Chapter 13
Asset pricing

Doubt is not a pleasant condition, but certainty is absurd.
— Voltaire

13.1. Introduction

This chapter explores how a fear of model misspecification affects prices of risky securities.[1] Without fear of misspecification, the price of a claim to a random future payoff equals the conditional expectation of the inner product of a stochastic discount factor and the random future payoff, evaluated using the representative agent's model.[2] When the representative agent fears misspecification of his approximating model, two such inner-product representations of asset prices are available. They differ in what they take as the model with respect to which the conditional expectation is evaluated. In the first one, the conditional expectation is evaluated with respect to the representative agent's worst-case model, a model that depends on the parameter θ that calibrates his fear of misspecification. A second representation of the same prices exists because the approximating model and the worst-case model put positive probabilities on the same events. This second representation evaluates the conditional expectation with respect to the approximating model. The first representation captures a concern about robustness by adjusting the probability distribution relative to the approximating model, while the second representation instead adjusts the stochastic discount factor (a.k.a. pricing kernel). In particular, to represent asset prices in terms of conditional expectations under the approximating model, the second representation multiplies the ordinary stochastic discount factor without fear of misspecification by the likelihood ratio, or Radon-Nikodym derivative, of the endogenous worst-case distorted model relative to the approximating model. When evaluated with respect to the worst-case probability distribution, the expected value of the logarithm of that likelihood ratio is the entropy measure that we used in chapter 2 to measure the proximity of models. It is also closely related to another entropy concept that describes bounds on the detection error probabilities of chapter 9 (see Anderson, Hansen, and Sargent (2003)).

[1] Studies of asset pricing under some form of model ambiguity include Dow and Werlang (1992), Epstein and Wang (1994), Chen and Epstein (2002), Maenhout (2004), Rigotti and Shannon (2005), and Hansen, Sargent, and Tallarini (1999).

[2] Without fear about misspecification, an agent can discard the adjective "approximating."

After reviewing asset pricing formulas in a standard model without a fear of misspecification, this chapter modifies those formulas to express a representative agent's fear of misspecification. As an example, we study asset pricing in the permanent income economy of chapter 10.

13.2. Approximating and distorted models

Chapters 11 and 12 describe planning problems and competitive equilibria for a class of linear-quadratic models. The consumption smoothing model of chapter 10 and the occupational choice model of section 12.7 are special cases of these models. The environment of chapter 11 is arranged so that without a fear of misspecification, the planning problem fits into the optimal linear regulator problem. Chapter 12 then uses a robust linear regulator to create a model in which the representative household's fear of misspecification is indexed by parameter $\theta > 0$. Equilibrium representations for prices and quantities can be determined from the solution of the robust linear regulator.

Chapter 11 describes matrices that portray the preferences, technology, and information structure of the economy. These can be assembled into matrices that define the robust linear regulator for a planning problem. The solution of the planning problem determines competitive equilibrium prices and quantities. Associated with the robust planning problem is the Bellman equation

$$-x'Px - p = \max_u \min_w \{r(x,u) + \theta\beta w'w + \beta E(-x^{*\prime}Px^* - p)\} \qquad (13.2.1)$$

where the extremization is subject to

$$x^* = Ax + Bu + C(\epsilon + w), \qquad (13.2.2)$$

where $\epsilon \sim \mathcal{N}(0, I)$ and $\theta \in (\underline{\theta}, +\infty]$. A Markov perfect equilibrium of this two-player zero-sum game is a pair of decision rules $u = -F(\theta)x, w = K(\theta)x$. The equilibrium determines the following two laws of motion for the state:

$$x_{t+1} = A^o x_t + C\epsilon_{t+1} \qquad (13.2.3)$$

and

$$x_{t+1} = (A^o + CK(\theta))x_t + C\epsilon_{t+1}, \qquad (13.2.4)$$

where $A^o = A - BF(\theta)$. For a given $\theta \in [\underline{\theta}, \infty)$, (13.2.3) is the approximating model under the robust rule for u, while (13.2.4) is the distorted worst-case model under the robust rule.

Where there is no fear of misspecification, $\theta = +\infty$. Chapter 11 describes a class of economies whose equilibria can be presented in the form (13.2.4) together with selector matrices that determine equilibrium prices and quantities as functions of the state x_t. In particular, quantities Q_t and scaled state-contingent prices p_t are linear functions of the state:

$$Q_t = S_Q x_t \tag{13.2.5a}$$

$$p_t = p_Q x_t. \tag{13.2.5b}$$

We shall soon remind the reader what we mean by scaled prices. We showed how to compute these in chapter 11 (see formulas (11.5.14), (11.5.21)).

