference estimation routine. My estimation procedure does not target the average risk-free rate, or the average equity premium, or the variance of either of these.

Instead I target the long-run risk and stochastic volatility in the data on consumption growth and equity returns. This distinction is crucial for interpreting the results in the paper. Previous papers have calibrated the persistence of consumption growth in order to account for asset price data. Bansal and Yaron (2004) and others have, in effect, answered the question: how do I need consumption to behave in order to resolve this set of asset-pricing puzzles? In this paper I first estimate the persistence of consumption growth in the data. Then I answer the questions: can my model account for this consumption behavior and, if so, does it resolve these asset-pricing puzzles?

Estimating the Auxiliary Model

Estimation of the auxiliary model is not a straightforward exercise since equations 5 through 7 are non-Gaussian, and all of the equations include variables that are unobservable in the data. I estimate the auxiliary model's parameters in two steps.

Step 1 In step one I ignore the stochastic volatility, and assume that the level of volatility is fixed. This yields the following three equations:

$$g^c_{t+1} = \mu_c + x_t + \sigma\epsilon_{t+1}$$

$$r^e_{t+1} = \mu_e + \phi_e x_t + \varphi_e \sigma \eta_{t+1}$$

$$x_{t+1} = \rho_x x_t + \varphi_x \sigma \nu_{t+1}.$$

These three equations describe a linear state space model. I follow Harvey (1989) to write the equations in the state space form (SSF) and use maximum likelihood

(ML) to estimate the parameters. I use the Kalman Filter to construct the prediction error decomposition form of the likelihood. I use the FFSQP Fortran algorithm to numerically maximize this likelihood function, yielding ML estimates for six of the nine parameters in my full auxiliary model. I also follow Harvey (1989) and construct the information matrix, providing me with an estimate of the covariance matrix of the ML estimates. I use this to compute the standard errors for my estimates.

Step 2 Next, I use a Kalman smoother to obtain an estimated time series of the unobserved state x_t. Using this time series for x_t and the estimated parameter values, I compute a time series of the residuals from the two measurement equations in Step 1. Let $c_{t+1} = \sigma_t \epsilon_{t+1}$ be the residuals from the consumption growth equation and $e_{t+1} = \varphi_e \sigma_t \eta_{t+1}$ be the residuals from the equity returns equation. Taking the natural log of the square of these residuals yields

$$log(c_{t+1}^2) = log(\sigma_t^2) + log(\epsilon_{t+1}^2)$$

$$log(e_{t+1}^2) = log(\varphi_e^2) + log(\sigma_t^2) + log(\eta_{t+1}^2).$$

This allows me to define another linear SSF

$$log(c_{t+1}^2) = E[log(\epsilon_{t+1}^2)] + h_t + \xi_{t+1}^c$$

$$log(e_{t+1}^2) = E[log(\eta_{t+1}^2)] + 2log(\varphi_e) + h_t + \xi_{t+1}^e$$

$$h_{t+1} = (1 - \rho_\sigma)\mu_\sigma + \rho_\sigma h_t + \sigma_\omega \omega_{t+1};$$

where

$$h_t = log(\sigma_t^2),$$

$$\xi_{t+1}^c = log(\epsilon_{t+1}^2) - E[log(\epsilon_{t+1}^2)],$$

and

$$\xi_{t+1}^e = log(\eta_{t+1}^2) - E[log(\eta_{t+1}^2)].$$

ξ_{t+1}^c and ξ_{t+1}^e are non-Gaussian, zero mean, white noises. Since ϵ_{t+1} and η_{t+1} are assumed to be $NID(0,1)$, the mean and variance of $log(\epsilon_{t+1}^2)$, as well as the mean and variance of $log(\eta_{t+1}^2)$, are known to be approximately -1.27 and $\pi^2/2$ respectively.[2]

Following Ruiz (1994) I treat ξ_t and ξ_t^e as though they are $NID(0, \pi^2/2)$ and estimate the parameter values by Quasi-Maximum Likelihood (QML). I once again use the Kalman Filter to estimate the unobserved stochastic volatility process, and I use prediction error decomposition to obtain the Quasi-Likelihood function. Numerically maximizing this function yields the QML estimates of the auxiliary model parameters in the stochastic volatility equation. Also following an appendix to Ruiz (1994) I compute the variance-covariance matrix for the QML estimates, and use this to compute standard errors.

A Recovery Exercise

I perform a recovery exercise in order to test the method I use to estimate the auxiliary model. I select a set of parameter values for $\boldsymbol{\beta}$. I treat these as the true values in the population. Using these values, I simulate the continuous processes using draws from a random standard normal distribution for each of the shocks. I simulate very

[2]see Abramovitz and Stegun (1970)

long time series for consumption growth and equity returns: 10,000 periods long. Next, I treat these long time series as my data, and I use this data to estimate the parameters of the auxiliary model following the method laid out in the previous subsection. Given this large amount of data the estimation technique should yield estimates of the parameter values that are close to the "true" values I selected at the start.

For each of the parameters in $\boldsymbol{\beta}$, Table 2 shows the true value and the estimated value. We can see that my estimation technique performs well in this exercise. Many of the parameter values are exactly recovered. There does appear to be some downward bias in the estimation of ρ_x, the persistence of the long-run component. However, it is important to note that the magnitude of this bias does not appear large enough to change any of the results presented in this paper.

Table 1.2: Auxiliary model parameters

	True	Estimated
μ_c	0.02	0.02
μ_e	0.03	0.03
μ_σ	-9.734	-9.732
ϕ_e	0.50	0.79
φ_e	0.75	0.75
ρ_x	0.95	0.82
φ_x	0.001	0.001
ρ_σ	0.45	0.45
σ_w	0.207	0.214

1.3.2 Indirect Inference Estimator

I estimate the key parameters of my model using Indirect Inference to target specific consumption and equity return dynamics. I use this method to estimate nine of the economic model's parameters. Two of the estimated parameters describe agent preferences; γ, which controls IES, and δ, which controls risk aversion. The third estimated parameter, ϕ, controls the magnitude of the capital adjustment costs. The next two estimated parameters, ρ_x and φ_x, control the persistence and the relative volatility of long-run productivity risk. Two more estimated parameters control the persistence, ρ_σ, and the standard deviation, σ_ν, of stochastic volatility in aggregate productivity. The final two estimated parameters, ρ_π and σ_υ, determine the volatility and the persistence of the shocks to the probability of a disaster event. Let $\boldsymbol{\xi}$ be a vector containing these nine parameters.

First, I estimate the auxiliary model's parameters using actual data. I call the vector of these estimates $\hat{\boldsymbol{\beta}}$. Next, I estimate the auxiliary model's parameters many more times using simulated data. Each time I solve the economic model given a different set of values for $\boldsymbol{\xi}$, and use this solution to obtain simulated data for consumption growth and equity returns. I use the simulated data from this particular solution to estimate the auxiliary model's parameters, and I call these estimates $\tilde{\boldsymbol{\beta}}(\boldsymbol{\xi})$.

Since I use post-war U.S. data to create my estimation targets $\hat{\boldsymbol{\beta}}$ and this data does not contain any episodes similar to the disasters in my economic model; I do not allow the simulated data to contain any realizations of a disaster event. This disaster risk is actually a Peso Problem, in the style of Danthine and Donaldson (1999). The agents in the model believe that there is the risk of a disaster that never actually occurs.

My Indirect Inference estimate of $\boldsymbol{\xi}$ minimizes the objective function

$$\left(\frac{\tilde{\boldsymbol{\beta}}(\boldsymbol{\xi}) - \hat{\boldsymbol{\beta}}}{\hat{\boldsymbol{\beta}}}\right)' \left(\frac{\tilde{\boldsymbol{\beta}}(\boldsymbol{\xi}) - \hat{\boldsymbol{\beta}}}{\hat{\boldsymbol{\beta}}}\right).$$

I use the DiRect (short for Dividing Rectangles) optimization algorithm to solve this optimization problem. This routine has a couple of features that make it ideal for this application. First, it requires the parameters space, over which it searches, to be bounded. Most of the parameters in $\boldsymbol{\xi}$ are bounded, so using an approach that naturally integrates this information about the parameters space into the search routine is useful. Second, the DiRect algorithm is globally convergent to the global optimum. The objective function is likely to contain many local minima, and so using a method that converges to the global minimum is absolutely critical.

1.4 Results

In this section of the paper I present my results. I begin with the empirical results, obtained by estimating the joint process for consumption growth and equity returns. Then I turn my focus to the theoretical results, obtained by Indirect Inference estimation of the economic model from Section 2 of the paper.

1.4.1 Empirical Estimation of the Auxiliary Model

I estimate the auxiliary model's parameters first using actual data. For my data on consumption I use the quarterly series of non-durables and services consumption from 1947 Q1 to 2015 Q3 from the BEA. For data on equity returns I use the monthly value-weighted return on the S&P 500 with dividends from CRSP, and aggregate these values to the quarterly level.

Table 3 contains the estimated parameter values and their standard errors. The data reject the idea that the process for consumption growth and equity returns is IID, as φ_x is significantly larger than zero. The growth rate shock is also rather persistent, ρ_x is significantly larger than zero. However, this persistence is significantly lower than the amount typically assumed in the literature on long-run risks. For example, Schorfheide, Song, and Yaron (2017) estimate the persistence to be 0.92 .

Table 1.3: Auxiliary Model Parameters

	Estimate	S.E.
μ_c	0.0047	0.00055
μ_e	0.029	0.0052
μ_σ	-11.67	1.62
ϕ_e	5.25	1.88
φ_e	20.05	2.22
ρ_x	0.71	0.1
φ_x	0.65	0.21
ρ_σ	0.78	0.53
σ_ω	0.27	0.28

There is also evidence of stochastic volatility in the data on consumption growth and equity returns. ρ_σ is estimated to be 0.78 and σ_ω is estimated to be 0.27. However, these parameters are not quite as precisely estimated, as their standard errors are somewhat larger than the standard errors for the other estimated parameters.

1.4.2 Indirect Inference Estimation of the Economic Model

For the first set of parameters values $\{\beta, \theta, \alpha, \delta_K\}$ I choose values that are common across the real business cycle literature. Next, I set the average growth rate of productivity, μ, equal to the average growth rate of consumption in the data. I choose the average volatility in the stochastic volatility process for aggregate productivity, μ_σ, to match the average volatility of consumption in the data. I use the values for the mean and variance of the size of the rare disaster, μ_ζ and σ_ζ, and the average probability of a disaster, $\bar{\pi}$, from Gourio (2012). Finally, I estimate the remaining parameters in my economic model using Indirect Inference as described in the previous section of the paper.

Table 4 contains a complete list of the economic model parameters from the baseline estimation of the model, with the economic model parameters estimated using Indirect Inference in the second column. I estimate the Arrow-Pratt coefficient of risk aversion to be about five, which is within the range accepted by the literature. I estimate that the IES is approximately one third. In stark contrast to the results in Bansal Yaron (2004) and other papers, which critically rely on a value for the IES that is greater than one.

Table 1.4: Economic model parameter values

β	0.99	γ	-2.21
θ	0.205	δ	5
δ_K	0.015	ϕ	6.47
α	0.34	ρ_x	0.3
μ	0.004758	φ_x	1.14
μ_σ	-9.734	ρ_σ	0.78
μ_ζ	-0.01123	σ_ω	0.2
σ_ζ	0.092	ρ_π	0.25
$log(\bar{\pi})$	-8.8537	σ_υ	0.34

Next, Table 5 contains estimates of the auxiliary model parameters. The first column reports the estimates obtained using the real data. As noted in the previous subsection, there is evidence for both long-run risk and stochastic volatility in the data on consumption and equity returns. However, the persistence of both the long-run process and the stochastic volatility process does not appear to be nearly as high as what is assumed in Bansal and Yaron (2004). The data suggests that their results may rely on shocks to the growth rates and volatilities that have effects that are too long-lasting.

Table 1.5: Auxiliary Model Parameters

	Data	Model
μ_c	0.0047	0.0053
	(0.00055)	
μ_e	0.029	0.016
	(0.0052)	
μ_σ	-11.67	-8.65
	(1.62)	
ϕ_e	5.25	4.76
	(1.88)	
φ_e	20.05	0.49
	(2.22)	
ρ_x	0.71	0.73
	(0.1)	
φ_x	0.65	0.33
	(0.21)	
ρ_σ	0.78	0.87
	(0.53)	
σ_ω	0.27	0.25
	(0.28)	

The second column of Table 5 contains the auxiliary model parameters estimated using data obtained from simulation of the estimated economic model. Overall, the estimates from the model data are similar to the estimates from the real data. In particular, the economic model does a good job of accounting for both the level of long-run risk and the level of stochastic volatility that we observe in the data on consumption and equity returns. The fact that the model, which we now know delivers empirically plausible comsumption dynamics, estimates a value for the IES of roughly one third casts serious doubt on the assumption that IES is greater than one.

The model is very clearly not able to account for the fact that equity returns are roughly 20 times more volatile than consumption growth in the data. This lack of volatility in returns is not an uncommon problem in the production-based asset-pricing literature. I defer a more thorough discussion of this finding until the next subsection of the paper.

For now, I turn my focus to the asset pricing implications of the estimated economic model. Table 6 contains the mean and the standard deviation of both the risk-free rate and excess returns. The first column reports the data moments, and the second column reports the moments implied by the estimated economic model. The third column contains the same moments from the results reported in Croce (2014). All of the moments are reported in annual percentage form.

