Chapter 9
Calibrating misspecification fears with detection error probabilities

The temptation to form premature theories upon insufficient data is the bane of our profession.
— Sherlock Holmes, in Sir Arthur Conan Doyle, The Valley of Fear, 1915

9.1. Introduction

This chapter proposes a strategy for calibrating the robustness parameter θ for some macroeconomic applications of multiplier robust control problems. Our procedure is to set θ so that, given the finite amount of data at his disposal, a decision maker would find it difficult statistically to distinguish members of a set of alternative models against which he seeks robustness (i.e., the models on and inside the entropy ball depicted in figure 1.7.1). We have in mind that, relative to the rate at which new data arrive, the decision maker's discount factor makes him sufficiently impatient that he cannot wait for those new data to resolve his model misspecification fears for him.

In chapter 14, we apply the approach of this chapter to calibrate an asset pricing example. There we demonstrate what we find to be a fascinating connection between the statistical detection error probabilities of this chapter and an object that is conventionally interpreted as the market price of risk, but that we suggest should instead be regarded as the market price of model uncertainty.

9.2. Entropy and detection error probabilities

Random disturbances in the transition law conceal the distortion of a perturbed model relative to the approximating model and can make the distortion difficult to detect statistically.[1] In this chapter, we illustrate how unconditional entropy governs statistics for distinguishing two models using moderate amounts of data. We use a statistical theory of model selection[2] to define a mapping from the parameter θ to a detection error probability for discriminating between the approximating model and an endogenous worst-case model associated with that θ. We use that detection error probability

[1] See chapter 3 for a formulation with stochastic shocks to the transition law.
[2] For example, see Burnham and Anderson (1998).

to determine a context-specific θ that is associated with a set of alternative models against which it is reasonable to seek robustness.[3]

9.2.1. The context-specific nature of θ

An outcome of the analysis in this chapter is a proposal to calibrate θ in a preliminary analysis the inputs of which include (1) a decision maker's approximating model, and (2) the decision maker's intertemporal objective function. In the course of describing detection error probabilities, we hope to clarify a mental experiment in which the decision maker is confronted with a model selection problem that differs markedly from the mental experiment involving *known* models with which Pratt (1964) confronted a decision maker when he wanted to extract measures of the decision maker's risk aversion. Chapter 14 draws out the differing natures of these mental experiments in the context of asset pricing.

9.2.2. Approximating and distorting models

For a given decision rule $u_t = -Fx_t$, we assume that the approximating model makes the state evolve according to the stochastic difference equation

$$x_{t+1} = A_o x_t + C\check{\epsilon}_{t+1}, \tag{9.2.1}$$

where now $\check{\epsilon}_{t+1}$ is an i.i.d. sequence of Gaussian disturbances with mean zero and identity contemporaneous covariance matrix. In turn, we will represent a distorted model as

$$\begin{aligned} x_{t+1} &= A_o x_t + C\left(\epsilon_{t+1} + w_{t+1}\right) \\ &= \hat{A} x_t + C\epsilon_{t+1} \end{aligned} \tag{9.2.2}$$

where $\hat{A} = A_o + C\kappa(\theta)$, $w_{t+1} = \kappa(\theta)x_t$, and ϵ_{t+1} is another i.i.d. Gaussian vector with mean 0 and identity covariance matrix. The transition densities associated with models (9.2.1) and (9.2.2) are absolutely continuous with respect to each other, i.e., they put positive probabilities on the same events.[4] Models that are not absolutely continuous with respect to each other are easy to distinguish empirically.

[3] For continuous-time models, Anderson, Hansen, and Sargent (2003) relate the penalty parameter and entropy to a bound on detection error probabilities as well as to alterations of market prices for risk associated with a concern about robustness.

[4] The two models (i.e., the two infinite-horizon stochastic processes) are absolutely continuous over finite intervals, a concept whose definition is reported by Hansen, Sargent, Turmuhambetova, and Williams (2006). The stochastic processes are not mutually absolutely continuous (over infinite intervals).

9.3. Detection error probabilities

Detection error probabilities can be calculated using likelihood ratio tests. Thus, consider two alternative models. Model A is the approximating model (9.2.1), and model B is the distorted model (9.2.2) associated with the context specific worst-case shock implied by θ. Consider a fixed sample of observations on the n-dimensional state vector x_t for $t = 0, \ldots, T-1$. Let L_i be the likelihood of that sample for model i. Form the log-likelihood ratio

$$\log \frac{L_A}{L_B}.$$

A likelihood ratio test selects model A when $\log \frac{L_A}{L_B} > 0$ and model B when $\log \frac{L_A}{L_B} < 0$. When model A generates the data, the probability of a model detection error is

$$p_A = \text{Prob}\left(\log \frac{L_A}{L_B} < 0 \middle| A \right).$$

In turn when model B generates the data, the probability of a model detection error is

$$p_B = \text{Prob}\left(\log \frac{L_A}{L_B} > 0 \middle| B \right).$$

Form the probability of a detection error by averaging p_A and p_B with prior probabilities over models A and B of .5:

$$p(\theta) = \frac{1}{2}(p_A + p_B).$$

Here, θ is the robustness parameter used to generate a particular model B by taking the associated worst-case perturbation of model A in light of a particular objective function for a decision maker. The following section shows in detail how to estimate the detection error probability by means of simulations. In a given context, we propose to set $p(\theta)$ to a reasonable number, then invert $p(\theta)$ to find a plausible value of θ.

9.4. Details

We now describe how to estimate detection error probabilities.