To determine equilibria under a fear of misspecification, we simply set $\theta < +\infty$ in (13.2.1). Formulas for equilibrium prices and quantities from chapter 11 (i.e., the S_Q, M_Q in (13.2.5)) apply directly. Associated with an equilibrium under a fear of misspecification are the approximating transition law (13.2.3) and the distorted transition law (13.2.4) for the state x_t, as well as auxiliary equations for prices and quantities of the form (13.2.5).

The approximating and distorted equilibrium laws of motion (13.2.3) and (13.2.4) induce Gaussian transition densities[3]

$$f(x_{t+1}|x_t) \sim \mathcal{N}(A^o x_t, CC') \tag{13.2.6a}$$

$$\hat{f}(x_{t+1}|x_t) \sim \mathcal{N}((A^o + CK)x_t, CC'), \tag{13.2.6b}$$

where we use f without a $(\hat{\cdot})$ to denote a transition density under the approximating model and f with a $(\hat{\cdot})$ to denote a probability associated with the distorted model (13.2.4). These transition densities induce joint densities $f^{(t)}(x^t)$ on histories $x^t = [x_t, x_{t-1}, \dots, x_0]$ via

$$f^{(t)}(x^t) = f(x_t|x_{t-1}) f(x_{t-1}|x_{t-2}) \dots f(x_1|x_0) f(x_0),$$

and similarly for $\hat{f}^{(t)}(x^t)$. Let $f_t(x_t|x_0)$ denote the t-step transition densities

$$f_t(x_t|x_0) \sim \mathcal{N}(A^{ot} x_0, V_t) \tag{13.2.7a}$$

$$\hat{f}_t(x_t|x_0) \sim \mathcal{N}((A^o + CK)^t x_0, \hat{V}_t), \tag{13.2.7b}$$

where V_t satisfies the recursion $V_t = A^{o\prime} V_{t-1} A_o + CC'$ initialized from $V_1 = CC'$, and $\hat{V}_t$ satisfies the recursion $\hat{V}_t = (A^o + CK)' \hat{V}_{t-1} (A_o + CK) + CC'$ initialized from $\hat{V}_1 = CC'$.

[3] An alternative formulation in chapter 3 allows for a broader set of perturbations of a Gaussian approximating model by letting the minimizing agent choose an arbitrary density. Under that formulation, the minimizing agent would still choose a Gaussian transition density with the same conditional mean as (13.2.6b) but with conditional covariance $\hat{C}\hat{C}' = C(I - \theta^{-1} C' P C)^{-1} C'$.

13.3. Asset pricing without robustness

In section 11.7, we explained how the value of claims on risky streams of returns can be represented as the inner product of price and payout processes, where both the price and payout are expressed as functions of the planner's state vector x_t. In portraying the household's problem in a recursive competitive equilibrium, we needed to distinguish between the individual household's x_t and its "market wide" counterpart X_t that drives prices. Nevertheless, we showed that for the purpose of computing asset prices, we can exclude X_t from the state vector and simply use x_t as the state vector. Accordingly, in the remainder of this chapter, we express prices in terms of x_t and histories x^t.[4]

When $\theta = +\infty$, there is no discrepancy between the distorted and worst-case models and the following standard representative agent asset pricing theory applies. Let c_t denote a vector of time-t consumption goods. The price of a unit vector of consumption goods in period t contingent on the history x^t is[5]

$$q^{(t)}(x^t|x_0) = \beta^t \frac{u'(c_t(x^t))}{e_1 \cdot u'(c_0(x_0))} f^{(t)}(x^t|x_0), \tag{13.3.1}$$

where $c_t(x^t)$ is a possibly history-dependent state-contingent consumption process, $u'(c)$ is the vector of marginal utilities of consumption, and e_1 is a selector vector that pulls off the first consumption good, the time-zero value of which we take as numeraire. To make (13.3.1) well defined, we assume that $e_1 \cdot u'(c_0(x_0)) \neq 0$ with probability one. If we assume that the consumption allocation is not history-dependent, so that $c_t(x^t) = c(x_t)$, as it is true in the models that occupy us, then we can use the t-step pricing kernel

$$q_t(x_t|x_0) = \beta^t \frac{u'(c(x_t))}{e_1 \cdot u'(c(x_0))} f_t(x_t|x_0). \tag{13.3.2}$$

Let the owner of an asset be entitled to $\{y(x_t)\}_{t=0}^{\infty}$, a stream of a vector of consumption goods whose state-contingent price is given by (13.3.2). The time-zero price of the asset is

$$a_0 = \sum_{t=0}^{\infty} \int_{x_t} q_t(x_t|x_0) \cdot y(x_t) d\, x_t$$

---

[4] The household in a competitive economy would face prices that are the same functions of X_t and X^t.