While the estimated economic model in this paper does improve upon the canonical business cycle model, which delivers an equity premium that is essentially zero, it is easy to see that the model does not resolve these asset-pricing puzzles. To put this result into context it is important to remember that my Indirect Inference estimation procedure did not target any of the moments in Table 6. The model in Croce

Table 1.6: Annual percentages

	Data	Model	Croce
$E[r^f]$	0.65	4.22	0.94
	(0.38)		
$E[r^e - r^f]$	4.06	2.02	4.31
	(2.25)		
$\sigma(r^f)$	1.86	2.11	0.94
	(0.32)		
$\sigma(r^e - r^f)$	20.89	4.88	12.47
	(2.21)		

(2014) is calibrated in order to account for this set of facts. This suggests that the results in Croce (2014) might rely on consumption behaviour that is counterfactual. It suggests that long-run risk and stochastic volatility in consumption growth cannot alone resolve asset pricing puzzles. The estimated economic model presented here is only able to account for about half of the equity premium we observe in the data.

1.4.3 Shutting Down Disaster Risk or Long-Run Productivity Risk

I estimate two alternative specifications of my economic model in addition to the baseline estimation presented in the previous subsection. In the first I assume that there is no risk of a rare disaster ever occurring. In the second I assume that there is no long-run risk or stochastic volatility in the process for aggregate productivity growth. I use these two estimates of the economic model to examine the role each of the features plays in accountinf for the data.

Table 7 contains the auxiliary model parameters that control long-run risk and stochastic volatility in consumption growth and equity returns. The first column presents the parameter values estimated using the actual data, the second column contains the parameter values estimated using data obtained from the baseline estimation of the economic model, and the third column contains the parameter values estimated using data obtained by simulating the economic model estimated without disaster risk. Shutting down the disaster risk in the model does not significantly diminsh the model's ability to account for this set of facts in the data.

Table 1.7: Auxiliary Model Parameters

	Data	Model	No D.R.
ρ_x	0.71	0.73	0.68
φ_x	0.65	0.33	0.28
ρ_σ	0.78	0.87	0.77
σ_ω	0.27	0.25	0.27

Table 8 compares the asset pricing implications of the economic model estimated with or without disaster risk. Again we can see that the model's performance in this regard is not significantly altered by the inclusion of disaster risk. This does not mean that the model with disaster risk is not *capable* of delivering a lower risk-free rate, larger equity returns, or more volatile excess returns than the model without disaster risk. It only means that these two particular models estimated in this fashion, to deliver certain consumption dynamics, do not differ significantly in this regard.

Table 1.8: Annual percentages

	Data	Model	No D.R.
$E[r^f]$	0.65	4.22	4.20
$E[r^e - r^f]$	4.06	2.02	2.00
$\sigma(r^f)$	1.86	2.11	1.34
$\sigma(r^e - r^f)$	20.89	4.88	4.41

Table 9 compares some business cycle quantity moments implied by the economic model estimated with and without disasters. It shows that the model estimated without disaster risk implies volatility in both output and investment that are significantly higher than the model with disaster risk. This result shows that the full economic model with disaster risk is better able to account for business cycle quantity data, without diminishing its ability to account for asset pricing data.

Table 1.9: Quarterly standard deviations

	Data	Model	No D.R.
$\Delta log(Y)$	1.04	1.40	1.84
$\Delta log(I)$	2.79	3.11	3.83
$\Delta log(n)$	0.91	0.93	0.86

Next, I consider the economic model without long-run risk or stochastic volatility in aggregate productivity growth. Table 10 contains the auxiliary model parameters estimated from the actual data, data simulated from the economic model with long-run productivity risk, and data simulated from the economic model without it. It shows that the model without long-run productivity risk is much less able to account for the long-run risk and stochastic volatility in consumption growth and equity returns. Together with the results from Table 7, we see that most of the model's ability to deliver these empirically plausible consumption dynamics comes through the model's long-run productivity risk.

Table 1.10: Auxiliary model parameters

	Data	Model	No L-R. R.
ρ_x	0.71	0.73	0.56
φ_x	0.65	0.33	1.28
ρ_σ	0.78	0.87	0.74
σ_w	0.27	0.25	0.64

1.4.4 Forcing the Model to Account for Asset Prices

In this subsection I perform an experiment. The baseline estimation of my economic model was not able to account for the fact that equity returns are roughly 20 times more volatilie than consumption growth. Is the model even *capable* of generating more volatile equity returns? I re-estimate the economic model several times, each time placing a larger weight in the object function for my Indirect Inference routine on the model's relative error in regarding the auxiliary model parameter φ_e.

Table 1.11: Auxiliary model parameters

	Data	Model	x2	x5	x10
φ_e	20.05	0.49	1.06	2.78	4.83
ρ_x	0.71	0.73	0.62	0.43	0.42
φ_x	0.65	0.33	0.61	1.84	3.00
ρ_σ	0.78	0.87	0.73	0.55	0.49
σ_w	0.27	0.25	0.40	0.48	0.49
IES		0.31	1.49	4.55	14.29

Table 11 shows the results from this experiment. We see that the model is, in fact, able to deliver more volatile equity returns. However, we see that the model's improvement along this dimension comes at the expense of its ability to account for long-run risk and stochastic volaility in consumption growth. As the volatility of equity returns increases, the persistence of the long-run risk shocks decreases, and the processes become more *iid*. It appears that impoving this model's ability to account for asset pricing data comes at the expense of its ability to deliver the empirically supported consumption dynamics.

1.4.5 Financial Leverage

Next, I estimate the model using a levered equity claim. Each period the representative firm finances a fixed fraction of their capital using debt. Dividends are now given by

$$d_t = Y_t - w_t n_t - I_t - \lambda r_t^f \bar{k}.$$

Since my auxiliary model targets equity returns I must now estimate the model again using this levered equity claim. Table 12 shows the asset pricing implications of this specification of the model, when I assume $\lambda = 0.018$ I observe that model is now better able to account for the high volatility in equity returns.

Table 1.12: Asset return moments

	Data	Baseline	Levered
$E[r^f]$	0.23	1.06	1.06
$E[r^e - r^f]$	2.68	0.5	0.53
$\sigma(r^f)$	0.8	0.53	0.81
σ_π	7.59	1.22	2.17

1.5 Conclusion

In this paper I estimate a joint process for consumption growth and equity returns, and show that there is evidence in the data that both time series contain long-run risk and stochastic volatility. Next I estimate a DSGE model with production that contains long-run risk and stochastic volatility in aggregate productivity and rare disasters that destroy a fraction of the capital stock and occur with time-varying probability. I use the previously estimated joint process for consumption growth and equity returns as my auxiliary model.

I show that this model can account for the level of long-run risk and stochastic volatility that we observe in the data on consumption growth and equity returns. Although the model's asset-pricing implications improve upon the canonical RBC model, they still fall short of solving many key asset pricing puzzles. This shows that previous research may have relied on counterfactual assumptions about the behavior of consumption and returns. This also opens the door to future research that seeks to explain this still puzzling set of facts.

The model's performance suffers greatly if we eliminate one of the model's two key features, long-run productivity risk or rare disasters. I also find that incorporating financial leverage into the model is able to improve some of its asset-pricing shortcomings.

Chapter 2

European Call Options in a Production Economy Model

2.1 Introduction

An option gives its owner the right to either buy or sell an underlying asset for a specific price at a specific date, or in some cases at specific dates, in the future. The option to buy is referred to as a call option, and the option to sell is referred to as a put option. The price the owner is entitled to buy or sell for in the future is called the strike price.

Observing the patterns over time across the prices of various options can tell us about how likely different future states of the world appear to be to investors, as well as how desirable or undesirable those states are. Options offer a glimpse into their buyers' beliefs about the distribution of possible outcomes for the underlying asset. They can also help us to measure how volatile investors perceive the underlying asset to be. For these reasons the information contained in options pricing data is of interest to those who wish to explain many features of real-world financial markets. In this chapter of my dissertation, I use the theoretical model presented in Chapter 1 to price several different options. I analyze several of the key patterns in the prices across different options, and evaluate the model's ability to account for some well-known phenomena in the data on options prices.

A European option, either a call or a put, can only be exercised on one particular day in the future. Whereas an American option can be exercised at any point up until a particular day in the future. A call option whose strike price is the current price of the underlying asset, known as an "at the money" (ATM) option, is a bet that the underlying asset will increase in value. If this takes place, the owner of the option can buy the asset at the strike price and sell it at the new higher price, earning a positive return. An ATM put option is precisely the opposite, a bet that the underlying asset will decrease in value. Options need not be purchase ATM. A call option with a strike price that is below the current price of the underlying asset is called "in the money" (ITM), and a call with a strike price above the current price is called "out of the money" (OTM). For a put option ITM and OTM are reversed. ITM and OTM calls and puts can be used to form hedges against a great number of positions that an investor can take in the underlying asset. I focus on European call options with varying strike prices exclusively in what follows.

In this chapter, I use the full version of the production economy model first presented in Chapter 1 of the dissertation. This model is estimated using the same Indirect Inference proceedure described there. The parameter values used in this chapter, therefore, correspond to the baseline parameter estimates from Chapter 1.

I price five new financial assets in addition to the risk-free asset and the equity asset from the previous chapter. These are European call options with a three-month time to expiration where the underlying financial asset is the return to the capital stock. These five call options cover a range of strike prices. Two are ITM, one with a strike price ten percent below the current market price and the other with a strike price five percent below the current market price. One is an ATM call option. The final two are OTM, one with a strike price five percent above the current market price and the other with a strike price ten percent above the current market price. Once the pricing functions have been estimated, I first examine the patterns of option value across levels of current capital for each of the five options. The ITM, ATM, and OTM call options display three distinct patterns of value across today's level of capital. The ITM call option with strike price five percent below the current market price exhibits a value that is increasing in today's level of capital. The ATM call option exhibits a value that is slightly decreasing in today's level of capital. Finally, the OTM call option with strike price five percent above the current market price is more sharply decreasing in today's level of capital. These patterns demonstrate the multiple forces at play in valuing these options. Since the price of the capital stock is nearly linear in today's capital, as the level of today's capital increases so does the price of the option's underlying asset. However, when today's capital level is well above the steady state level the capital stock is likely to decrease, and therefore the price of capital is likely to be lower, decreasing the ATM or OTM option's value. The opposite is true for levels of today's capital that are far below the steady state level. The strike price also plays a crucial role in determining the number of states in which the option is likely to be exercised at expiry. A deep ITM option is almost surely going to be exercised, and thus its value in relationship to today's level of capital is more similar to the relationship for underlying financial asset. In the case of the ATM and OTM options, the other effects dominate and the relationship with today's level of capital is reversed.

Black and Scholes (1973) introduced a seminal model for the pricing of options, which continues to be used widely today. The Black-Scholes options pricing equation gives the theoretical value of a particular option as a function of, among other things, the volatility of the underlying asset. Inverting this function, holding all else constant, gives us the volatility of the underlying asset as a function of the Black-Scholes theoretical option value. Plugging the market price of this option in for the theoretical value yields the Black-Scholes implied volatility. This implied volatility is the unique value for the volatility of the underlying asset for which the Black-Scholes option pricing equation yields the market price of the option. There is no closed form equation for this inverse of the Black-Scholes equation, however we can solve for this implied volatility numerically. Solving for the implied volatility using the price of options with the same underlying asset but different strike prices and expiries yields an implied volatility surface. In this paper I focus on the implied volatility computed using the return to capital as the underlying asset for options with various strike prices. The Black-Scholes model from which their equation is derived assumes that the underlying asset follows a Brownian motion process with a constant volatility σ. If the Black-Scholes model was "correct" and the market price for each possible option with this same underlying asset was computed using the Black-Scholes equation then implied volatility, computed in the manner discussed previously, would yield the constant volatility term σ; no matter what the expiry or the strike price was. However, there is a great deal of empirical evidence that contradicts this flat implied volatility surface[1]. Options on individual stocks or options with a relatively short maturity tend to yield a pattern of implied volatilities across strike prices known as a volatility

[1] See Rubinstein (1985) and Mixon (2002) for some examples

smile. This pattern produces higher values for implied volatility from options that are either deep ITM or OTM, as opposed to options that are ATM. Cao (1992) and Bates (1996) provide evidence that foreign currency options also tend to display an implied volatility smile. Options where either the underlying asset is broad, such as a market index, or the time to expiration is relatively long tend to yield a pattern known either as a volatility smirk or a volatility skew. These terms refer to a pattern of implied volatilities that are generally decreasing in the strike price. Bates (2000) shows that since the U.S. stock market crash of 1987, S&P 500 options have exhibited an implied volatilty skew.

I simulate the estimated model to obtained data from this theoretical economy. I use the simulated options price data to solve numerically for measures of the Black-Scholes implied volatility at five different strike prices. Since the underlying asset is analogous to a broad index asset in the data we would hope to observe a volatility smirk or skew. The simulated data from the estimated model does, in fact, produce this pattern. I also propose a statistic to measure the average level of the volatiity skew in the simulated data.

A long-run risk is a persistent shock to the growth rate of a time series. This type of shock stands in contrast to the traditional idiosyncratic shocks in the business cycle framework that affect the level of aggregate productivity. Bansal and Yaron (2004) first suggested that long-run risk in consumption might be critical to accounting for various asset pricing puzzles, such as the equity premium puzzle and the low risk-free rate puzzle. Their paper examined an endowment economy in which the exogenous process for consumption growth and dividend growth contained time-varying levels of volatility, known as stochastic volatility, as well as long-run risk. They showed that by assuming consumption growth and dividend growth displayed these particular dynamics, their model was able to account for many of the puzzling aspects of asset price data. Croce (2014) added long-run risk in aggregate productivity growth to a production economy model. He showed that this feature improved the model's asset pricing implications.