9.4.1. Likelihood ratio under the approximating model

Define w^A as the mean of the worst-case shock assuming that the actual data generating process is the approximating model, i.e., $w^A = \kappa x^A$ where x^A is generated under (9.2.1). Define $\hat{A} = A_o + C\kappa$. Then we can express the innovation under the worst-case model as

$$\begin{aligned} \epsilon_{t+1} &= (C'C)^{-1} C' \left(x_{t+1} - \hat{A} x_t \right) \\ &= \check{\epsilon}_{t+1} - \kappa x_t \\ &= \check{\epsilon}_{t+1} - w^A_{t+1}. \end{aligned} \tag{9.4.1}$$

The log-likelihood function under the approximating model is

$$\log L_A = -\frac{1}{T} \sum_{t=0}^{T-1} \{ n \log \sqrt{2\pi} + \frac{1}{2} \left(\check{\epsilon}_{t+1} \cdot \check{\epsilon}_{t+1} \right) \}$$

The log-likelihood function for the distorted model is

$$\begin{aligned} \log L_B &= -\frac{1}{T} \sum_{t=0}^{T-1} \{ n \log \sqrt{2\pi} + \frac{1}{2} \left(\epsilon_{t+1} \cdot \epsilon_{t+1} \right) \} \\ &= -\frac{1}{T} \sum_{t=0}^{T-1} \{ n \log \sqrt{2\pi} + \frac{1}{2} \left(\check{\epsilon}_{t+1} - w^A_{t+1} \right)' \left(\check{\epsilon}_{t+1} - w^A_{t+1} \right) \}. \end{aligned} \tag{9.4.2}$$

The log-likelihood ratio is therefore

$$r|A = \frac{1}{T} \sum_{t=0}^{T-1} \{ \frac{1}{2} w^{A'}_{t+1} w^A_{t+1} - w^{A'}_{t+1} \check{\epsilon}_{t+1} \}, \tag{9.4.3}$$

assuming that the approximating model is the data generating process. The second term in the above expression will vanish as $T \to \infty$, so that the log-likelihood ratio converges to the unconditional average value of $.5 w^{A\prime}_{t+1} w^A_{t+1}$, the measure of model discrepancy used throughout chapter 2, for example.

We can estimate the detection error probability conditional on model A by simulating a large number for x_t of length T under model A and counting the fraction of realizations for which $r|A$ computed as in (9.4.3) is negative.

9.4.2. Likelihood ratio under the distorted model

Now suppose that the data generating process is actually the distorted model (9.2.2). The innovations in the approximating model are linked to those in the distorted model by $\check{\epsilon}_{t+1} = \epsilon_{t+1} + w^B_{t+1}$, where $w^B_{t+1} = \kappa x^B_t$ and x^B_t is generated under (9.2.2).

Assuming that the distorted model generates the data, the log-likelihood function $\log L_B$ for the distorted model is

$$\log L_B = -\frac{1}{T} \sum_{t=0}^{T-1} \{ n \log \sqrt{2\pi} + \frac{1}{2} \left(\epsilon_{t+1} \cdot \epsilon_{t+1} \right) \}. \tag{9.4.4}$$

The log-likelihood function $\log L_A$ for the approximating model is

$$\begin{aligned}\log L_A &= -\frac{1}{T}\sum_{t=0}^{T-1}\{n\log\sqrt{2\pi}+\frac{1}{2}\left(\check{\epsilon}_{t+1}\cdot\check{\epsilon}_{t+1}\right)\} \\ &= -\frac{1}{T}\sum_{t=0}^{T-1}\{n\log\sqrt{2\pi}+\frac{1}{2}\left(\epsilon_{t+1}+w_{t+1}^B\right)'\left(\epsilon_{t+1}+w_{t+1}^B\right)\}.\end{aligned} \tag{9.4.5}$$

Hence, assuming that the distorted model B is the data-generating process, the log-likelihood ratio is

$$r|B = \frac{1}{T}\sum_{t=0}^{T-1}\{\frac{1}{2}w_{t+1}^{B\prime}w_{t+1}^B + w_{t+1}^{B\prime}\epsilon_{t+1}\}. \tag{9.4.6}$$

As $T\to\infty$, $r|B$ converges to the unconditional average value of one-period entropy $.5w_{t+1}^{B\prime}w_{t+1}^B$. Again, we can estimate p_B, the detection error probability conditioned on model B, by simulating a large number of paths of length T under model B and counting the fraction of realizations for which $r|B$ is positive.

9.4.3. The detection error probability

If we attach equal prior weights to models A and B, the overall detection error probability is

$$p(\theta) = \frac{1}{2}\left(p_A + p_B\right), \tag{9.4.7}$$

where $p_i = \text{freq}(r|i \le 0)$ for $i = A, B$. [5]

9.4.4. Breakdown point examples revisited

Figures 9.4.1 and 9.4.2 display estimated detection error probabilities for examples 2 and 3 from section 8.7, where we studied the effects of driving θ downwards toward the breakdown point $\underline{\theta}$. The figures record detection error probabilities for samples of length $T=50$ and $T=200$. We estimated the detection error probabilities for each value of $\sigma=-\theta^{-1}$ by averaging detection error rates over 100,000 simulations of length T.

The figures indicate that for $T=200$, the detection error probability for θ near the breakdown point is essentially zero for both examples. But a sample size of $T=50$ is small enough to leave the detection error probabilities as high as .05 near the breakdown point.

[5] The Matlab program `detection2.m` computes detection error probabilities.

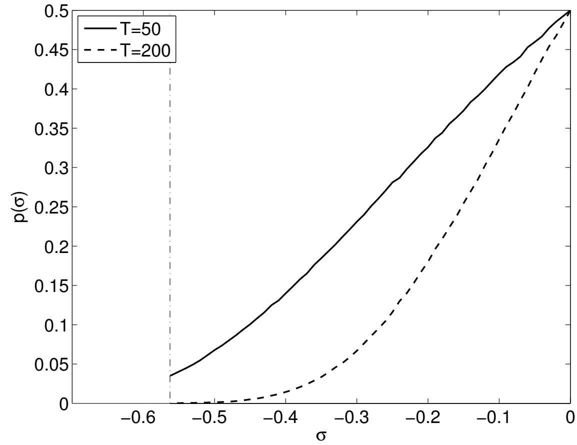

Figure 9.4.1: Detection error probability as a function of $\sigma = -\theta^{-1}$ for example 2 of section 8.7. The dotted vertical line denotes the breakdown point.