[5] We denote by $u'(c_t)$ the vector of marginal utilities of the consumption vector c_t. In our model, $u'(c_t) = M_c x_t$.

or

$$a_0 = \sum_{t=0}^{\infty} \int_{x_t} \beta^t \frac{u'(c(x_t))}{e_1 \cdot u'(c(x_0))} y(x_t) f_t(x_t|x_0) d\, x_t. \tag{13.3.3}$$

We can represent (13.3.3) as

$$a_0 = \frac{E_0 \sum_{t=0}^{\infty} \beta^t u'(c(x_t)) \cdot y(x_t)}{e_1 \cdot u'(c(x_0))}. \tag{13.3.4}$$

In linear-quadratic general equilibrium models, $u'(c(x_t))$ and $y(x_t)$ are both linear functions of the state. This means that the price of an asset is the conditional expectation of a geometric sum of a quadratic form, as portrayed in (13.3.4). Equation (13.3.4) implies a Sylvester equation (see page 97). Thus, let

$$p_c(x_t) = \frac{u'(c(x_t))}{e_1 \cdot u'(c(x_0))}.$$

Then the asset price can be represented as

$$a_0 = E_0 \sum_{t=0}^{\infty} \beta^t p_c(x_t) \cdot y(x_t). \tag{13.3.5}$$

We can regard p_c as a scaled Arrow-Debreu price. We scale the Arrow-Debreu state price by dividing it by β^t times the pertinent conditional probability. Scaling the price system in this way facilitates computation of asset prices as conditional expectations of an inner product of state prices and payouts. Often $\beta^t p_c(x_t)$ is called a t-period *stochastic discount factor*. Below we shall also denote the stochastic discount factor as $m_{0,t} \equiv \beta^t p_c(x_t)$, so that (13.3.5) becomes

$$a_0 = E_0 \sum_{t=0}^{\infty} m_{0,t} \cdot y(x_t).$$

Hansen and Sargent (2008) provide a more complete treatment of asset pricing within linear-quadratic general equilibrium models. They show that (1) equilibrium scaled Arrow-Debreu prices and quantities have representations (13.2.5); (2) the information required to form the matrix S_Q is embedded in F, A, B from the optimal linear regulator problem; and (3) the matrices M_p that pin down the scaled Arrow-Debreu prices can be extracted from the matrix P in the value function $-x'Px - p$ and the matrix $A^o = A - BF$ that emerge from the planner's problem (see formulas (11.5.14), (11.5.21)). Thus, in such models

$$p_c(x_t) = M_c x_t / e_1 M_c x_0. \tag{13.3.6}$$

See (11.5.11), (11.5.13) in chapter 11 for a formula for M_c and more details.

13.4. Asset pricing with robustness

We activate a fear of misspecification by setting $\theta < +\infty$, which causes the transition densities (13.2.6a), (13.2.6b) under the approximating and distorted models to disagree. In addition, the formulas for S_Q and M_Q in (13.2.5) respond to the setting for θ, via the dependence of S_Q on $F(\theta)$ and the dependence of M_Q on the P that solves the Bellman equation (13.2.1). Again, see (11.5.14), (11.5.21). We give an example in section 12.7.

The price system that supports a competitive equilibrium can be represented in the forms (13.3.1) and (13.3.2), with the distorted densities $\hat{f}^{(t)}$ and $\hat{f}_t$ replacing the corresponding densities for the approximating model in (13.3.1) and (13.3.2). Thus, with a fear of misspecification, the time 0 price of the asset corresponding to (13.3.3) is

$$a_0 = \sum_{t=0}^{\infty} \int_{x_t} \beta^t p_c(x_t) \cdot y(x_t) \hat{f}_t(x_t|x_0) d\, x_t. \tag{13.4.1}$$

We can represent (13.4.1) as

$$a_0 = \hat{E}_0 \sum_{t=0}^{\infty} \beta^t p_c(x_t) \cdot y_t \tag{13.4.2}$$

where $\hat{E}$ denotes mathematical expectation using the distorted model (13.2.4), and $u'(c(x_t))$ must be computed using the M_Q in representation (13.2.5b) associated with θ.