There are also a number of papers that incorporate some form of disaster risk in order to resolve certain asset pricing puzzles[2]. Gourio (2012) models a production economy with disasters that both negatively affect productivity and destroy a fraction of the capital stock. In this model the probability of a disaster varies persistently over time and the size of this disaster is also a random variable. Whereas in the standard business cycle framework risk premia move in concert with the cycle, they calibrate the model in their paper to account for the countercyclical nature of the equity premium in the data. In the first chapter of this dissertation I demonstrated how the inclusion of both long-run productivity risk and disaster risk in a production economy model allowed the model to better match both business cycle quantity and asset pricing moments simultaneously.

Models containing these ingredients have proliferated in the macroeconomic literature devoted to examining asset pricing puzzles. However, this literature tends to focus on a very narrow set of financial assets, generally a risk-free asset and a single levered or unlevered equity asset. In this chapter I expand on this set of financial assets. I examine the price of several European call options for strike prices ranging from ten percent ITM to ten percent OTM, each with the return to capital as their underlying financial asset. I explore how these modelling features that have been put forward to resolve asset pricing puzzles affect the behavior of options prices.

The model used here, and originally presented in Chapter 1 of the dissertation,

[2]See Rietz (1988), Barro (2006), and Gabaix (2012) for some examples.

builds upon the standard business cycle framework by incorporating recursive preferences, convex capital adjustment costs, long-run risk and stochastic volatility in productivity growth, and disaster risk. I investigate how each of the options pricing phenomena described above are affected when certain model features are shut down.

I explore how shutting down some of the key model features affects the pattern in the relationship between option price and today's level of capital across the various strike prices. I show that the pattern in the direction of the relationship between the two across strike prices does not change when these model features are shut down. However, the magnitude of the negative relationship for ATM and OTM options is increased when long-run productivity risk is shut down. The same effect is not observed when disaster risk is shut down. This result highlights the difference between the effect that long-run productivity risk, which increases the probability of both extremely good and extremely bad states, and disaster risk, which only increases the probability of extremely bad states, have on call options.

I also explore how many of the model's key features affect my measure of the average level of implied volatility skew. I show that a model that roughly corresponds to the standard business cycle framework does exhibit an implied volatility skew. However the addition of either long-run productivity risk or disaster risk increases the average level of implied volatility skew, and the addition of both model features increases the average level of implied volatility skew even farther. I also show that the addition of the first, whichever it might be, of these two features has a greater impact on the average level of implied volatility skew than the addition of the second.

The rest of this chapter is organized as follows. Section 2 re-introduces the full version of the production economy model and introduces the Euler equation used to price the European call options. Section 3 describes the estimation procedure as well as the algorithm used to solve for the Black-Scholes implied volatility numerically. Section 4 presents this chapter's main results, and Section 5 concludes.

2.2 Model

In this section of Chapter 2, I re-introduce the DSGE model of a production economy from Chapter 1. This representative agent model builds upon the standard real business cycle (RBC) framework. I begin by describing the representative consumer's preferences.

2.2.1 Consumer Preferences

Following Epstein and Zin (1989), consumers have recursive preferences

$$U_t(C_t, \ell_t, \psi_t) = \{(1-\beta)\tilde{C}_t^{\gamma} + \beta\psi_t^{\gamma}\}^{1/\gamma},$$

where $\tilde{C}_t$ is the per period utility from consumption, C_t, and leisure, ℓ_t,

$$\tilde{C}_t = C_t^{\theta}(Z_{t-1}\ell_t)^{1-\theta}$$

and ψ_t is the certainty equivalent of future uncertain utility

$$\psi_t = (E[U_{t+1}(C_{t+1}, \ell_{t+1}, \psi_{t+1})^{1-\delta}])^{1/1-\delta}.$$

I assume leisure is multiplied by the previous period's level of aggregate productivity, Z_{t-1}, in order to ensure balanced growth.

In a model where the representative consumer has expected utility preferences the intertemporal elasticity of substitution (IES) and the coefficient of relative risk aversion are inverses of one another. Recursive preferences include one additional parameter in order to break this relationship apart. In this model, IES is $(1-\gamma)^{-1}$ and the Arrow-Pratt coefficient of risk aversion is δ. Using a model with recursive preferences Tallarini (2000) was able to show that, while holding IES constant at one, increasing the level of risk aversion improved the model's ability to account for asset-pricing data while having almost no impact on the model's implications for business cycle quantities. This demonstrated the value of separating risk aversion from IES.

Empirical estimates of IES vary widely. However, many papers, including Bansal and Yaron (2004), have highlighted mechanisms that rely on IES to be greater than one in order to generate a large equity risk premium. In this chapter, as in Chapter 1, I estimate the parameters that control IES and risk aversion. I do not restrict risk aversion to be large, nor the IES to be greater than one.

2.2.2 Production

In this model, the representative firm produces output, Y_t, according to a Cobb-Douglas production function

$$Y_t = K_t^{\alpha}(Z_t n_t)^{1-\alpha},$$

where K_t is capital, n_t is labor input, and Z_t is aggregate productivity. The capital stock evolves according to the law of motion

$$K_{t+1} \leq (1-\delta_K)K_t + I_t - G_t K_t,$$

where capital depreciates at the rate δ_K, and I_t is investment. Following Jermann (1998) this law of motion for capital contains a convex capital adjustment cost, given by G_t. I assume a quadratic form for the cost

$$G_t = \frac{\phi}{2}\left(\frac{K_{t+1}}{K_t} - 1\right)^2.$$

In a model without capital adjustment costs, the supply of capital is perfectly elastic. Capital adjustment costs are of crucial importance as they create fluctuations in the price of capital. I estimate ϕ, which controls the size of the cost.

The economy in the model is subject to two aggregate resource constraints

$$C_t + I_t \leq Y_t,$$

and

$$l_t + n_t \leq 1.$$

2.2.3 Productivity Process

The process for aggregate productivity growth includes both long-run risk and stochastic volatility, as in Croce (2014). Let Z_t be aggregate productivity and z_t be the natural log of its value. Then

$$\Delta z_{t+1} = log(Z_{t+1}) – log(Z_t)$$

approximates the growth rate of aggregate productivity. Aggregate productivity growth evolves according to

$$\Delta z_{t+1} = \mu + x_t + \sigma_t \epsilon_{t+1}, \tag{2.1}$$

where μ is the average growth rate of aggregate productivity and ϵ_{t+1} is an *iid* shock to the level of aggregate productivity growth, with standard deviation σ_t at time t. x_t is a shock to the average growth rate of aggregate productivity. This persistent shock to the average growth rate, which evolves according to the AR(1) process

$$x_{t+1} = \rho_x x_t + \varphi_x \sigma_t \eta_{t+1}, \tag{2.2}$$

is referred to as long-run risk. The long-run component, x_t is mean zero, its persistence is captured by ρ_x, and the relative volatility of this shock is given by φ_x. The modelling of these two separate shocks allows aggregate productivity growth to contain both an *iid* short-run component and a slow moving long-run component. Most of the long-run risk literature assume that the long-run component is small but extremely persistent. This means that the effects of a large realization of the shock to this long-run component would be extremely long-lasting. In this chapter I estimate the persistence and relative volatility of this long-run component.

The natural log of the variance, $log(\sigma_t^2)$, of both shocks is time-varying. This property is referred to as stochastic volatility. The stochastic volatility evolves according to another AR(1) process

$$log(\sigma_{t+1}^2) = (1 - \rho_\sigma)\mu_\sigma + \rho_\sigma log(\sigma_t^2) + \sigma_\nu \nu_{t+1}. \tag{2.3}$$

Here, μ_σ is the unconditional mean of the process and ρ_σ is its persistence. σ_ν is the standard deviation of the shock to the volatility process. Because of this stochastic volatility, the processes for the evolution of aggregate productivity growth and its long-run component are non-Gaussian. Through Monte Carlo simulation experiments it can be shown that an increase in σ_ν increases the kurtosis in the distributions of both Δz_{t+1} and x_t. The stochastic volatility can be interpreted as fluctuation in the overall level of uncertainty surrounding the fundamental variables of the economy. I also estimate the persistence and volatility of this stochastic volatility process in this chapter.

With some simple algebraic manipulation of the terms in equation 1 we can arrive at

$$log(Z_{t+1}) = \mu + x_t + log(Z_t) + \sigma_t \epsilon_{t+1},$$

which makes it obvious that the evolution of aggregate productivity contains a unit root. As a result, the process is non-stationary and the effects of this shock are permanent. In order for me to be able to solve the model I follow Linde (2009) and scale each variable by the level of aggregate productivity today. In what follows I use a hat symbol to denote the transformed quantities. $\hat{C}_t = C_t/Z_t$, $\hat{Y}_t = Y_t/Z_t$, and so on. Since K_{t+1} is a date t choice variable $\hat{K}_{t+1} = K_{t+1}/Z_t$, and thus $\hat{K}_t = K_t/Z_{t-1}$. Hours worked, n_t, need not be transformed since the functional form for per period utility is specified to ensure that it is stationary. Along the balanced growth path in this economy all of the other variables will grow at the common rate μ from equation 1.

2.2.4 Disasters

Each period there is some probability, π_t, that a disaster event will occur. These disasters are assumed to be rare, meaning that the average probability of a disaster is quite small, but very severe. This stands in stark contrast to the type of risk, typically to aggregate productivity, modelled by the standard real business cycle framework. The more traditionally modelled risks are smaller in magnitude, but negative realizations of these shocks occur much more frequently. As far back as Reitz (1988), rare disaster risk has been thought to explain asset-pricing puzzles. In this chapter, similar to the model in Gourio (2012), I assume that a disaster event destroys a fraction of the capital stock in the economy. The size of this disaster event is random with

$$\zeta_t \sim NID(\mu_\zeta, \sigma_\zeta^2),$$

where ζ_t determines the fraction of the capital stock destroyed. To incorporate these disasters, I can re-write the law of motion for capital in the economy as

$$K_{t+1} \leq \bigg((1-\delta_K)K_t + I_t - G_t K_t\bigg)e^{D_{t+1}\zeta_{t+1}},$$

where D_{t+1} is an indicator function that takes on a value of one if a disaster event occurs. I assume that a disaster is realized only after the representative agent has made all decision for the current period t. The agent makes all decisions for the period fully aware of the current probability of a disaster π_t. However, only after all decisions have been made does D_{t+1} take on a value of zero or one; and if $D_{t+1} = 1$, only then is the size of the disaster event, ζ_{t+1} , realized.

The probability of a disaster event, π_t, changes over time according to

$$\log(\pi_{t+1}) = (1-\rho_\pi)\log(\bar{\pi}) + \rho_\pi \log(\pi_t) + \sigma_\upsilon \upsilon_{t+1}, \tag{2.4}$$

where $\bar{\pi}$ is the average probability of a disaster. σ_υ is the standard deviation of the shock to the probability of a disaster, and ρ_π governs the persistence of this shock. Gourio (2012) showed that if the probability of a disaster, where disasters were modelled similarly, was constant then these disaster events were observationally equivalent to a preference shock. The results from that paper suggested that time-variation in the probability of disaster events was crucial to resolving asset-pricing puzzles. In this chapter I estimate both ρ_π and σ_υ.

2.2.5 Solution Method

It is straightforward to show that the solution to the planner's problem is equivalent to the competitive equilibrium in this economy. I now define the planner's dynamic programming problem. The aggregate state of the economy is described by five state variables: aggregate capital stock K_t, aggregate productivity Z_t, long-run component x_t, volatility σ_t, and disaster probability π_t. Let $\boldsymbol{\kappa}$ be a vector containing these state variables. The planner optimizes over four choice variables: consumption C_t, hours worked n_t, investment I_t , and tomorrow's capital stock K_{t+1}. Let $\boldsymbol{\chi}$ be a vector of the choice variables. The planner's value function is therefore

$$V(\boldsymbol{\kappa}) = \max_{\boldsymbol{\chi}} \left((1-\beta)\left(C_t^{\theta}(Z_{t-1}(1-n_t))^{1-\theta} \right)^{\gamma} + \beta \left(E[V(\boldsymbol{\kappa'})^{1-\delta}] \right)^{\frac{\gamma}{1-\delta}} \right)^{\frac{1}{\gamma}},$$

where $1-n_t$ has been substituted in for ℓ_t since the time constraint always binds. The planner maximizes this value function subject to two non-linear inequality constraints

$$C_t + I_t \leq K_t^{\alpha}(Z_t n_t)^{1-\alpha}$$

and

$$K_{t+1} \leq \left((1-\delta_K)K_t + I_t - \tfrac{\phi K_t}{2}\left(\tfrac{K_{t+1}}{K_t} - 1\right)^2\right)e^{D_{t+1}\zeta_{t+1}}.$$

The first constraint is obtained by simply substituting the production function into the aggregate resource constraint. Aggregate productivity, Z_t , evolves according to equations 1 through 3. The second constraint is the law of motion for capital, where the probability that $D_{t+1} = 1$ is π_{t+1} from equation 4, and ζ_{t+1} is the previously specified *iid* normal random variable.