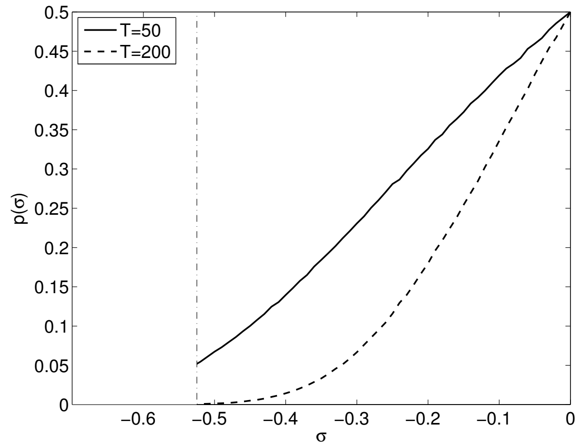

Figure 9.4.2: Detection error probability as a function of $\sigma = -\theta^{-1}$ for example 3 of section 8.7. The dotted vertical line denotes the breakdown point

9.5. Ball's model

We illustrate the use of detection error probabilities to discipline the choice of θ in the context of the simple dynamic model that Ball (1999) designed to study alternative rules by which a monetary policy authority might set an interest rate.[6] The model is

$$y_t = -\beta r_{t-1} - \delta e_{t-1} + \epsilon_t \tag{9.5.1}$$

$$\pi_t = \pi_{t-1} + \alpha y_{t-1} - \gamma \left(e_{t-1} - e_{t-2}\right) + \eta_t \tag{9.5.2}$$

$$e_t = \theta r_t + \nu_t, \tag{9.5.3}$$

where y is the log of real output, r is the real interest rate, e is the log of the real exchange rate, π is the inflation rate, and ϵ, η, ν are serially uncorrelated and mutually orthogonal disturbances. Ball assumed that the monetary authority wants to maximize

$$C = -E\left(\pi_t^2 + y_t^2\right).$$

The government sets the interest rate r_t as a function of the current state at t, which Ball shows can be reduced to y_t, e_t.

Ball motivates (9.5.1) as an open-economy IS curve and (9.5.2) as an open-economy Phillips curve. He uses (9.5.3) to capture effects of the interest rate on the exchange rate. Ball set the parameters $\gamma, \theta, \beta, \delta$ at the values $.2, 2, .6, .2$. Following Ball, we set the standard deviation of the innovation equal to $1, 1, \sqrt{2}$.

To discipline the choice of the parameter expressing a concern about robustness, we calculated the detection error probabilities for distinguishing Ball's model from the worst-case models associated with various values of $\sigma \equiv -\theta^{-1}$. We calculated these taking Ball's parameter values as the approximating model and assuming that $T = 142$ observations are available, which corresponds to 35.5 years of annual data for Ball's quarterly model. Figure 9.5.1 shows these detection error probabilities $p(\sigma)$ as a function of σ. Notice that the detection error probability is .5 for $\sigma = 0$, which verifies that the approximating model and the worst-case model are identical. The detection error probability falls to .1 for $\sigma \approx -.085$. If we think that a reasonable concern about robustness is to want rules that work well for alternative models whose detection error probabilities are .1 or greater, then $\sigma = -.085$ is a reasonable choice of this parameter.

We can use Ball's model to illustrate the robustness attained by alternative settings of the parameter θ. In particular, we compute a robust decision rule for Ball's model with $\sigma = -.085$ and compare its performance to the $\sigma = 0$ rule. For Ball's model, figure 9.5.2 shows that while robust rules do worse when the approximating model actually generates the data, their

[6] See Sargent (1999b) for further discussion of Ball's model from the perspective of robust decision theory.

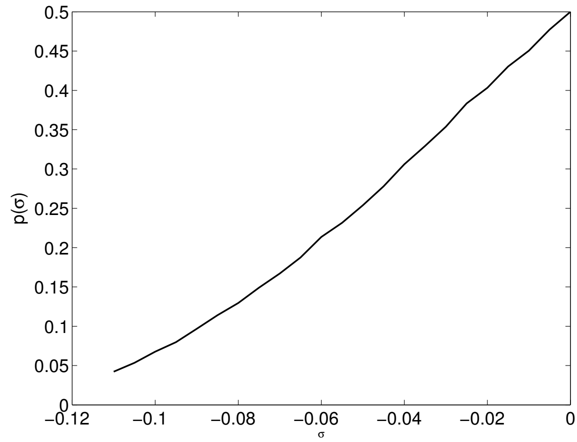

Figure 9.5.1: Detection error probability (coordinate axis) as a function of $\sigma \equiv -\theta^{-1}$ for Ball's model.

performance deteriorates more slowly with departures of the data-generating mechanism from the approximating model.

Figure 9.5.2 plots the value $C = -E(\pi^2 + y^2)$ attained by three rules under the alternative data-generating model associated with the worst-case model for the value of σ on the ordinate axis. The rules correspond to values $\sigma = 0, -.04, -.085$, respectively. Recall how the detection error probabilities computed above associate a value of $\sigma = -.085$ with a detection error probability of about .1. Notice how the robust rules (those computed with robustness parameter $\sigma = -.04$ or $-.085$) yield criterion values that deteriorate at a lower rate with model misspecification (they are flatter). Notice that the rule for $\sigma = -.085$ does worse than the $\sigma = 0$ or $\sigma = -.04$ rules when $\sigma = 0$, but is more robust in the sense that it deteriorates less when the model becomes more misspecified.

9.6. Concluding remarks

We shall use detection error probabilities to discipline the choice of θ again when we study a permanent income model of Hansen, Sargent, and Tallarini (1999) in chapter 10 and an asset pricing model of Tallarini (2000) in chapter 14.[7]

[7] Anderson, Hansen, and Sargent (2003) and Hansen (2007) analyzed some mathematical connections among entropy, market prices of model uncertainty, and bounds on detection error probabilities.

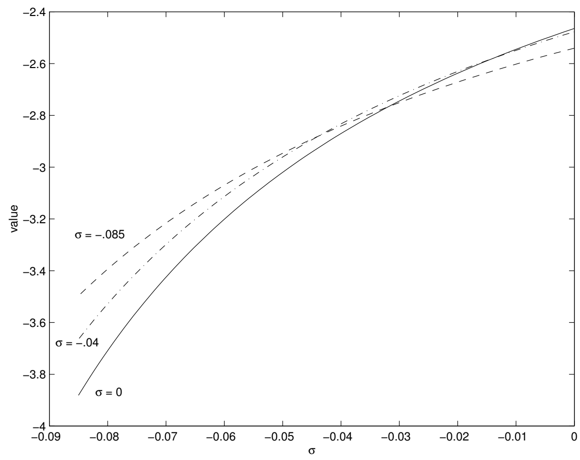

Figure 9.5.2: Value of criterion function $C = -E(\pi^2 + y^2)$ for three decision rules when the data are generated by the worst-case model associated with the value of σ on the horizontal axis: $\sigma = 0$ rule (solid line), $\sigma = -.04$ rule (dashed-dotted line), $\sigma = -.085$ (dashed) line.