13.4.1. Adjustment of stochastic discount factor for fear of model misspecification

Formula (13.4.2) represents the asset price in terms of the distorted measure that the planner uses to evaluate future utilities in the Bellman equation (13.2.1). To compute asset prices using this formula, we must solve a Sylvester equation using transition matrix $A^o + CK(\theta)$ from equation (13.2.4) to reflect that we are evaluating the expectation using the *distorted* transition law. We can also evaluate asset prices by computing expectations under the *approximating* model, but this requires that we adjust the stochastic discount factor to make the asset price satisfy (13.4.1). By dividing and multiplying by $f_t(x_t|x_0)$, we can represent (13.4.1) as

$$a_0 = \sum_{t=0}^{\infty} \int_{x_t} \beta^t p_c(x_t) \left(\frac{\hat{f}_t(x_t|x_0)}{f_t(x_t|x_0)} \right) \cdot y(x_t) f_t(x_t|x_0) d\, x_t \tag{13.4.3}$$

or

$$a_0 = E_0 \sum_{t=0}^{\infty} \beta^t p_c(x_t) \left(\frac{\hat{f}_t(x_t|x_0)}{f_t(x_t|x_0)} \right) \cdot y(x_t), \tag{13.4.4}$$

where the absence of a $(\hat{\cdot})$ from E denotes that the expectation is evaluated with respect to the approximating model (13.2.3).[6]

In summary, with a fear of misspecification, if we want to evaluate asset prices under the approximating model, we have to adjust the ordinary t-period stochastic discount factor $m_{0,t} = \beta^t p_c(x_t)$ for a concern about model misspecification and to use the modified stochastic discount factor[7]

$$m_{0,t}\left(\frac{\hat{f}_t(x_t|x_0)}{f_t(x_t|x_0)}\right).$$

For our linear-quadratic-Gaussian setting, the likelihood ratio is

$$L_t = \frac{\hat{f}_t(x_t|x_0)}{f_t(x_t|x_0)} = \exp\left[\sum_{s=1}^{t}\{\epsilon_s' w_s - .5 w_s' w_s\}\right].$$

13.4.2. Reopening markets

This section describes how to extend our asset pricing formulas to allow us to price "tail assets" that are traded at time t and that pay vectors of consumption $\{y_\tau\}_{\tau=t}^{\infty}$ for $t > 0$. We want the price to be stated in time-t units of the numeraire good.

Letting the t-step discount factor at time 0 be $m_{0,t} \equiv \beta^t p_c(x_t)$, (13.4.2) can be portrayed as

$$a_0 = \hat{E}_0 \sum_{t=0}^{\infty} m_{0,t} \cdot y_t \tag{13.4.5}$$

where $m_{0,t}$ is a vector of time-0 stochastic discount factors for pricing a vector of time-t payoffs. Define $m_{t,\tau}$ as the vector of corresponding time-t stochastic discount factors for pricing time $\tau \geq t$ payoffs[8]

$$m_{t,\tau} = \beta^{\tau-t} p_c(x_\tau)/e_1 p_c(x_t). \tag{13.4.6}$$

Then in time t units of the numeraire consumption good, the vector of payoffs $\{y_\tau\}_{\tau=0}^{\infty}$ is

$$a_t = \hat{E}_t \sum_{\tau=t}^{\infty} m_{t,\tau} y_\tau. \tag{13.4.7}$$

---

6 Notice the appearance of the same likelihood ratio in (13.4.4) used to define entropy in chapters 2 and 3 and to describe detection error probabilities in chapter 9.

7 Such a multiplicative adjustment to the stochastic discount factor $m_{0,t}$ carries over to nonlinear models.

8 We assume that $e_1 p_c(x_t) \neq 0$ with probability 1.

Equation (13.4.7) is equivalent to

$$a_t = E_t \sum_{\tau=t}^{\infty} (m_{t,\tau} m^u_{t,\tau}) \cdot y_\tau, \tag{13.4.8}$$

where the appropriate multiplicative adjustment $m^u_{t,\tau}$ to the stochastic discount factor is the likelihood ratio

$$\begin{aligned} m^u_{t,\tau} &= \frac{\hat{f}_{\tau-t}(x_\tau | x_t)}{f_{\tau-t}(x_\tau | x_t)} \\ &= \exp\left[\sum_{s=t}^{\tau} \{\epsilon_s' w_s - .5 w_s' w_s\}\right]. \end{aligned} \tag{13.4.9}$$