I follow Linde (2009) and transform this dynamic programming problem to ensure stationarity. The full algebraic derivation of the following equations can be found in an appendix to Chapter 1. Recall that variables with a hat have been divided by aggregate productivity ($\hat{C}_t = C_t/Z_t$ for example). Let $g_t = Z_t/Z_{t-1}$, then the planner's value function can be rewritten as

$$\hat{V}(\hat{\boldsymbol{\kappa}}) = \max_{\hat{\boldsymbol{\chi}}}\left((1-\beta)\left(\hat{C}_t^{\theta}(g_t^{-1}(1-n_t))^{1-\theta}\right)^{\gamma} + \beta\left(E[\hat{V}(\hat{\boldsymbol{\kappa}}')^{1-\delta}]\right)^{\frac{\gamma}{1-\delta}}\right)^{\frac{1}{\gamma}}.$$

The two non-linear inequality constraints can be rewritten as

$$\hat{C}_t + \hat{I}_t \leq g_t^{-\alpha}\hat{K}_t^{\alpha} n_t^{1-\alpha}$$

and

$$\hat{K}_{t+1} \leq \left(g_t^{-1}(1-\delta_K)\hat{K}_t + \hat{I}_t - \frac{\phi \hat{K}_t}{2g_t}\left(g_t \frac{\hat{K}_{t+1}}{\hat{K}_t} - 1 \right)^2 \right) e^{D_{t+1}\zeta_{t+1}}.$$

$\Delta z_t = log(g_t)$ evolves as previously described in equations 1 through 3.

I use value function iteration to solve the planner's stationary dynamic programming problem. I discretize the aggregate state space by constructing grids for each of the aggregate state variables. I use 20 grid points in the grid for the aggregate capital stock, the values for these grid points range from 75% below to 75% above the steady state level. In order to evaluate the mathematical expectation in the planner's value function I use Markov chain approximations to the continuous processes that describe the evolution of the other state variables. I describe my discrete approximation method for the processes in equations 1 through 3 at length in the next subsection of the chapter. I use the method from Rouwenhorst (1995) to construct a Markov chain approximation to the continuous process for the probability of a disaster. This method has been shown to approximate autoregressive processes well, even those that are very highly persistent. The grid points and transition matrix are chosen in a manner such that the discrete approximation will match the unconditional mean and variance of the continuous process exactly. I construct a grid for the probability of a disaster containing five grid points that span the range

$$\left[log(\bar{\pi}) - \sqrt{\frac{4}{1-\rho_\pi^2}}\sigma_v, log(\bar{\pi}) + \sqrt{\frac{4}{1-\rho_\pi^2}}\sigma_v \right]$$

I use the FFSQP FORTRAN algorithm for numerical optimization to maximize the value function during each iteration in the value function iteration routine. Since this numerical optimization procedure optimizes over the continuous state space, as opposed to my discretized representation, I need to interpolate function evaluations of the value function. I construct piecewise continuous cubic hermite splines to interpolate the value function. Finally, I use Howard's improvement algorithm to decrease the number of iterations necessary for the value function iteration routine to converge to the solution.

Discrete Approximation

In this section of the chapter I re-introduce the method used to discretely approximate the continuous process for aggregate productivity growth described by equations 1 through 3. I begin by using the method from Rouwenhorst (1995) to construct a Markov chain approximation for equation 3. This volatility process is simply a Gaussian AR(1) process. As discussed in the previous subsection the Rouwenhorst method has been shown to approximate this type of continuous process well, even when the process in close to unit root. I create a grid for the values of $log(\sigma_t^2)$ containing M grid points, and an MxM transition matrix containing the probability of transitioning to each of the M values for $log(\sigma_{t+1}^2)$ tomorrow given today's value. This transition matrix depends only on the number of grid points M and the persistence of this process ρ_σ. The values for the M grid points are spread evenly across the range

$$\left[\mu_\sigma - \sqrt{\tfrac{M-1}{1-\rho_\sigma^2}}\sigma_\nu, \mu_\sigma + \sqrt{\tfrac{M-1}{1-\rho_\sigma^2}}\sigma_\nu\right],$$

which depends upon σ_ν. In all of the results presented in this chapter I use $M = 5$ grid points.

Next, I must approximate the continuous processes in equations 1 and 2. However, these processes are not Gaussian, due to the stochastic volatility. To discretely approximate these continuous processes, I extend the Rouwenhorst method. It is important to note that the specification of the transition matrix, which governs the movement across states over time, does not depend on σ_t. It only depends on the number of grid points to be used, and the persistence parameter of the process to be approximated. As such, I can create a single transition matrix for the aggregate productivity growth process and a single transition matrix for the long-run component process, neither of which depend on the current value of σ_t. Each of the two transition matrices is NxN. I next create M grids, each grid containing N grid points for the value of aggregate productivity growth. The N grid points in each grid are spread evenly across the range

$$\left[\mu - \sqrt{N-1}\sigma_i, \mu + \sqrt{N-1}\sigma_i\right] \qquad i = 1, 2, ..., M,$$

one grid for each of the M discrete values of σ_t. A large value of σ_i means the grid covers a wider range of values, and a smaller value of σ_i means the grid covers a narrower range of values. I use the same method to create M grids, each containing N grid points for the value of the long-run component. The N grid points in each grid are spread evenly across the range

$$\left[-\sqrt{\tfrac{N-1}{1-\rho_x^2}}\varphi_x\sigma_i, \sqrt{\tfrac{N-1}{1-\rho_x^2}}\varphi_x\sigma_i\right] \qquad i = 1, 2, ..., M.$$

In all of the results presented in this chapter I use $N = 5$ grid points.

In order to test whether or not this technique reasonably approximates the underlying continuous process I perform a Monte Carlo experiment. I simulate the continuous processes and the discrete Markov approximations each for a large number of periods. Table 1 contains the first four unconditional moments obtained from simulating the continuous processes for Δz_t, x_t, and σ_t, and the same moments obtained from simulating my discrete approximations. The first two columns report these moments for the processes without any stochastic volatility, $\sigma_\nu = 0$. Under these circumstances Rouwenhorst's method guarantees that the mean and variance of the discrete approximation for each of the processes exactly matches its counterpart from the continuous process. We notice also that this method approximates the skewness and kurtosis of these processes well. The results in the third and the fourth columns report the same set of moments, now with a considerable amount of stochastic volatility. The results from Rouwenhorst (1995) continue to guarantee that the discrete approximation will exactly match the mean and variance of the process for $log(\sigma_t^2)$. However, we can no longer rely on that to be the case for the, now non-Gaussian, processes for Δz_t and x_t. The results in Table 1 suggest that the discrete approximation method for these continuous processes performs reasonably well. We also notice that the inclusion of stochastic volatility here dramatically increases the kurtosis in the distribution of aggregate productivity growth.

I also test that my approximation technique is able to match the conditional moments of these distribution as well. The full results from those Monte Carlo experiments can be found in an appendix to Chapter 1. The results confirm that this discrete Markov technique approximates the true continuous processes reasonably well. I utilize the technique when solving the model to obtain all of the results in this chapter.

Table 2.1: Unconditional Moments

	Continuous	Discrete	Continuous	Discrete
$E[\Delta z_t]$	0.012	0.012	0.012	0.012
$\sigma(\Delta z_t)$	0.010	0.010	0.006	0.004
$skew(\Delta z_t)$	0.000	0.000	-0.002	0.003
$kurt(\Delta z_t)$	2.96	2.90	34.0	30.9
$E[x_t]$	0.000	0.000	0.000	0.000
$\sigma(x_t)$	0.007	0.007	0.004	0.003
$skew(x_t)$	0.000	0.000	0.003	0.003
$kurt(x_t)$	2.95	2.83	35.3	32.8
$E[log(\sigma_t^2)]$	-15.5	-15.5	-15.5	-15.5
$\sigma(log(\sigma_t^2))$	0.000	0.000	3.10	3.10
$skew(log(\sigma_t^2))$	0.000	0.001	0.039	0.000
$kurt(log(\sigma_t^2))$	2.71	2.66	2.68	2.66

2.2.6 Options Pricing

A convenient feature of consumption-based asset pricing models is that they permit us to price a wide variety of assets using the same methodology. Options can be priced using the same Euler equation as stocks and bonds. The price of each asset is simply given by

$$q_t = \beta E\big[M_{t+1}R_{t+1}\big]$$

where M_{t+1} is the asset pricing kernel from Chapter 1

$$M_{t+1} = \left(\frac{C_{t+1}}{C_t}\right)^{-1}\left(\frac{\tilde{C}_{t+1}}{\tilde{C}_t}\right)^{\gamma}\left(\frac{U_{t+1}}{\mu_t}\right)^{1-\gamma-\delta},$$

and R_{t+1} is the asset's return. As in Chapter 1, this method can be used to price an equity asset who's return is the return to capital by using the equation

$$p_t = \beta E\big[M_{t+1}(d_{t+1} + p_{t+1})\big]$$

where p_{t+1} is the price of the equity asset tomorrow and d_{t+1} is tomorrow's dividend payment. Dividends are given by

$$d_t = Y_t - w_t n_t - I_t.$$

This is what is left over from the sales of output after paying for labor and investing. As usual the wages are equal to the marginal product of labor

$$w_t = (1 - \alpha) K_t^{\alpha} (Z_t n_t)^{-\alpha}.$$

This follows from the firm's profit maximization condition.

Next I price several European call options, each with a time to expiration of one period. Call options, when exercised, entitle their owner to purchase an underlying financial asset at a predetermined strike price. European call options may only be exercised at the time of expiration, as opposed to American options that may be exercised at any point up until expiration. The current market price of the underlying financial asset is referred to as the spot price. The five European call options that I price range in strike price from ten percent below the spot price to ten percent above the spot price. Options with a strike price below the spot price are known as in-the-money (ITM). Options with a strike price equal to the spot price are said to be at-the-money (ATM). And those with a strike price above the spot price are known as out-of-the-money (OTM). The underlying financial asset for each of the options is the equity asset defined above. The returns R_{t+1} in the Euler equations for each of the five European call options can be expressed as

$$max(p_{t+1} - 0.90p_t, 0)$$

$$max(p_{t+1} - 0.95p_t, 0)$$

$$max(p_{t+1} - p_t, 0)$$

$$max(p_{t+1} - 1.05p_t, 0)$$

$$max(p_{t+1} - 1.10p_t, 0).$$

Here, the max function captures the decision to exercise the option or not. If next period the market price of the equity asset is higher than the strike price, the option will be exercised and its owner will earn the difference between the two prices. If, however, next period the market price of the equity asset is lower than the strike price then the option will not be exercised an it's owner earns nothing.

In the next section of the chapter, when working with the Black-Scholes equation, I require a measure of the interest rate in the economy. I will use the one period net risk free rate of return given by

$$r_t^f = \frac{1}{q_t} - 1.$$

By using the equation

$$((1 + r_t^f)^4) - 1$$

I obtain an annualized measure of the interest rate.

2.3 Estimation

Smith (1993) and Gourieroux, Monfort, and Renault (1993) first developed the Indirect Inference estimator. At its core this estimation procedure is similar in strategy to the more common Method of Moments estimation procedure. Method of Moments requires one to choose a set of moments in the data to target. Indirect Inference, instead, requires one to specify an auxiliary model. The parameters of this auxiliary model, when estimated using data, make up the set of moments that are targeted. To avoid any confusion with the auxiliary model, I will refer to the model presented in Section 2 as the economic model for the remainder of this chapter. Indirect Inference places very few restrictions on what can serve as the auxiliary model for estimation. The auxiliary model need not be linear, or Gaussian, or have any analytical solution. The only requirement is that one can solve for the auxiliary model's parameters using actual data or using data obtained by simulating the economic model. The auxiliary model parameters are first estimated using the real data. Then the auxiliary model parameters are estimated many more times using simulated data, where each time the simulated data has been obtained by simulating the economic model given a different set of economic model parameters. Indirect Inference selects the set of economic model parameters such that the auxiliary model parameters estimated using data obtained by simulating the economic model, given this set of economic model parameters, are as close as possible to the auxiliary model parameters estimated using the real data.

2.3.1 Auxiliary Model

For my auxiliary model in this chapter I use a stochastic process similar to the one used to specify consumption and dividend growth in Bansal and Yaron (2004). My auxiliary model is a stochastic process that describes the joint behavior of the growth rate of consumption, g_t^c, and equity returns, r_t^e. The auxiliary model is described by the four equation system

$$g_{t+1}^c = \mu_c + x_t + \sigma_t \epsilon_{t+1} \tag{2.5}$$

$$r_{t+1}^e = \mu_e + \phi_e x_t + \varphi_e \sigma_t \eta_{t+1} \tag{2.6}$$

$$x_{t+1} = \rho_x x_t + \varphi_x \sigma_t \nu_{t+1} \tag{2.7}$$

$$log(\sigma_{t+1}^2) = (1 - \rho_\sigma)\mu_\sigma + \rho_\sigma log(\sigma_t^2) + \sigma_\omega \omega_{t+1}, \tag{2.8}$$

where x_t is a persistent shock to average consumption growth and equity returns and σ_t is a shock to the volatility at time t. ϵ_{t+1}, η_{t+1}, ν_{t+1}, and ω_{t+1} are all assumed to be independent and identically distributed $N(0,1)$. The four equations contain a total of nine auxiliary model parameters. Let $\boldsymbol{\beta} = (\mu_c, \mu_e, \mu_\sigma, \phi_e, \varphi_e, \rho_x, \varphi_x, \rho_\sigma, \sigma_\omega)$. These nine auxiliary model parameters, estimated using actual data, are the targets of my Indirect Inference estimation routine. My estimation procedure does not target the average risk-free rate, or the average equity premium, or the variance of either of these. Instead I target the long-run risk and stochastic volatility in the data on consumption growth and equity returns.

Estimating the Auxiliary Model

Estimation of the auxiliary model is not a straightforward exercise since equations 5 through 7 are non-Gaussian, and all of the equations include variables that are unobservable in the data. I estimate the auxiliary model's parameters in two steps.