Chapter 10
A permanent income model

If you would be wealthy, think of saving as well as getting.
— Benjamin Franklin

10.1. Introduction

The permanent income model is a good laboratory for exploring the consequences of a consumer's fears about misspecification of the stochastic process governing his labor income. We shall see that a consumer who distrusts his specification of the labor income or endowment process engages in a kind of precautionary savings that comes from his worst-case slanting of the probability law for the endowment process.[1] We use the Stackelberg multiplier game of chapter 7 to help us interpret how this probability slanting manifests itself in the permanent income model.

The permanent income model is also a good vehicle for gathering intuition from the frequency domain approach of chapter 8. A permanent income consumer is patient enough to smooth high-frequency fluctuations in income. That means that he automatically acquires robustness with respect to misspecification of the high-frequency details of the stochastic process for his labor income. But he is not patient enough to smooth low-frequency (i.e., very persistent) income fluctuations. Recognizing that the latter fluctuations cause the consumer the most trouble, the minimizing agent makes the worst-case shocks more persistent, an outcome that informs the consumer that his decision rule is most fragile with respect to low-frequency misspecifications of his income process. The robust permanent income consumer responds to those more persistent worst-case shocks by saving more than he would if he had no doubts about his endowment process. Thus, he engages in a form of precautionary savings that prevails even when he has quadratic preferences, which distinguishes this from the conventional form of precautionary savings that emerges for preferences that have convex marginal utilities.[2]

We apply the label "precautionary" because the effect increases with the volatility of innovations to endowments under the consumer's approximating model and because it also depends on the parameter θ that indexes his concern about robustness. Our model of precautionary savings exhibits the

[1] We can regard this context-specific slanting as corresponding to that mentioned by Fellner in the passage cited on page 38 of chapter 1.

[2] Leland (1968) and Miller (1974) are classic references on precautionary savings. See footnote 21 in this chapter.

usual feature that it modifies the certainty equivalence present in the linear-quadratic permanent income model. However, the model keeps the marginal propensity to save out of financial wealth equal to that out of human wealth, in contrast to models like those of Caballero (1990), where precautionary saving makes the marginal propensity to save out of human wealth exceed that out of financial wealth.[3]

To explore these issues, this chapter uses an equilibrium version of a permanent income model that Hansen, Sargent, and Tallarini (1999) (HST) estimated for U.S. consumption and investment data.[4] We restate (and extend in appendix B) an observational equivalence result of HST, who showed that activating a concern about robustness increases savings in the same way that increasing the discount factor would: the discount factor can be changed to offset the effect of a change in the robustness parameter θ on consumption and investment. HST thereby established that consumption and investment data alone are insufficient to identify both the robustness parameter θ and the subjective discount factor β.[5] We use the Stackelberg multiplier game from chapter 7 to shed more light on this observational equivalence proposition and the impact on decision rules of distortions in the conditional expectations under the worst-case model. We state another observational equivalence result for a new baseline model and use it to show that activating a concern about robustness still equalizes the marginal propensities to save out of human and nonhuman wealth.[6]

In addition, this chapter illustrates how the detection error probabilities described in chapter 9 can discipline plausible choices of θ and provides some numerical examples of how much robustness can be achieved by rules designed

3 See Wang (2003) for a treatment of how precautionary savings without robustness separates the marginal propensities to consume out of financial and nonfinancial wealth.

4 Hall (1978), Campbell (1987), Heaton (1993), and Hansen, Roberds, and Sargent (1991) applied versions of this model to aggregate U.S. time series data on consumption and investment.

5 Despite their failure to affect the consumption allocation, HST showed that such variations in (σ, β) do affect the relevant stochastic discount factor and therefore the valuation of risky assets. We shall take up asset pricing implications of the robust permanent income model in chapter 13.

6 Kasa (1999) constructs an observational equivalence result for the optimal linear regulator problem and its robust counterpart for the single-state, single-control case. He shows that for a given H_∞ decision rule there is a strictly convex function relating values of the H_∞ norm to the variable summarizing the relative cost of state versus control variability. Orlik (2006) establishes a general observational equivalence result between the standard optimal control and robust control problems. In an example application of the result, she shows that the same interest rate will be set by the policy maker who fully trusts his model as well as by the robust central banker provided that the preferences of the latter one with respect to inflation-output gap stabilization are appropriately specified.

with various settings of θ. In chapter 12, we describe how to decentralize the allocation chosen by the planner in the economy of this chapter. Then in chapter 13, we use that decentralized economy as a laboratory for studying ways to represent the effects on asset prices of a concern about robustness.

10.2. A robust permanent income theory

HST's model features a planner with preferences over consumption streams $\{c_t\}_{t=0}^{\infty}$, intermediated through service streams $\{s_t\}_{t=0}^{\infty}$.[7] Let b be a preference shifter in the form of a utility bliss point. The Bellman equation for the robust planner is

$$-x'Px - p = \sup_c \inf_w \left\{ -(s-b)^2 + \beta \left(\theta w^{*\prime} w^* - E x^{*\prime} P x^* - p \right) \right\} \tag{10.2.1}$$

where the maximization is subject to

$$\begin{aligned} s &= (1+\lambda) c - \lambda h && (10.2.2a) \\ h^* &= \delta_h h + (1-\delta_h) c && (10.2.2b) \\ k^* &= \delta_k k + i && (10.2.2c) \\ c + i &= \gamma k + d && (10.2.2d) \\ \begin{bmatrix} d \\ b \end{bmatrix} &= Uz && (10.2.2e) \\ z^* &= A_{22} z + C_2 (\epsilon^* + w^*) && (10.2.2f) \\ x' &= [\, h' \quad k' \quad z' \,]. && (10.2.2g) \end{aligned}$$

Here * denotes next period's value, $'$ denotes transpose, $\epsilon^* \sim \mathcal{N}(0, I)$, E is the expectation operator, c is consumption, s denotes a scalar service measure, and the law of motion mapping this period's state x into next period's state will be defined below. As before, the penalty parameter $\theta > 0$ governs concern about robustness to misspecification of the endowment process d and the preference shock process b embedded in (10.2.2e) and (10.2.2f). HST assumed that the eigenvalues of A_{22} are bounded in modulus by unity. We transform θ to the risk-sensitivity parameter $\sigma = -\theta^{-1}$. In (10.2.1), a scalar household service s_t is produced by the scalar consumption c_t via the household technology (10.2.2a) and (10.2.2b) where $\lambda > 0$ and $\delta_h \in (0,1)$. The household technology (10.2.2a),(10.2.2b) accommodates habit persistence or durability as in Ryder and Heal (1973), Becker and Murphy (1988), Sundaresan (1989), Constantinides (1990), and Heaton (1993). By construction, h_t is

[7] The model fits within the framework described in chapter 11. See page 257 for an additional stability condition that must be imposed.

a geometric weighted average of current and past consumption. Setting $\lambda > 0$ induces intertemporal complementarities. Consumption services depend positively on current consumption, but negatively on a weighted average of past consumption, a reflection of habit persistence.