13.5. Pricing single-period payoffs

We now use the permanent income model of chapter 10 to shed light on the implications of a fear of misspecification for the equity premium. Let consumption be a scalar process and y_{t+1} be a scalar random payoff at time $t+1$. Without a fear of misspecification, the price at time t of a time $t+1$ payout is

$$a_t = E_t m_{t,t+1} y_{t+1}. \tag{13.5.1}$$

We follow Hansen and Jagannathan (1991) by applying the definition of a conditional covariance to (13.5.1) and using the Cauchy-Schwarz inequality to obtain

$$\left(\frac{a_t}{E_t m_{t,t+1}}\right) \geq E_t y_{t+1} - \left(\frac{\sigma_t(m_{t,t+1})}{E_t m_{t,t+1}}\right) \sigma_t(y_{t+1}). \tag{13.5.2}$$

The bound is attained by payoffs on the efficient frontier. The left side is the price of the risky asset relative to the price $E_t m_{t,t+1}$ of a risk-free asset that pays out 1 for sure next period. The term $\left(\frac{\sigma_t(m_{t,t+1})}{E_t m_{t,t+1}}\right)$ is the "market price of risk": it indicates the rate at which the price ratio $a_t / E_t m_{t,t+1}$ deteriorates with increases in the conditional standard deviation of the payout y_{t+1}.

Without imposing any theory about $m_{t,t+1}$, various studies have estimated the market price of risk $\left(\frac{\sigma_t(m_{t,t+1})}{E_t m_{t,t+1}}\right)$ from data on (a_t, y_{t+1}). For post World War II quarterly data, estimates of the market price of risk hover around .25. Hansen and Jagannathan's (1991) characterization of the equity premium puzzle is that .25 is much higher than would be implied by many theories that explicitly link $m_{t,t+1}$ to aggregate consumption. A standard benchmark is the theory $m_{t,t+1} = \beta u'(c_{t+1})/u'(c_t)$, where $u(\cdot)$ is a power utility function with power γ. That specification makes $m_{t,t+1} = \beta \left(\frac{c_{t+1}}{c_t}\right)^{\gamma}$. But aggregate consumption is a smooth series, so that the growth rate of consumption has a standard deviation so small that unless γ is implausibly large, the market price of risk implied by this theory of the stochastic discount factor $m_{t,t+1}$ remains far below the observed value of .25. Similarly, the permanent income model of chapter 10 that sets $m_{t,t+1} = M_c x_{t+1}/M_x x_t$ also implies too low a value of the market price of risk, again because the volatility of consumption growth is too small.[9]

How does imputing a concern about robustness to the representative agent impinge on these calculations? When the representative household is concerned about robustness, we have

$$a_t = E_t(m_{t,t+1} m^u_{t,t+1}) y_{t+1} \tag{13.5.3}$$

where from (13.4.9)

$$m^u_{t,t+1} = \exp\left[\epsilon'_{t+1} w_{t+1} - .5 w'_{t+1} w_{t+1}\right]. \tag{13.5.4}$$

By construction, $E_t m^u_{t,t+1} = 1$. Hansen, Sargent, and Tallarini (1999) (HST) computed that $E_t(m^u_{t,t+1})^2 = \exp(w'_{t+1} w_{t+1})$ so that

$$\sigma_t(m^u_{t,t+1}) = \sqrt{\exp(w'_{t+1} w_{t+1} - 1)} \approx |w'_{t+1} w_{t+1}|. \tag{13.5.5}$$

HST refer to $\sigma_t(m^u_{t,t+1})$ as the one-period market price of model uncertainty. Similarly, the $(\tau - t)$-period market price of model uncertainty is the conditional standard deviation of $m^u_{t,\tau}$ defined by (13.4.9). A fear of misspecification can boost the market price of risk by increasing these multiplicative adjustments to stochastic discount factors.

13.5.1. Calibrated market prices of model uncertainty

At this point, it might be useful for the reader to review the observational equivalence result in chapter 10. There we discussed the fact that there is a locus of (σ, β) pairs, all of which imply the same equilibrium quantities, i.e., the same consumption, investment, and output.[10] As in chapter 10, we follow HST and use the parameterization $\sigma \equiv -\theta^{-1}$. HST computed one-period market prices of risk for a calibrated version of the permanent income model described in chapter 10. In particular, they proceeded as follows:

---

[9] We return to these issues in chapter 14.