Step 1 In step one I ignore the stochastic volatility, and assume that the level of volatility is fixed. This yields the following three equations:

$$g^c_{t+1} = \mu_c + x_t + \sigma\epsilon_{t+1}$$

$$r^e_{t+1} = \mu_e + \phi_e x_t + \varphi_e \sigma \eta_{t+1}$$

$$x_{t+1} = \rho_x x_t + \varphi_x \sigma \nu_{t+1}.$$

These three equations describe a linear state space model. I follow Harvey (1989) to write the equations in the state space form (SSF) and use maximum likelihood (ML) to estimate the parameters. I use the Kalman Filter to construct the prediction error decomposition form of the likelihood. I use the FFSQP Fortran algorithm to numerically maximize this likelihood function, yielding ML estimates for six of the nine parameters in my full auxiliary model. I also follow Harvey (1989) and construct the information matrix, providing me with an estimate of the covariance matrix of the ML estimates. I use this to compute the standard errors for my estimates.

Step 2 Next, I use a Kalman smoother to obtain an estimated time series of the unobserved state x_t. Using this time series for x_t and the estimated parameter values, I compute a time series of the residuals from the two measurement equations in Step

1. Let $c_{t+1} = \sigma_t \epsilon_{t+1}$ be the residuals from the consumption growth equation and $e_{t+1} = \varphi_e \sigma_t \eta_{t+1}$ be the residuals from the equity returns equation. Taking the natural log of the square of these residuals yields

$$log(c_{t+1}^2) = log(\sigma_t^2) + log(\epsilon_{t+1}^2)$$

$$log(e_{t+1}^2) = log(\varphi_e^2) + log(\sigma_t^2) + log(\eta_{t+1}^2).$$

This allows me to define another linear SSF

$$log(c_{t+1}^2) = E[log(\epsilon_{t+1}^2)] + h_t + \xi_{t+1}^c$$

$$log(e_{t+1}^2) = E[log(\eta_{t+1}^2)] + 2log(\varphi_e) + h_t + \xi_{t+1}^e$$

$$h_{t+1} = (1 - \rho_\sigma)\mu_\sigma + \rho_\sigma h_t + \sigma_\omega \omega_{t+1};$$

where

$$h_t = log(\sigma_t^2),$$

$$\xi_{t+1}^c = log(\epsilon_{t+1}^2) - E[log(\epsilon_{t+1}^2)],$$

and

$$\xi_{t+1}^e = log(\eta_{t+1}^2) - E[log(\eta_{t+1}^2)].$$

ξ_{t+1}^c and ξ_{t+1}^e are non-Gaussian, zero mean, white noises. Since ϵ_{t+1} and η_{t+1} are assumed to be $NID(0,1)$, the mean and variance of $log(\epsilon_{t+1}^2)$, as well as the mean and variance of $log(\eta_{t+1}^2)$, are known to be approximately -1.27 and $\pi^2/2$ respectively.[3]

Following Ruiz (1994) I treat ξ_t and ξ_t^e as though they are $NID(0, \pi^2/2)$ and

[3] see Abramovitz and Stegun (1970)

estimate the parameter values by Quasi-Maximum Likelihood (QML). I once again use the Kalman Filter to estimate the unobserved stochastic volatility process, and I use prediction error decomposition to obtain the Quasi-Likelihood function. Numerically maximizing this function yields the QML estimates of the auxiliary model parameters in the stochastic volatility equation. Also following an appendix to Ruiz (1994) I compute the variance-covariance matrix for the QML estimates, and use this to compute standard errors.

A Recovery Exercise

I perform a recovery exercise in order to test the method I use to estimate the auxiliary model. I select a set of parameter values for $\boldsymbol{\beta}$. I treat these as the true values in the population. Using these values, I simulate the continuous processes using draws from a random standard normal distribution for each of the shocks. I simulate very long time series for consumption growth and equity returns: 10,000 periods long. Next, I treat these long time series as my data, and I use this data to estimate the parameters of the auxiliary model following the method laid out in the previous subsection. Given this large amount of data the estimation technique should yield estimates of the parameter values that are close to the "true" values I selected at the start.

For each of the parameters in $\boldsymbol{\beta}$, Table 2 shows the true value and the estimated value. We can see that my estimation technique performs well in this exercise. Many of the parameter values are exactly recovered. There does appear to be some downward bias in the estimation of ρ_x, the persistence of the long-run component. However, it is important to note that the magnitude of this bias does not appear large enough to change any of the results.

Table 2.2: Auxiliary model parameters

	True	Estimated
μ_c	0.02	0.02
μ_e	0.03	0.03
μ_σ	-9.734	-9.732
ϕ_e	0.50	0.79
φ_e	0.75	0.75
ρ_x	0.95	0.82
φ_x	0.001	0.001
ρ_σ	0.45	0.45
σ_w	0.207	0.214

2.3.2 Indirect Inference Estimator

I estimate the key parameters of my model using Indirect Inference to target specific consumption and equity return dynamics. I use this method to estimate nine of the economic model's parameters. Two of the estimated parameters describe agent preferences; γ, which controls IES, and δ, which controls risk aversion. The third estimated parameter, ϕ, controls the magnitude of the capital adjustment costs. The next two estimated parameters, ρ_x and φ_x, control the persistence and the relative volatility of long-run productivity risk. Two more estimated parameters control the persistence, ρ_σ, and the standard deviation, σ_ν, of stochastic volatility in aggregate productivity. The final two estimated parameters, ρ_π and σ_υ, determine the volatility and the persistence of the shocks to the probability of a disaster event. Let $\boldsymbol{\xi}$ be a vector containing these nine parameters.

First, I estimate the auxiliary model's parameters using actual data. I call the vector of these estimates $\hat{\boldsymbol{\beta}}$. Next, I estimate the auxiliary model's parameters many more times using simulated data. Each time I solve the economic model given a different set of values for $\boldsymbol{\xi}$, and use this solution to obtain simulated data for consumption growth and equity returns. I use the simulated data from this particular solution to estimate the auxiliary model's parameters, and I call these estimates $\tilde{\boldsymbol{\beta}}(\boldsymbol{\xi})$.

Since I use post-war U.S. data to create my estimation targets $\hat{\boldsymbol{\beta}}$ and this data does not contain any episodes similar to the disasters in my economic model; I do not allow the simulated data to contain any realizations of a disaster event. This disaster risk is actually a Peso Problem, in the style of Danthine and Donaldson (1999). The agents in the model believe that there is the risk of a disaster that never actually occurs.

My Indirect Inference estimate of $\boldsymbol{\xi}$ minimizes the objective function

$$\left(\frac{\tilde{\boldsymbol{\beta}}(\boldsymbol{\xi}) - \hat{\boldsymbol{\beta}}}{\hat{\boldsymbol{\beta}}}\right)' \left(\frac{\tilde{\boldsymbol{\beta}}(\boldsymbol{\xi}) - \hat{\boldsymbol{\beta}}}{\hat{\boldsymbol{\beta}}}\right).$$

I use the DiRect (short for Dividing Rectangles) optimization algorithm to solve this optimization problem. This routine has a couple of features that make it ideal for this application. First, it requires the parameters space, over which it searches, to be bounded. Most of the parameters in $\boldsymbol{\xi}$ are bounded, so using an approach that naturally integrates this information about the parameters space into the search routine is useful. Second, the DiRect algorithm is globally convergent to the global optimum. The objective function is likely to contain many local minima, and so using a method that converges to the global minimum is absolutely critical.

2.3.3 Implied Volatility

In this subsection, I describe the method used to solve numerically for the Black-Scholes implied volatility for each of the call options priced using the model described in the previous section. By generating a long sequence of random numbers from the standard normal distribution and using the estimated policy and pricing functions I simulate the economy in the model. I obtain time series for each of the choice variables as well as for each of the asset prices. These option prices are treated as the true market prices for each of the call options in each period. The volatility implied by each of these options prices is the value for σ, the volatility of the underlying financial asset, such that the Black-Scholes options pricing equation yields a theoretical value for the option that is equal to the true market price.

Black and Scholes published their seminal paper on the theoretical pricing of options in 1973. The key intuition that leads to their famous pricing equation is the idea that one should not be able to create a riskless position out of long and short positions in the option and its underlying asset. From this no arbitrage condition and a few other assumptions Black and Scholes derived a partial differential equation, the solution of which yielded the theoretical value of the option. This value C^o is given by

$$C_t^o = \Phi(d_1)p_t - \Phi(d_2)Xe^{-r_t^f T}$$

where Φ is the standard normal cumulative density function, X is the strike price, T is the time to expiration,

$$d_1 = \frac{1}{\sigma\sqrt{T}}\Big[ln\big(\frac{p_t}{X}\big) + \big(r_t^f + \frac{\sigma^2}{2}\big)T\Big],$$

p_t is the spot price, r_t^f is the annualized interest rate, and

$$d_2 = d_1 - \sigma\sqrt{T}.$$

This equation does not depend of the expected return of the underlying financial asset. The option's value is, all else equal, increasing in the volatility of the underlying asset. The theoretical value of an option, according to the Black-Scholes equation also depends on the time until expiration, the strike price, the risk-free interest rate, and the current market price of the underlying financial asset. Since the Black-Scholes equation is continuous and increasing in the volatility of the underlying financial asset the Intermediate Value Theorem guarantees the existence of a unique σ^* such that the theoretical option value is equal to its market price. This σ^* is the Black-Scholes implied volatility for a particular strike price. Plotting this for various strike prices yields an implied volatility curve. The Black-Scholes options pricing model yields the closed form equation for the theoretical value of an option given above. However, there isn't a closed form equation for its inverse with respect to the volatility of the underlying financial asset. The value for the implied volatility must be solved for numerically. I use the Newton Raphson algorithm to solve for the root

$$C_t^o(\sigma) - o_t^i = 0$$

where o^i is the market value of option i.

Brenner and Subrahmanyan (1988) propsed a simple closed-form solution for the Black-Scholes implied volatility. I use this closed form approximation to select the starting value for the implied volatility in my iterative routine. I let

$$\sigma_0 = \sqrt{\frac{2\pi}{T}} \frac{C_t^o}{p_t}$$

where C^o is the theoretical value of the option obtained using the Black-Scholes option pricing equation with $\sigma = 1$, X is the strike price, and T is the time to expiration expressed in years. I then iterate using the Newton-Raphson equation

$$\sigma_{n+1} = \sigma_n - \frac{C_t^o(\sigma_n) - o_t^i}{\nu(\sigma_n)}$$

where ν is the derivative of the Black-Scholes option pricing equation with respect to the volatility of the underlying financial asset. This derivative is known as the Black-Scholes Vega and is given by

$$\nu(\sigma) = p_t \phi(d_1)\sqrt{T}$$

where $\phi(\cdot)$ is the standard normal probability density function and d_1 is defined as above. I continue to iterate using the above equation until the new value for the implied volatility is within 1x10^{-4} of the previous iteration's value.

I use the same method to solve for the volatility implied by the market price for each of the five options I price. Finally, I take the average of the implied volatility for each of the options over all periods and plot this against the strike price for each of the options. This yields the implied volatility curve for the model presented in the previous section. The next section will present the main results from the simulation of the model.

2.4 Results

2.4.1 Options Pricing Functions

I use the model presented in an earlier section of this chapter to price five European call options. Each option gives its owner the right to purchase the underlying financial asset for a fixed strike price three months from the time of purchase. If the option is ATM or OTM then it is a bet that the price of the underlying financial asset will go up. If three months from the time of purchase the market price for the underlying financial asset is higher than the option's strike price, then the owner will exercise his option to buy at the strike price and sell at the market price, earning a positive return. If the market price is not higher than the strike price at expiry then the owner can simply choose not to exercise the option, earning zero return. If the option is ITM then as long as the market price does not fall below the strike price the owner will choose to exercise the option at expiry. Two of the options I price are ITM; one with a strike price ten percent below the current market price and one with a strike price five percent below the current market price. One is an ATM option. The other two are OTM; one with a strike price five percent above the current market price and one with a strike price ten percent above the current market price. For each of the five options priced the underlying financial asset is the return to capital. I will refer to this underlying financial asset as the equity asset. For each option the time to expiration will be three months. Since the model was estimated to account for quarterly data, this represents one period in the model.

The solution to the model yields pricing functions for each of the five call options. These give the price of each of the call options as a function of the state variables. For a given combination of the state variables the price of a call option is decreasing in its strike price. This result is very intuitive since the probability that the market price of the equity asset three months from now is higher than the strike price increases as the strike price decreases. Taken to one extreme if the strike price were zero than the price of the equity asset next period would almost surely be higher than the strike and we would almost be guaranteed a positive return. Agents should be willing to pay a high price for such an option.

Examining the option pricing functions across values for today's capital stock, holding all other state variables constant; a distinct pattern emerges across the various call options. This pattern highlights two opposing effects of different levels for today's capital stock on the value of an option in the model economy. The first effect follows from the fact that if the level of today's capital stock is high, then the current price of the equity asset is also high. The pricing function for the equity asset is very nearly linear in the current capital stock. If the price of the option is increasing in the price of the underlying financial asset, then this would cause the option value to be higher when today's capital stock is higher.

The second effect follows from the fact that if the capital stock is currently above the steady state level it will tend to shrink and if the current capital stock is above the steady state level it will tend to grow. This follows from the concavity of the decision rule for tomorrow's level of capital in today's level of capital. Since, as previously stated, the price of the equity asset is very nearly linear in the current level of capital the price of the equity asset will tend to rise when the current level of capital is below the steady state and tend to shrink when the current level of capital is above the steady state. This would cause the call option's value to be lower when today's capital stock is higher, and vice versa.