There is a linear production technology (10.2.2d) where the capital stock k^* at the end of period t evolves according to (10.2.2c), where i is time t gross investment, and $\{d_t\}$ is an exogenously specified endowment process. The parameter γ is the (constant) marginal product of capital, and δ_k is the depreciation factor for capital. HST specified a bivariate ("two-factor") stochastic endowment process: $d_t = \mu_d + \tilde{d}_t + \hat{d}_t$.[8] They assumed that the two endowment processes are orthogonal and that both obey second-order autoregressions

$$(1-\phi_1 L)(1-\phi_2 L)\,\tilde{d}_t = c_{\tilde{d}}\left(\epsilon_t^{\tilde{d}} + w_t^{\tilde{d}}\right)$$

$$(1-\alpha_1 L)(1-\alpha_2 L)\,\hat{d}_t = c_{\hat{d}}\left(\epsilon_t^{\hat{d}} + w_t^{\hat{d}}\right)$$

where the vector ϵ_t is i.i.d. Gaussian with mean zero and identity covariance matrix, and $w_t^{\tilde{d}}, w_t^{\hat{d}}$ are distortions to the means of $\epsilon_t^{\tilde{d}}, \epsilon_t^{\hat{d}}$. HST estimated values of the ϕ_j's and α_j's that imply that the $\tilde{d}_t$ process is more persistent than the $\hat{d}_t$ process, as we see below.

Solving the capital evolution equation for investment and substituting into the linear production technology gives

$$c_t + k_t = Rk_{t-1} + d_t, \tag{10.2.3}$$

where

$$R \equiv \delta_k + \gamma,$$

which is the physical gross return on capital, taking into account that capital depreciates over time.[9]

Let the state vector be $x_t' = \left[h_{t-1}' \; k_{t-1}' \; d_{t-1} \; 1 \; d_t \; \tilde{d}_t \; \tilde{d}_{t-1}\right]$ (see Hansen, Sargent, and Wang (2002)). There is a set of state transition equations indexed by a $\{w_{t+1}\}$ process:

$$x_{t+1} = Ax_t + Bu_t + C\left(w_{t+1} + \epsilon_{t+1}\right) \tag{10.2.4}$$

where $u_t = c_t$ and $w_{t+1}' = [\,w_{t+1}^{\tilde{d}} \quad w_{t+1}^{\hat{d}}\,]'$ is the distortion to the conditional mean of ϵ_{t+1}. Let J_t be the sigma algebra induced by $\{x_0, \epsilon_s, 0 \le s \le t\}$.

[8] For two observed time series (c_t, i_t), HST's econometric specification needed at least two shock processes to avoid stochastic singularity.

[9] For HST's decentralized economy, R coincided with the gross return on a risk-free asset.

We require that the components of the solution for $\{c_t, h_t, k_t\}$ belong to L_0^2, the space of stochastic processes $\{y_t\}$ defined as

$$L_0^2 = \{y : y_t \text{ is in } J_t \text{ for } t = 0, 1, \dots \text{ and } E \sum_{t=0}^{\infty} R^{-t} (y_t)^2 \mid J_0 < +\infty\}.$$

Given x_0, the planner chooses a process $\{c_t, k_t\}$ with components in L_0^2 to solve the Bellman equation (10.2.1) subject to versions of (10.2.2a)-(10.2.2d) and (10.2.3).[10] In what follows we shall discuss HST's parameter values and some properties of their numerical solution. But first we show that in terms of its effects on consumption and investment, more concern about robustness works, *ceteris paribus*, like an increase in the discount factor.[11]

10.3. Solution when $\sigma = 0$

We apply results from chapter 7 to show that the robust decision rule for $\sigma < 0$ also solves a $\sigma = 0$ version of the model in which the maximizing agent in (10.2.1) replaces the approximating model with a particular distorted model for $[\, d_t' \quad b_t' \,]$. We shall eventually use that insight to study the identification of σ and β. To begin, this section solves the $\sigma = 0$ model.

10.3.1. The $\sigma = 0$ benchmark case

This subsection computes a solution of the planning problem in the $\sigma = 0$ case. Though we shall soon focus on the case when $\beta R = 1$, we also want the solution when $\beta R \neq 1$. Thus, for now we allow $\beta R \neq 1$. When $\sigma = 0$, the decision maker's objective reduces to

$$E_0 \sum_{t=0}^{\infty} \beta^t \{-(s_t - b_t)^2\}. \tag{10.3.1}$$

Formulate the planning problem as a Lagrangian by putting random Lagrange multiplier processes $2\beta^t \mu_{st}$ on (10.2.2a), $2\beta^t \mu_{ht}$ on (10.2.2b), and $2\beta^t \mu_{ct}$ on (10.2.3). First-order necessary conditions are

$$\mu_{st} = b_t - s_t \tag{10.3.2a}$$

[10] We can convert this problem into a special case of the control problem posed in chapter 7 as follows. Form a composite state vector x_t as described above, and let the control be given by $s_t - b_t$. Solve (10.2.2a) for c_t as a function of $s_t - b_t$, b_t, and h_{t-1} and substitute into equations (10.2.2b) and (10.2.3). Stack the resulting two equations along with the state evolution equation for z_t to form the evolution equation for x_{t+1}.