[10] Such observational equivalence seems also to be an excellent approximation in the non LQ model of Tallarini (2000).

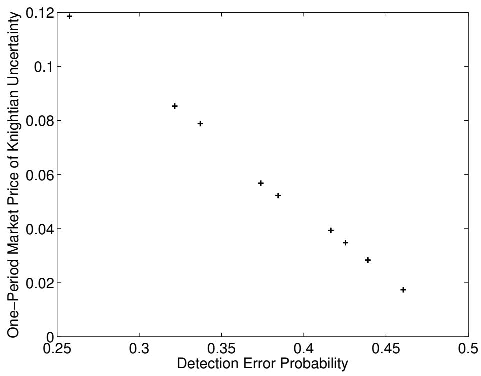

Figure 13.5.1: Market price of model uncertainty for one-period securities $\sigma_t(m_{t,t+1})^u$ as a function of detection error probability in the HST model.

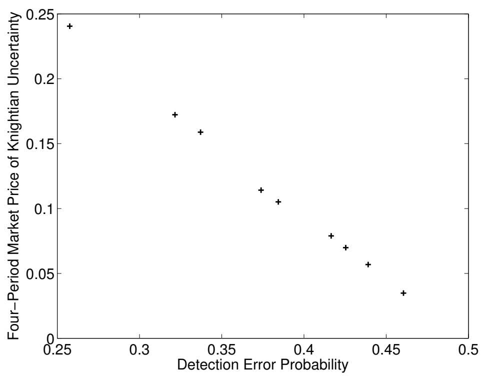

Figure 13.5.2: Market price of model uncertainty for four-period securities $\sigma_t(m_{t,t+4})^u$ as a function of detection error probability in the HST model.

1. Setting $\sigma = 0$ and $\beta R = 1$, HST used the method of maximum likelihood to estimate the remaining free parameters of the permanent income model of chapter 10.

2. HST used those maximum likelihood parameter estimates as the approximating model of the endowment processes $d_t^*, \hat{d}_t$ for a representative agent whose continuation values they used to price risky assets. Thus, HST took a stand on how the representative agent created his approximating model, something that robust control theory is silent about.

3. To study the effects of a fear of misspecification on asset prices while leav-

ing the consumption-investment allocation (c_t, i_t) intact, HST lowered σ below zero, but adjusted the discount factor according to the relation $\beta = \hat{\beta}(\sigma)$ given by equation (10.3.18), which defines a locus of (σ, β) pairs that freeze $\{c_t, i_t\}$. For each (σ, β) thereby selected, HST calculated market prices of model uncertainty and the detection error probabilities associated with distinguishing the approximating model from the worst-case model associated with σ. Figure 10.7.1 in chapter 10 reports those detection error probabilities as a function of σ. We are interested in the relation between the detection error probabilities and the j-period market prices of model uncertainty.

4. For one- and four-period horizons, figures 13.5.1 and 13.5.2 report the calculated market prices of model uncertainty plotted against the detection error probabilities. These graphs reveal two salient features. First, there appear to be approximately linear relationships between the detection error probabilities and the market prices of model uncertainty. For a continuous-time diffusion specification, Anderson, Hansen, and Sargent (2003) establish an exact linear relationship between the market price of risk and a bound on the detection error probabilities. To the extent that their bound is informative, their finding explains the striking pattern in these figures. Second, the market price of model uncertainty is substantial even for values of the detection error probability sufficiently high that it seems plausible to seek robustness against models that close to the approximating model. Thus, a detection error probability of .3 leads to a one-period market price of uncertainty of about .15, which can explain about half of the observed equity premium.

13.6. Concluding remarks

The asset pricing example of HST indicates how a little bit of concern about model misspecification can potentially substitute for a substantial amount of risk aversion when it comes to boosting theoretical values of market prices of risk. The boost in the market price of risk emerges from pessimism relative to the representative agent's approximating model. The form that the pessimism takes is endogenous, depending both on the transition law and the representative agent's discount factor and one-period return function. Pessimism has been proposed by several researchers as an explanation of asset pricing puzzles, e.g., Reitz (1988) and Abel (2002). The contribution of the robustness framework is to discipline the appeal to pessimism by restricting the direction in which the approximating model is twisted, and by how much, through the detection probability statistics that we use to restrict θ.