Figure 2.1: Option price by current level of capital, ITM

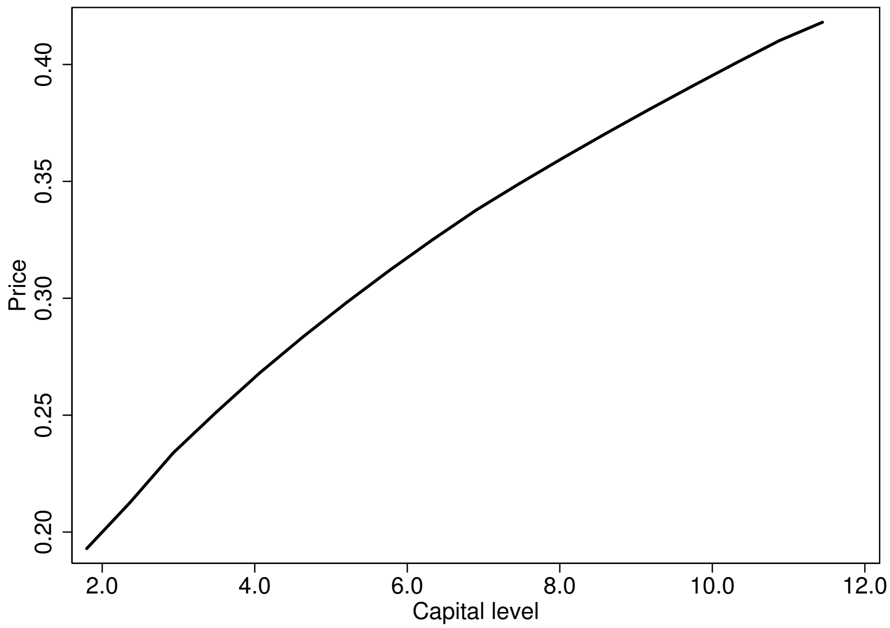

Figures 1 through 3 show that which of the two effects will dominate depends on the strike price. Changes in the strike price represent a sort of extensive margin; increasing or decreasing the number of future states in which the option is likely to be exercised. For the ATM and OTM options, the second effect dominates and the option price is decreasing in today's level of capital. For the ITM option, where the option is extremely likely to be exercised, the first effect dominates and the option price is increasing in today's level of capital.

2.4.2 Black-Scholes Implied Volatility Surface

As discussed in the preceding section of this chapter, there exists a unique value σ, the volatility of the underlying financial asset, such that the theoretical value of a call option according to the Black-Scholes options pricing equation is equal to

Figure 2.2: Option price by current level of capital, ATM

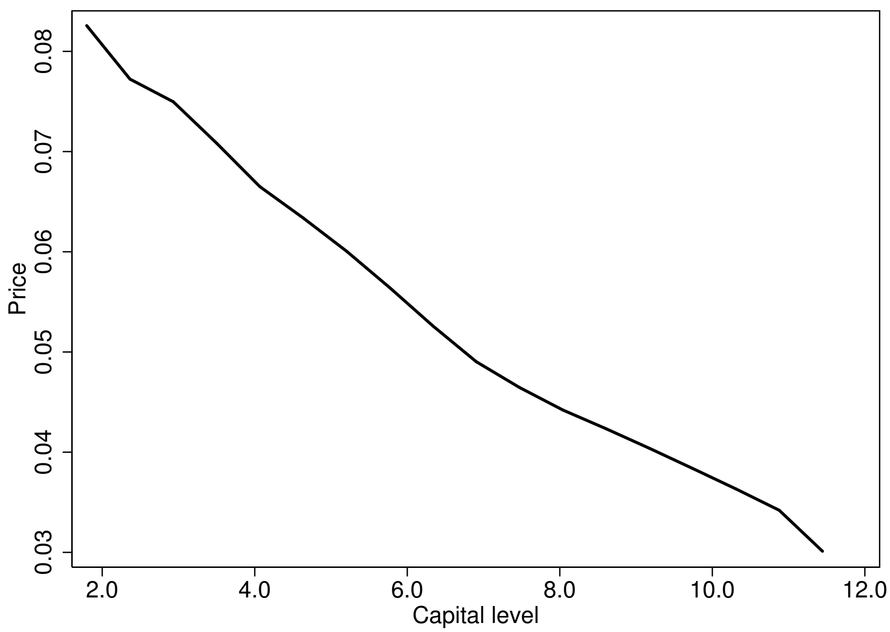

Figure 2.3: Option price by current level of capital, OTM

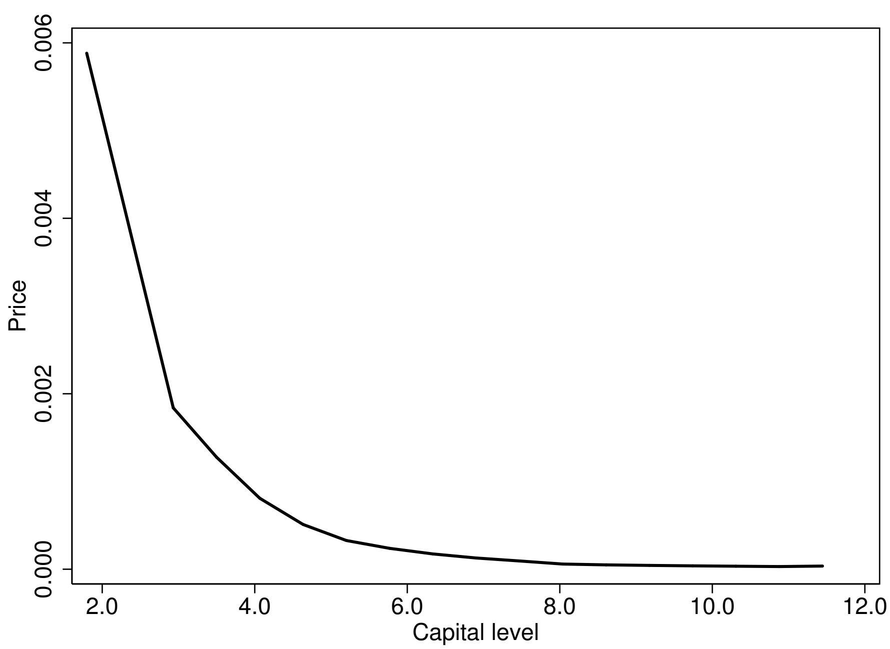

the market price for that option. This σ^* is known as the Black-Scholes implied volatility. This implied volatility can be solved for numerically for options with the same underlying financial asset but different strike prices and expiries. Plotting these implied volatilities against the different strike prices and expiries yields a Black-Scholes implied volatility surface.

In this paper, I focus on options that not only share the same underlying financial asset but also the same time to expiration. I plot a Black-Scholes implied volatility curve as a function of strike price. One assumption of the Black-Scholes options pricing environment is that the underlying asset follows a Brownian motion with a constant dispersion parameter σ. The Black-Scholes implied volatility for a given strike price is the value of this constant dispersion parameter σ such that the Black-Scholes theoretical option value is equal to the market price for the option. Therefore, if the market prices for each of the call options were generated using the Black-Scholes options pricing model, then the implied volatility computed using each of the call options would be σ regardless of the strike price. So the Black-Scholes options pricing model implies a flat implied volatility curve.

However, computing the implied volatilities using actual data does not yield a flat implied volatility curve. Bates (2000) provides empirical evidence that since the U.S. stock market crash of 1987, the implied volatility curve for S&P 500 options has been downward sloping. Since the underlying financial asset for all of the options being priced in this paper is a broad market return, we should expect to observe a volatility skew if the model is consistent with the empirical evidence. Figure 4 shows an implied volatility curve computed using data simulated from the model. The figure shows that the model does, in fact, generate a volatility skew. The implied volatility is monotonically decreasing in strike price.

Figure 2.4: Implied Volatility Curve

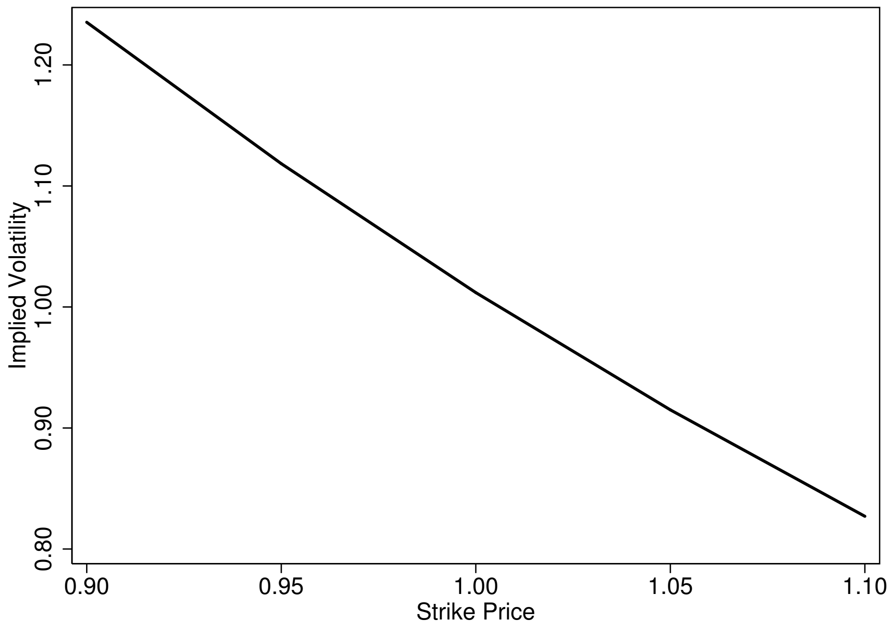

Next, I propose a method for measuring the magnitude of this volatility skew. I simulate the economy for 275 periods and I compute

$$\frac{\sigma^*(ITM) - \sigma^*(OTM)}{\sigma^*(ATM)}$$

for each period. This measure is the absolute difference between the implied volatility computed using the ten percent ITM call option and the implied volatility computed using the ten percent OTM call option, scaled by the implied volatility computed using the ATM call option. This measure is similar in style to the one proposed in Mixon (2010), which that paper shows to be empirically useful. I then take the average over all of the periods. I perform this simulation exercise 1,000 times and compute the mean of this average over all of the simulations. If the Black-Scholes option pricing model were correct then this mean should be zero. This measure of the magnitude of the volatility skew, computed from the baseline version of the model is found to be 0.3490. The baseline version of the model is therefore able to account for this feature of the options price data

2.4.3 Shutting Down Features of the Model

The results presented in the previous two subsections form the baseline set of results that I measure against in this subsection of the paper. I now analyze how the various model features contribute to these baseline results. I explore shutting down the long-run productivity risk or the disaster risk individually. Then I explore shutting them down simultaneously. Finally, I explore eliminating the non-expected utility preferences.

Figures 1 through 3 plot the pricing functions for the five percent ITM, ATM, and five percent OTM call options as a function of the level of today's capital. The pattern of the relationship between the capital stock and the option value across the different strike prices was described in detail in the last subsection of the paper. The pattern of the direction of the relationship between the option price and the level of capital across the different strike prices survives shutting down each of the key model features. However, the pattern of slopes across the different strike prices changes notably as features of the full model are shut off.

Figures 5 through 7 show the relationship between price and today's level of capital for the same three options as before, now estimated from a version of the model without long-run risk or stochastic volatility in aggregate productivity growth. The long-run risk and stochastic volatility fatten the tails of the distribution of future aggregate productivity. This does not have much effect on the five percent ITM

Figure 2.5: Option price by current level of capital, ITM

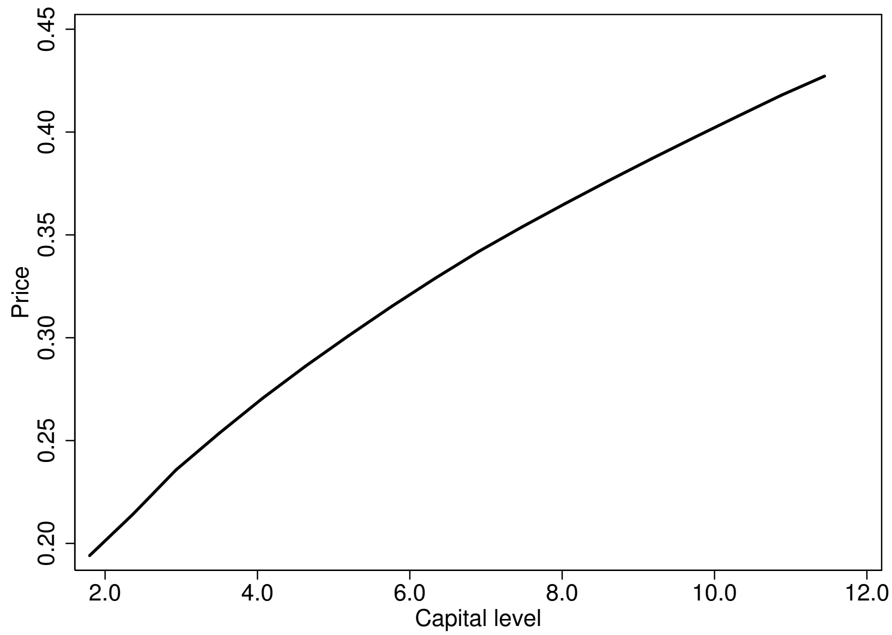

option, as we can see from Figure 5. For a given level of capital today, the fatter tails in the distribution of tomorrow's productivity increase the probability of an extremely positive draw from the productivity distribution that would raise tomorrow's capital stock, and therefore increase tomorrow's price of the underlying asset. This is clearly shown in Figure 6. The pricing function is swung upward by the inclusion of long-run productivity risk. We can also see this effect in Figure 7. The probability that the capital stock, and therefore the price of the underlying asset, rises by more than five percent is negligible for all but the lowest levels of today's capital. But the inclusion of long-run productivity risk increases the range of values for today's capital stock for which the five percent OTM option takes on a positive price.