[11] However, in chapter 13, we shall show that (σ, β) pairs that imply observationally equivalent consumption and investment plans nevertheless imply different prices for risky assets. This finding is the basis of what Lucas (2003, p. 7) calls Tallarini's (2000) finding of "an astonishing separation of quantity and asset price determination."

$$\mu_{ct} = (1+\lambda)\,\mu_{st} + (1-\delta_h)\,\mu_{ht} \tag{10.3.2b}$$

$$\mu_{ht} = \beta E_t\left[\delta_h \mu_{ht+1} - \lambda \mu_{st+1}\right] \tag{10.3.2c}$$

$$\mu_{ct} = \beta R E_t \mu_{ct+1} \tag{10.3.2d}$$

and also (10.2.2a)-(10.2.2b) and (10.2.3). Equation (10.3.2d) implies that $E_t\mu_{ct+1} = (\beta R)^{-1}\mu_{ct}$. Then (10.3.2$b$) and (10.3.2$c$) solved forward imply that μ_{st}, μ_{ht} must satisfy $E_t\mu_{st+1} = (\beta R)^{-1}\mu_{st}$ and $E_t\mu_{ht+1} = (\beta R)^{-1}\mu_{ht}$. Therefore, μ_{st} has the representation

$$\mu_{st} = (\beta R)^{-1}\,\mu_{st-1} + \nu'\epsilon_t \tag{10.3.3}$$

for some vector ν. The endogenous volatility vector ν will play an important role below, and we shall soon tell how to compute it. The effects of the endogenous state variables h_{t-1}, k_{t-1} on consumption and investment are intermediated through the one-dimensional endogenous state vector μ_{st}, the marginal valuation of services.

Use (10.3.2a) to write $s_t = b_t - \mu_{st}$, substitute this into the household technology (10.2.2a)-(10.2.2b), and rearrange to get the system

$$c_t = \frac{1}{1+\lambda}\left(b_t - \mu_{st}\right) + \frac{\lambda}{1+\lambda} h_{t-1} \tag{10.3.4a}$$

$$h_t = \tilde{\delta}_h h_{t-1} + \left(1-\tilde{\delta}_h\right)\left(b_t - \mu_{st}\right) \tag{10.3.4b}$$

where $\tilde{\delta}_h = \frac{\delta_h+\lambda}{1+\lambda}$. Equation (10.3.4$a$) shows that knowledge of μ_{st}, b_t, h_{t-1} allows us to compute c_t, so that μ_{st} plays the role of the essential scalar endogenous state variable in the model. Equation (10.3.4b) can be used to compute

$$E_t \sum_{j=0}^{\infty} R^{-j} h_{t+j-1} = \left(1 - R^{-1}\tilde{\delta}_h\right)^{-1} h_{t-1} + \frac{R^{-1}\left(1-\tilde{\delta}_h\right)}{\left(1-R^{-1}\tilde{\delta}_h\right)} E_t \sum_{j=0}^{\infty} R^{-j}\left(b_{t+j} - \mu_{st+j}\right). \tag{10.3.5}$$

For the purpose of solving the first-order conditions (10.3.2), (10.2.2a), (10.2.2b), (10.2.3) subject to the side condition that $\{c_t, k_t\} \in L_0^2$, treat the technology (10.2.3) as a difference equation in $\{k_t\}$, solve forward, and take conditional expectations on both sides to get

$$k_{t-1} = \sum_{j=0}^{\infty} R^{-(j+1)} E_t\left(c_{t+j} - d_{t+j}\right). \tag{10.3.6}$$

Use (10.3.4a) to eliminate $\{c_{t+j}\}$ from (10.3.6), then use (10.3.3) and (10.3.5). Solve the resulting system for μ_{st} to get

$$\mu_{st} = \Psi_1 k_{t-1} + \Psi_2 h_{t-1} + \Psi_3 \sum_{j=0}^{\infty} R^{-j} E_t b_{t+j} + \Psi_4 \sum_{j=0}^{\infty} R^{-j} E_t d_{t+j}, \quad (10.3.7)$$

where

$$\begin{aligned} \Psi_1 &= -(1+\lambda) R \left(1 - R^{-2}\beta^{-1}\right) \left[\frac{1 - R^{-1}\tilde{\delta}_h}{1 - R^{-1}\tilde{\delta}_h + \lambda\left(1 - \tilde{\delta}_h\right)} \right] \\ \Psi_2 &= \frac{\lambda\left(1 - R^{-2}\beta^{-1}\right)}{1 - R^{-1}\tilde{\delta}_h + \lambda\left(1 - \tilde{\delta}_h\right)} \\ \Psi_3 &= \left(1 - R^{-2}\beta^{-1}\right) \\ \Psi_4 &= R^{-1}\Psi_1. \end{aligned} \quad (10.3.8)$$

Equations (10.3.7), (10.3.4), and (10.2.3) represent the solution of the planning problem when $\sigma = 0$.[12]

To compute ν in (10.3.3), it is useful to notice that formula (10.3.7) can be rewritten as

$$\begin{aligned} \mu_{st} = (\beta R)^{-1} \mu_{st-1} &+ \Phi_3 \sum_{j=0}^{\infty} R^{-j} \left(E_t b_{t+j} - E_{t-1} b_{t+j}\right) \\ &+ \Phi_4 \sum_{j=0}^{\infty} R^{-j} \left(E_t d_{t+j} - E_{t-1} d_{t+j}\right) \end{aligned} \quad (10.3.9)$$

where

$$\mu_{st-1} = \Phi_1 k_{t-1} + \Phi_2 h_{t-1} + \Phi_3 \sum_{j=0}^{\infty} R^{-j} E_{t-1} b_{t+j} + \Phi_4 \sum_{j=0}^{\infty} R^{-j} E_{t-1} d_{t+j}.$$

The third and fourth terms of equation (10.3.9) are scalars Ψ_3 and Ψ_4 multiplied by the innovations at t in the present values of b_t and d_t, respectively. Let the moving average representations for b_t and d_t be

$$b_t = \zeta_b(L)\epsilon_t \quad (10.3.10)$$

$$d_t = \zeta_d(L)\epsilon_t, \quad (10.3.11)$$

[12] When $\beta R = 1$, (10.3.7) makes μ_{st} depend on a geometric average of current and future values of b_t. Therefore, both the optimal consumption service process and optimal consumption depend on the difference between b_t and a geometric average of current and expected future values of b. So there is no "level effect" of the preference shock on the optimal decision rules for consumption and investment. However, the level of b_t will affect equilibrium asset prices.

where $\zeta_b(L) = U_b(I - A_{22}L)^{-1}C_2$ and $\zeta_d(L) = U_d(I - A_{22}L)^{-1}C_2$ from (10.2.2e).