However, the effect is not the same when we shut down the disaster risk in the model. Figures 8 through 10 show the relationship between price and today's level of

Figure 2.6: Option price by current level of capital, ATM

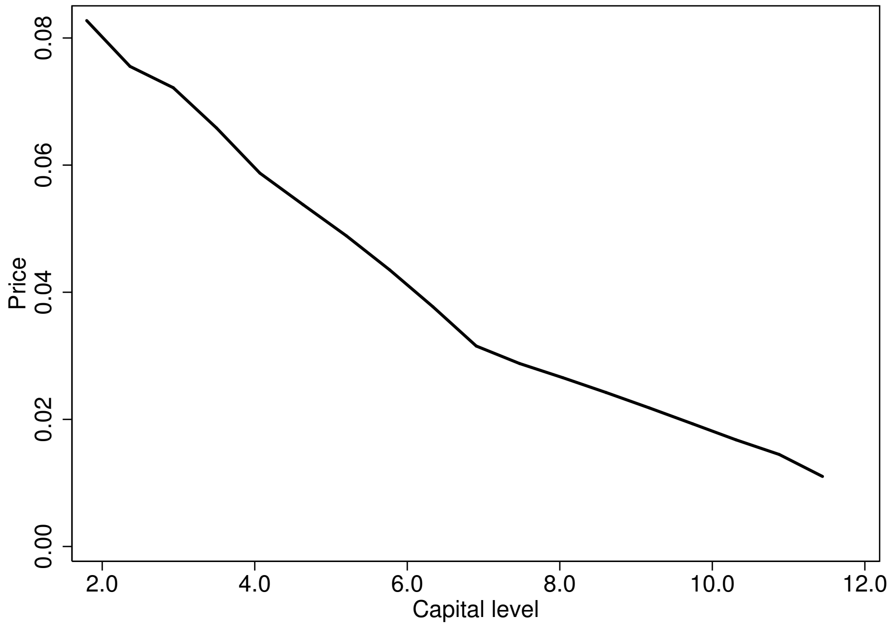

Figure 2.7: Option price by current level of capital, OTM

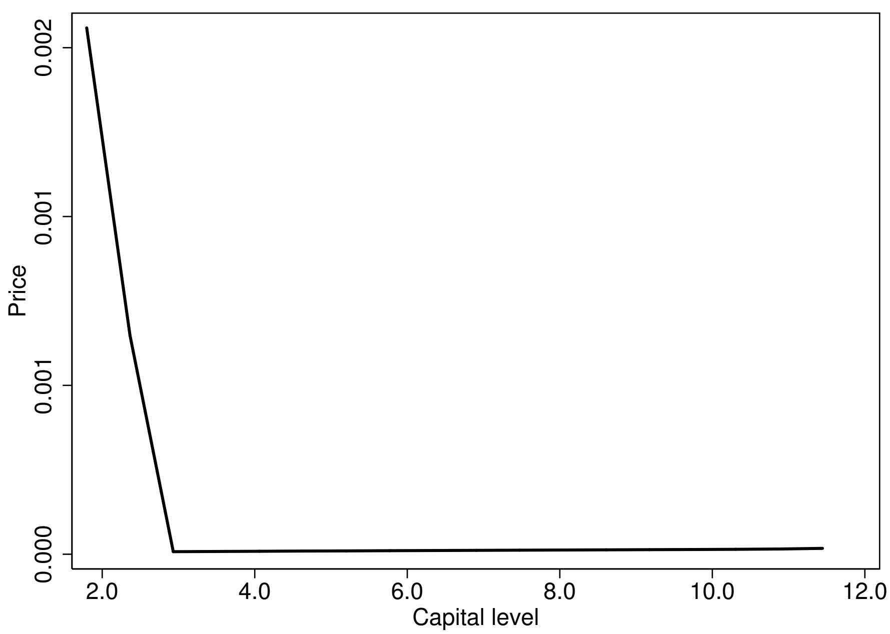

Figure 2.8: Option price by current level of capital, ITM

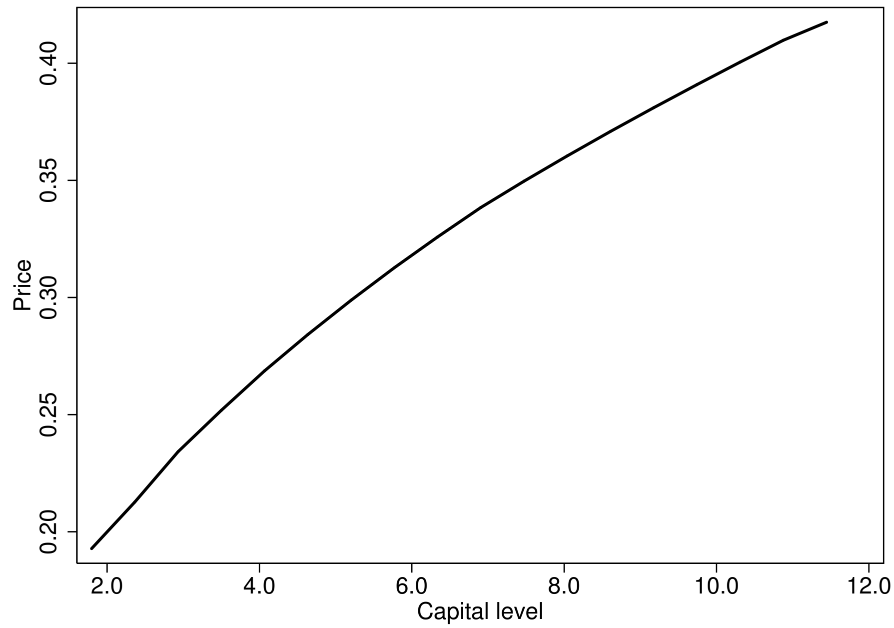

capital for the same three options as before is largely unchanged from the relationship estimated from the full version of the model. This paper explores only the prices of European call options. It may be the case that the patterns observed across the pricing functions for put options would be more affected by this disaster risk. Disasters present only downside risk without the possibility of upside risk. They would increase the number of future states in which an ATM or OTM put option might be exercised, thus affecting the tradeoffs described in the previous subsection.

Next I examine the effect of shutting down the key model features on my measure of the level of implied volatility skew. Recall from the previous subsection that this measure of the implied volatility skew for the baseline model was the absolute difference between the implied volatility computed using the ten percent ITM call option and the implied volatility computed using the ten percent OTM call option,

Figure 2.9: Option price by current level of capital, ATM

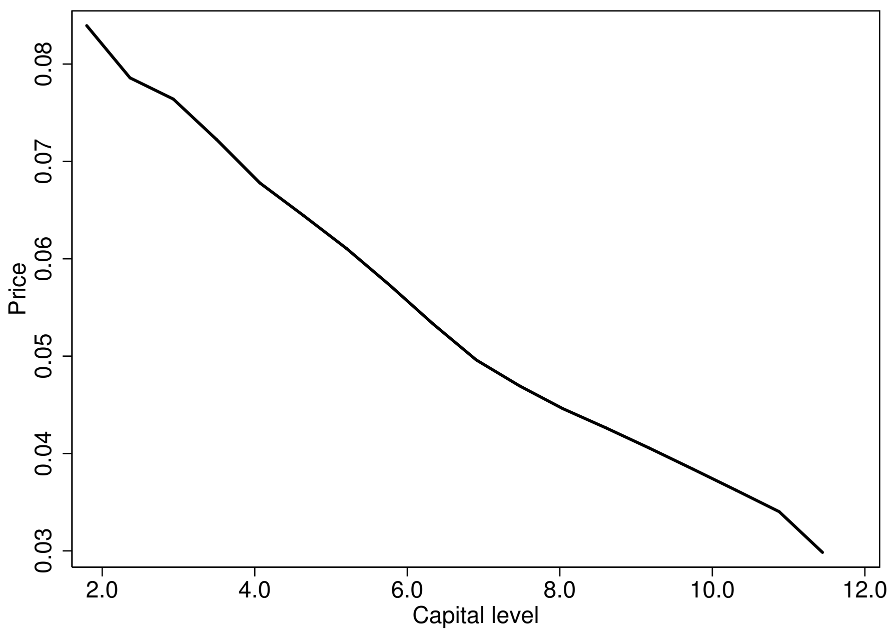

Figure 2.10: Option price by current level of capital, OTM

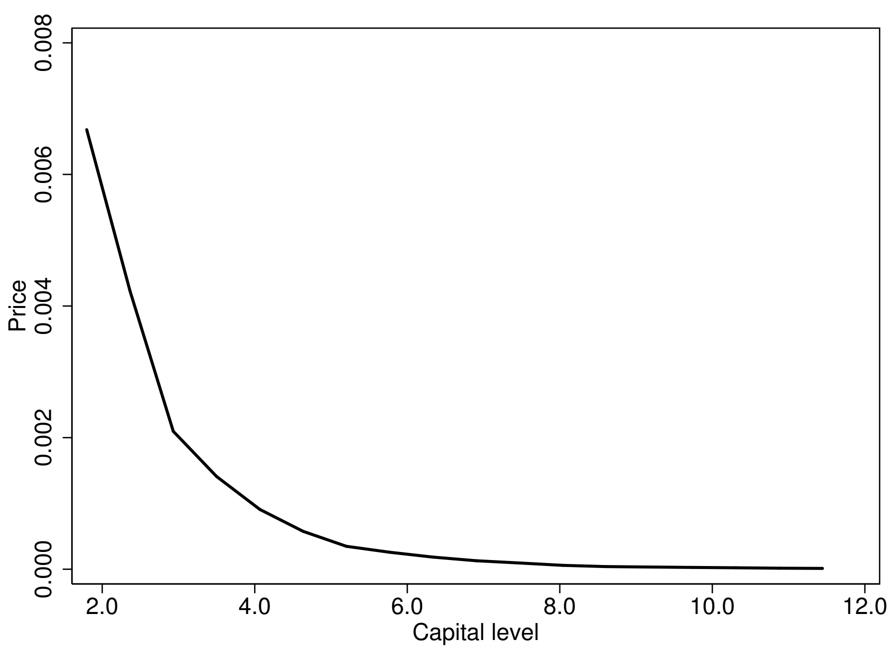

scaled by the implied volatility computed using the ATM call option. Values for this measure that are closer to zero represent a flatter implied volatility curve.

First, I compute this measure for the model without long-run risk or stochastic volatility in aggregate productivity growth ($\sigma_w = 0$ and $\varphi_x = 0$), using exactly the same method as descibed previously. The average level of implied volatility skew decreases to 0.3489. Next, I compute the measure using data from the model without disaster risk ($\mu_\pi = 0$). The average level of implied volatility skew decreases to 0.3488. Next, I compute the measure using data from the model without either long-run productivity risk or disaster risk. The average level of implied volatility skew decreases farther to 0.3430. Finally, I additionally shut down the capital adjustment costs and the non-expected utility preferences ($\phi = 0$ and $\rho = 1 - \sigma$). This version of the model roughly corresponds to the standard real business cycles framework. The average level of implied volatility skew is again decreased to 0.3415.

This evidence suggests that the standard real business cycles framework does produce an implied volatility smirk, consistent with the pattern in the data. The addition of either long-run productivity risk or disaster risk make the average volatility smirk steeper, and the addition of both model features make the average volatility smirk steeper still. However the effect of adding a second of these model features is smaller than the effect of adding the first.

2.5 Conclusion

In this chapter I explore the pricing of several European-style call options. For each of the options the underlying financial asset is the return to capital in the model. The model economy is the one first set out in Chapter 1, which contains both long-run productivity risk as well as rare disaster risk. The behavior of options prices can give us a glimpse of how agents in the economy view the distribution of future states of the world.

I observe a distinct distinct patterns for the price of each option as a function of today's capital stock, depending on if the option in ITM, ATM, or OTM. The ITM option is increasing in today's level of capital, the ATM option is slightly decreasing in capital stock, and the OTM is greatly decreasing in capital stock. Although these patterns are shown to exist in the standard business cycle framework version of the model, the inclusion of both long-run risk and stochastic volatility in aggregate productivity growth as well as rare disaster risk make the patterns more pronounced.

I also observe that the baseline version of the model generates an implied volatility skew. This is consistent with the data on S&P 500 call options of similar time to expiration. I propose a measure of the amount of skew in the model simulated data. I show that the inclusion of stochastic volatility and long-run risk in aggregate productivity growth increase the average steepness of the volatility skew, whereas the inclusion of rare disaster risk decreases the average volatility skew.

Obtaining more data on S&P 500 options, both calls and puts, would open up avenues to further this research. With this data the moments in the options data could serve as targets for the estimation of the economic model's parameters. Indeed, one could estimate the model to precisely match the level of the volatility skew observed in the data. This is left for future work.

References

Bansal, Ravi, Dana Kiku, and Amir Yaron. (2012) *An empirical evaluation of the long-run risks model for asset prices.* Critical Finance Review, 1: 183-221.

Bansal, Ravi, and Amir Yaron. (2004) *Risks for the long run: a potential resolution of asset pricing puzzles.* The Journal of Finance, Vol. 59, No. 4, pp. 1481-1509.

Barro, Robert, J. (2006) *Rare Disasters and Asset Markets in the Twentieth Century.* The Quarterly Journal of Economics, Vol. 121, No. 3, pp. 823–866.

Bates, David. (1996) *Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options.* Review of Financial Studies, Vol. 9, pp. 69-107.

Bates, David. (2000) *Post-'87 Crash Fears in S&P 500 Futures Options.* Journal of Econometrics, Vol. 94, pp. 181-238.

Black, Fischer, and Scholes, Myron S. (1973) *The Pricing of Options and Corporate Liabilities.* Journal of Political Economy, Vol. 81, No. 3, pp. 637-654.