By applying a formula of Hansen and Sargent (1980), it is easy to show that the innovations in the present values of b_t and d_t, respectively, equal the present values of the coefficients in these moving average representations.[13] Therefore, representation (10.3.9) can be rewritten as

$$\mu_{st} = (\beta R)^{-1} \mu_{st-1} + \left[\Psi_3 \zeta_b \left(R^{-1}\right) + \Psi_4 \zeta_d \left(R^{-1}\right)\right] \epsilon_t. \tag{10.3.12}$$

Comparing this with (10.3.3), we see that

$$\nu' = \Psi_3 \zeta_b \left(R^{-1}\right) + \Psi_4 \zeta_d \left(R^{-1}\right). \tag{10.3.13}$$

An equivalent way to compute ν is to note that formula (10.3.7) for μ_{st} can be represented in matrix notation as

$$\mu_{st} = M_s x_t \tag{10.3.14}$$

$$x_t = A_o x_{t-1} + C\epsilon_t \tag{10.3.15}$$

where x_t is the state vector k_{t-1}, h_{t-1}, z_t, where $z_t = [\, d_{t-1} \quad 1 \quad d_t \quad \tilde{d}_t \quad \tilde{d}_{t-1} \,]'$ the matrix M_s is determined by equation (10.3.7) and A_o, C and the laws of motion for b_t, d_t determine the law of motion for the entire state under the optimal rule for c_t.[14] It follows that $\mu_{st} = M_s A_o x_{t-1} + M_s C \epsilon_t$, which must agree with (10.3.3), so that $\mu_{s,t-1} \equiv M_s A_o x_{t-1}$ and

$$\nu' \equiv M_s C. \tag{10.3.16}$$

The scalar $\alpha = \sqrt{\nu' \nu}$ plays an important role in the argument below. It obeys

$$\alpha = \sqrt{M_s C C' M_s'}. \tag{10.3.17}$$

In the widely studied special case that $\lambda = \delta_h = 0$, so that $s_t = c_t$ and $\mu_{st} = b_t - c_t$, (10.3.7), (10.3.8) imply that the marginal propensity to consume out of "non-human wealth" $R k_{t-1}$ and the marginal propensity to consume out of "human wealth" $\sum_{j=0}^{\infty} R^{-j} E_t d_{t+j}$ both equal $-\Psi_1$. It is a well-known feature of the linear-quadratic model that these marginal propensities to consume are equal. Notice that human wealth is formed by discounting expected future endowments at the risk-free rate.

---

[13] The present value of the moving average coefficients plays an important role in linear-quadratic permanent income models. See Flavin (1981), Campbell (1987), and Hansen, Roberds, and Sargent (1991).

[14] Here C is the matrix that appears in (10.2.4) above. See Hansen and Sargent (2008, chapter 10) for fast ways to compute A_o, M_s, C for a class of models that includes that of this chapter.

10.3.2. Observational equivalence for quantities of $\sigma = 0$ and $\sigma \neq 0$

In the $\sigma = 0$ case, HST followed Hall (1978) and imposed that $\beta R = 1$. HST then showed that for fixed values of all other parameters, there is a set of (β, σ) pairs that leave the consumption-investment plan unaltered. In particular, if as we vary σ we also vary β according to[15]

$$\hat{\beta}(\sigma) = \frac{1}{R} + \frac{\sigma \alpha^2}{R-1}, \tag{10.3.18}$$

then we leave unaltered the decision rules for (c_t, i_t). Here $\alpha^2 = \nu' \nu$, where ν, as defined in (10.3.13), is a vector in the following martingale representation for the marginal utility of services μ_{st} that prevails as a special case of (10.3.3) when $\sigma = 0$ and $\beta R = 1$:

$$\mu_{st} = \mu_{st-1} + \nu' \epsilon_t.$$

(Also see equation (10.3.12).) The following subsection explains how HST constructed the locus identified by (10.3.18).

10.3.3. Observational equivalence: intuition

Here is the basic idea underlying the observational equivalence proposition. As already mentioned, a single factor μ_{st} summarizes the endogenous state variables h_{t-1}, k_{t-1}. When $\beta R = 1$ and $\sigma = 0$, it has the law of motion

$$\mu_{st} = \mu_{st-1} + \nu' \epsilon_t,$$

which can also be represented as

$$\mu_{st} = \mu_{st-1} + \alpha \tilde{\epsilon}_t \tag{10.3.19}$$

where $\tilde{\epsilon}_t$ is a scalar i.i.d. process with zero mean and unit variance and where $\alpha = \sqrt{\nu' \nu}$ verifies $\alpha \tilde{\epsilon}_t = \nu' \epsilon_t$. We generate our observational equivalence result by reverse engineering. We activate a concern about robustness by setting $\sigma < 0$, but insist that (10.3.19) continue to describe μ_{st} under the approximating model in order to make sure that the (c_t, i_t) allocation remains the same when $\sigma < 0$. For $\sigma < 0$ and a new value $\hat{\beta}$ that is to be determined, the worst-case model for μ_{st} is

$$\mu_{st} = \mu_{st-1} + \alpha(\tilde{\epsilon}_t + \tilde{w}_t) \tag{10.3.20}$$

[15] See footnote 23 of this chapter.

or

$$\mu_{st} = \left(1 + \alpha K\left(\sigma, \hat{\beta}\right)\right)\mu_{st-1} + \alpha\tilde{\epsilon}_t \tag{10.3.21}$$

where $\tilde{w}_t = K(\sigma, \hat{\beta})\mu_{st-1}$. Evidently, (10.3.21) implies that $\hat{E}_t\mu_{st+1} = (1 + \alpha K(\sigma, \hat{\beta}))\mu_{st}$, where $\hat{E}$ is the mathematical expectation with respect to the distorted model. Notice that we once again use the modified certainty equivalence principle. With a concern about robustness, the decision maker's choices conform to the following version of the Euler equation (10.3.3):

$$\hat{E}_t\mu_{st+1} = \left(\hat{\beta}R\right)^{-1}\mu_{st},$$

where $\hat{E}_t$ is evaluated with respect to the worst-case model (10.3.21) and $\hat{\beta}$ is a new value for β that we design to offset the effects of setting $\sigma < 0$. That is, if possible, we want to choose $\hat{\beta}$ to compensate for using the worst-case distribution to evaluate expectations in the above Euler equation. And we want the distorted model to be associated with the same approximating model (10.3.19) that generates the original c_t, i_t allocation. But according to (10.3.21), if the approximating model is to be (10.3.19), then $\hat{E}_t\mu_{st+1} = (1 + K(\sigma, \hat{\beta})\alpha)\mu_{st}$. Thus, for a given $\sigma < 0$, we want to find a replacement $\hat{\beta}$ for β that enables us to verify $(\hat{\beta}R)^{-1} = (1 + \alpha K(\sigma, \hat{\beta}))$, where $K(\sigma, \hat{\beta})$ solves the minimization problem that gives rise to the worst-case shock. In summary, we want to solve $1 = (\hat{\beta}R)(1 + \alpha K(\sigma, \hat{\beta}))$ for $\hat{\beta}$ as a function of σ. The proof of our observational equivalence Theorem 10.3.1 shows that a solution for $\hat{\beta}$ exists, that it is unique, and that it satisfies (10.3.18).