Bollerslev, T., and H. Zhou. (2002) *Estimating stochastic volatility diffusions using conditional moments of integrated volatility.* Journal of Econometrics, 109: 33-65.

Brenner, M. and Subrahmanyam, M.G. (1988) *A simple formula to compute the implied standard deviation.* Financial Analysts Journal, Vol. 44, pp. 80-83.

Cao, C. (1992) *Pricing Foreign Currency Options with Stochastic Volatility.* Working paper, University of Chicago.

Croce, Mariano M. (2014) *Long-run productivity risk: a new hope for production-based asset pricing?* Journal of Monetary Economics, 66: 13-31.

Danthine, J.P. and Donaldson, J. (1999) *Non-falsified expectations and general equilibrium asset pricing: the power of the peso.* The Economic Journal, Vol. 109, pp. 607-635.

Gabaix, Xavier. (2012) *Variable Rare Disasters: An Exactly Solved Framework for*

Ten Puzzles in Macro-Finance. The Quarterly Journal of Economics, Vol. 127, No. 2, pp. 645-700.

Gourieroux, C., A. Monfort, and E. Renault. (1993) *Indirect Inference.* Journal of Applied Econometrics, 8: S85-S118.

Gourio, Francois. (2012) *Disaster risk and business cycles.* American Economic Review, 102(6): 2734-2766.

Harvey, Andrew, C. (1989) *Forecasting, structural time series models and the Kalman filter.* Cambridge University Press

Khelifa, Ben Z. (1991) *Parametric and Nonparametric Tests of the Pure Diffusion Model Adjusted for the Early Exercise Premium Applied to Foreign Currency Options.* in Essays in International Finance, Wharton School dissertation, University of Pennsylvania, chap. 1.

Lucas, R.E., Jr. (1978) *Asset prices in an exchange economy.* Econometrica 46: 1429-1445.

Mehra, R. and E.C. Prescott. (1985) *The equity premium: a puzzle.* Journal of Monetary Economics, 15, pp 145-161.

Mixon, Scott. (2002) *Factors Explaining Movements in the Implied Volatility Surface.* Journal of Futures Markets, Vol. 22, No. 10, pp. 915-937.

Mixon, Scott. (2010) *What Does Implied Volatility Skew Measure?.* The Journal of Derivatives, Vol. 18.

Rietz, Thomas A. (1988) *The equity risk premium: a solution.* Journal of Monetary Economics, 22: 117-131.

Rubinstein, Mark. (1985) *Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Option Classes from August 23, 1976 through August 31, 1978.* The Journal of Finance, Vol. 40, pp. 455-480.

Ruiz, Esther. (1994) *Quasi-maximum likelihood estimation of stochastic volatility models.* Journal of Econometrics, 63: 289-306.

Schorfheide, Frank, Dongho Song, and Amir Yaron (2017) *Identifying long-run risks: a bayesian mixed-frequency approach.* , Econometrica (forthcoming)

Smith, A. A. (1993) *Estimating nonlinear time-series models using simulated vector autoregressions.* Journal of Applied Econometrics, 8: S83-S84

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

Appendices

Appendix A

Stationarity Transformation

This appendix contains the derivation of the equations for the stationary version of the planner's recursive problem. I begin with the value function

$$V(\boldsymbol{\kappa}) = \max_{\boldsymbol{\chi}} \left((1-\beta)\left(C_t^{\theta}(Z_{t-1}(1-n_t))^{1-\theta} \right)^{\gamma} + \beta \left(E[V(\boldsymbol{\kappa'})^{1-\delta}] \right)^{\frac{\gamma}{1-\delta}} \right)^{\frac{1}{\gamma}}$$

$$= \max_{\boldsymbol{\chi}} \left((1-\beta)\left((\frac{Z_t}{Z_t}C_t)^{\theta}((\frac{Z_t}{Z_t}Z_{t-1})(1-n_t))^{1-\theta} \right)^{\gamma} + \beta \left(E[V(\boldsymbol{\kappa'})^{1-\delta}] \right)^{\frac{\gamma}{1-\delta}} \right)^{\frac{1}{\gamma}}$$

Let $\hat{C}_t = C_t/Z_t$ and $g_t = Zt/Z_{t-1}$ then

$$V(\boldsymbol{\kappa}) = \max_{\boldsymbol{\chi}} \left((1-\beta)\left(Z_t^{\theta}\hat{C}_t^{\theta}Z_t^{1-\theta}(g_t^{-1}(1-n_t))^{1-\theta} \right)^{\gamma} + \beta \left(E[V(\boldsymbol{\kappa'})^{1-\delta}] \right)^{\frac{\gamma}{1-\delta}} \right)^{\frac{1}{\gamma}}$$

$$= \max_{\boldsymbol{\chi}} \left((1-\beta)\left(Z_t\hat{C}_t^{\theta}(g_t^{-1}(1-n_t))^{1-\theta} \right)^{\gamma} + \beta \left(E[V(\boldsymbol{\kappa'})^{1-\delta}] \right)^{\frac{\gamma}{1-\delta}} \right)^{\frac{1}{\gamma}}$$

Let $\hat{V}(\cdot) = V(\cdot)/Z_t$

$$Z_t\hat{V}(\hat{\boldsymbol{\kappa}}) = \max_{\boldsymbol{\chi}}\left((1-\beta)\left(Z_t\hat{C}_t^{\theta}(g_t^{-1}(1-n_t))^{1-\theta}\right)^{\gamma} + \beta\left(E[(Z_t\hat{V}(\hat{\boldsymbol{\kappa}}'))^{1-\delta}]\right)^{\frac{\gamma}{1-\delta}}\right)^{\frac{1}{\gamma}}$$

$$Z_t\hat{V}(\hat{\boldsymbol{\kappa}}) = \max_{\boldsymbol{\chi}}\left((1-\beta)Z_t^{\gamma}\left(\hat{C}_t^{\theta}(g_t^{-1}(1-n_t))^{1-\theta}\right)^{\gamma} + \beta Z_t^{\gamma}\left(E[\hat{V}(\hat{\boldsymbol{\kappa}}')^{1-\delta}]\right)^{\frac{\gamma}{1-\delta}}\right)^{\frac{1}{\gamma}}$$

$$Z_t\hat{V}(\hat{\boldsymbol{\kappa}}) = \max_{\boldsymbol{\chi}} Z_t\left((1-\beta)\left(\hat{C}_t^{\theta}(g_t^{-1}(1-n_t))^{1-\theta}\right)^{\gamma} + \beta\left(E[\hat{V}(\hat{\boldsymbol{\kappa}}')^{1-\delta}]\right)^{\frac{\gamma}{1-\delta}}\right)^{\frac{1}{\gamma}}$$

Which leads to the final equation used in the paper

$$\hat{V}(\hat{\boldsymbol{\kappa}}) = \max_{\hat{\boldsymbol{\chi}}}\left((1-\beta)\left(\hat{C}_t^{\theta}(g_t^{-1}(1-n_t))^{1-\theta}\right)^{\gamma} + \beta\left(E[\hat{V}(\hat{\boldsymbol{\kappa}}')^{1-\delta}]\right)^{\frac{\gamma}{1-\delta}}\right)^{\frac{1}{\gamma}}.$$

Next I transform the resource constraint

$$C_t + I_t \leq K_t^{\alpha}(Z_t n_t)^{1-\alpha}$$

$$(\frac{Z_t}{Z_t}C_t) + (\frac{Z_t}{Z_t}I_t) \leq (\frac{Z_{t-1}}{Z_{t-1}}K_t)^{\alpha}((\frac{Z_{t-1}}{Z_{t-1}}Z_t)n_t)^{1-\alpha}$$

Let $\hat{I}_t = I_t/Z_t$ and $\hat{K} = K_t/Z_{t-1}$ then

$$Z_t\hat{C}_t + Z_t\hat{I}_t \leq (Z_{t-1}\hat{K}_t)^{\alpha}(Z_{t-1}g_t n_t)^{1-\alpha}$$

$$Z_t(\hat{C}_t + \hat{I}_t) \leq Z_{t-1}\hat{K}_t^{\alpha}(g_t n_t)^{1-\alpha}$$

$$\hat{C}_t + \hat{I}_t \leq g_t^{-\alpha}\hat{K}_t^{\alpha}n_t^{1-\alpha}$$

Finally, I transform the law of motion for capital. Ignoring the disaster risk we have

$$K_{t+1} \leq (1-\delta_K)K_t + I_t - \tfrac{\phi K_t}{2}\Big(\tfrac{K_{t+1}}{K_t} - 1\Big)^2$$

$$(\frac{Z_t}{Z_t}K_{t+1}) \leq (1-\delta_K)(\frac{Z_{t-1}}{Z_{t-1}}K_t) + (\frac{Z_t}{Z_t}I_t) - \tfrac{\phi}{2}(\frac{Z_{t-1}}{Z_{t-1}}K_t)\Big(\frac{Z_t}{Z_t}\frac{Z_{t-1}}{Z_{t-1}}\tfrac{K_{t+1}}{K_t} - 1\Big)^2$$

$$Z_t\hat{K}_{t+1} \leq (1-\delta_K)Z_{t-1}\hat{K}_t + Z_t\hat{I}_t - \tfrac{\phi}{2}Z_{t-1}\hat{K}_t\Big(g_t\tfrac{\hat{K}_{t+1}}{\hat{K}_t} - 1\Big)^2$$

$$\hat{K}_{t+1} \leq (1-\delta_K)g_t^{-1}\hat{K}_t + \hat{I}_t - \tfrac{\phi}{2}g_t^{-1}\hat{K}_t\Big(g_t\tfrac{\hat{K}_{t+1}}{\hat{K}_t} - 1\Big)^2$$

Appendix B

Discrete Approximation Monte Carlo Simulation

In this appendix I report more results from Monte Carlo simulation exercises designed to test how well the discrete method presented in this dissertation approximates the continuous process for aggregate productivity growth. The method is described in the body of the dissertation. I first compare the first four moments of the unconditional distributions for δz, x, and σ^2. In the table below the first two columns report these moments without stochastic volatility and the next two columns report the moments with stochastic volatility included.

We can see that this method approximates the unconditional moments of the continuous process reasonably well. Next, I compare the first four moments of several of the conditional distributions. The columns of each table are divided as before, the first two without stochastic volatility and the next two with it. The rows of each table report moments from the distribution of $x|\sigma^2$ and from the distribution of $\delta z|x, \sigma^2$. The first table reports these moments conditional on small values for x and σ^2.

Table B.1: Unconditional Moments

	Continuous	Markov	Continuous	Markov
$E[\Delta z_{t+1}]$	0.012	0.012	0.012	0.012
$\sigma(\Delta z_{t+1})$	0.010	0.010	0.006	0.004
$skew(\Delta z_{t+1})$	0.000	0.000	-0.002	0.003
$kurt(\Delta z_{t+1})$	2.96	2.90	34.0	30.9
$E[x_t]$	0.000	0.000	0.000	0.000
$\sigma(x_t)$	0.007	0.007	0.004	0.003
$skew(x_t)$	0.000	0.000	0.003	0.003
$kurt(x_t)$	2.95	2.83	35.3	32.8
$E[\sigma_t^2]$	-15.5	-15.9	-15.5	-15.9
$\sigma(\sigma_t^2)$	0.000	0.000	3.10	2.93
$skew(\sigma_t^2)$	0.000	0.001	0.039	0.000
$kurt(\sigma_t^2)$	2.71	2.66	2.68	2.66

Table B.2: Conditional on low values

	Continuous	Markov	Continuous	Markov
$E[\Delta z_{t+1}]$	-0.016	-0.016	0.012	0.012
$\sigma(\Delta z_{t+1})$	0.007	0.007	0.000	0.000
$skew(\Delta z_{t+1})$	0.000	0.000	0.000	0.000
$kurt(\Delta z_{t+1})$	2.96	2.82	2.96	2.82
$E[x_t]$	-0.011	-0.011	0.000	0.000
$\sigma(x_t)$	0.007	0.007	0.000	0.000
$skew(x_t)$	0.000	0.227	0.000	0.227
$kurt(x_t)$	2.96	2.82	2.96	2.87

The next table reports these moments conditional on mid-range values for x and σ^2.

Table B.3: Conditional on mid values

	Continuous	Markov	Continuous	Markov
$E[\Delta z_{t+1}]$	0.012	0.012	0.012	0.012
$\sigma(\Delta z_{t+1})$	0.007	0.007	0.000	0.000
$skew(\Delta z_{t+1})$	0.000	0.000	0.000	0.000
$kurt(\Delta z_{t+1})$	2.96	2.82	2.96	2.82
$E[x_t]$	0.000	0.000	0.000	0.000
$\sigma(x_t)$	0.007	0.007	0.000	0.000
$skew(x_t)$	0.000	0.000	0.000	0.000
$kurt(x_t)$	2.96	2.87	2.96	2.87

The last table reports these moments conditional on large values for x and σ^2

Table B.4: Conditional on high values

	Continuous	Markov	Continuous	Markov
$E[\Delta z_{t+1}]$	0.040	0.040	0.647	0.647
$\sigma(\Delta z_{t+1})$	0.007	0.007	0.156	0.156
$skew(\Delta z_{t+1})$	0.000	0.000	0.000	0.000
$kurt(\Delta z_{t+1})$	2.96	2.82	2.96	2.82
$E[x_t]$	0.011	0.011	0.250	0.250
$\sigma(x_t)$	0.007	0.007	0.155	0.156
$skew(x_t)$	0.000	-0.227	0.000	-0.227
$kurt(x_t)$	2.96	2.87	2.96	2.87

Once again, we can see that this method approximates the unconditional moments of the continuous process reasonably well.