10.3.4. Observational equivalence: formal argument

Following HST, we begin by assuming that $\beta R = 1$ when $\sigma = 0$. We state

Theorem 10.3.1. *(Observational Equivalence, I) Fix all parameters, including R, except (σ, β). Suppose $\beta R = 1$ when $\sigma = 0$. There exists a $\underline{\sigma} < 0$ such that for any $\sigma \in (\underline{\sigma}, 0)$, the optimal consumption-investment plan for $(0, \beta)$ is also chosen by a robust decision maker when parameter values are $(\sigma, \hat{\beta}(\sigma))$ and where $\hat{\beta}(\sigma) < \beta$ satisfies (10.3.18).*

Proof. The proof is constructive. Begin with an allocation $\{\bar{s}_t, \bar{c}_t, \bar{k}_t, \bar{h}_t\}$ for a benchmark $\sigma = 0, \beta R = 1$ economy, then form a comparison economy with a $\sigma \in [\underline{\sigma}, 0]$, where $\underline{\sigma}$ is the lowest value for which the solution of (10.3.25) reported below is real. The comparison economy fixes all parameters except (σ, β) at their values for the benchmark economy. We then construct a discount factor $\hat{\beta} < \beta$ for which $\{\bar{s}_t, \bar{c}_t, \bar{k}_t, \bar{h}_t\}$ is also the allocation for the $\sigma < 0$ economy.

When $\beta R = 1$, (10.3.3) becomes

$$\mu_{st} = \mu_{st-1} + \nu' \epsilon_t. \tag{10.3.22}$$

The optimality of the allocation under the original $(0, \beta)$ implies that (10.3.22) is satisfied, which in turn implies that $E_t \mu_{st+1} = \mu_{st}$ and (10.3.7) are satisfied where E_t is the expectation operator under the approximating model. We seek a new value $\sigma < 0$ and an associated value $\hat{\beta}(\sigma)$ for which: (1) (10.3.22) remains satisfied under the approximating model; (2) the robust decision maker chooses the $(\bar{\cdot})$ allocation, which requires that $\hat{\beta} R \hat{E}_t \mu_{st+1} = \mu_{st}$, where $\hat{E}$ is the expectation with respect to the worst-case model associated with $(\sigma, \hat{\beta})$ when the approximating model obeys (10.3.22). However, when the approximating model satisfies (10.3.22), the worst-case model associated with $(\sigma, \hat{\beta})$ implies that $\hat{E}_t \mu_{st+1} = \hat{\zeta}(\hat{\beta}) \mu_{st}$, where $\hat{\zeta}(\hat{\beta}) = (1 + \alpha K(\sigma, \hat{\beta})) > 1$ can be found by solving the pure forecasting problem[16] associated with law of motion $\mu_{st} = \mu_{st-1} + \nu'(\epsilon_t + w_t)$, (10.3.22), one-period return function $-\mu_{st}^2 = -(b_t - s_t)^2$, and discount factor $\hat{\beta}$. If the σ-robust decision maker is to choose a decision rule that sustains (10.3.22) under the approximating model, so that (1) and (2) both prevail, $\hat{\beta}$ must verify

$$\hat{\beta} R \hat{\zeta} \left(\hat{\beta} \right) = 1. \tag{10.3.23}$$

To complete the argument, we compute $\hat{\zeta}(\hat{\beta})$ by solving a pure forecasting problem to find the distorted expectation operator $\hat{E}_t$. We use the recipe given in formulas (7.C.10) on page 168 and (7.C.26) and (7.C.27) on page 171. Taking (10.3.22) as given under the approximating model and noting that $\mu_{st}^2 = (b_t - s_t)^2$, the evil agent in the pure forecasting problem seeks to minimize $-\sum_{t=0}^{\infty} \hat{\beta}^t (\mu_{st}^2 + \hat{\beta} \frac{1}{\sigma} w_{t+1}^2)$ under the distorted law $\mu_{st} = \mu_{st-1} + \alpha w_t$, where $\alpha = \sqrt{\nu' \nu}$ (see (10.3.22)). Taking μ_s as the state, the evil agent's Bellman equation (7.C.27) is[17]

$$-P \mu_s^2 = -\mu_s^2 + \hat{\beta} \min_w \left(-\frac{1}{\sigma} w^2 - P \left(\mu_s + \alpha w \right)^2 \right). \tag{10.3.24}$$

The scalar P that solves (10.3.24) is

$$-P \left(\hat{\beta} \right) = \frac{\hat{\beta} - 1 + \sigma \alpha^2 + \sqrt{\left(\hat{\beta} - 1 + \sigma \alpha^2 \right)^2 + 4 \sigma \alpha^2}}{-2 \sigma \alpha^2}. \tag{10.3.25}$$

---

16 See page 171 for the definition of a pure forecasting problem.

17 We exploit a version of certainty equivalence and ignore the stochastic parts of the Bellman equation and the law of motion for μ_s.

Let $\hat{\zeta}(\hat{\beta}) = A + CK(\sigma, \hat{\beta}) = 1 + \alpha K(\sigma, \hat{\beta})$, where $w = K(\sigma, \hat{\beta})\mu_s$ is the formula for the worst-case shock and $A + CK$ is the state transition matrix for the distorted law of motion as in chapter 7. Applying formula (7.C.21) for $K(\sigma, \hat{\beta})$ in chapter 7 to the current problem gives

$$\hat{E}_t \mu_{st+1} = \hat{\zeta} \mu_{st} \tag{10.3.26}$$

where

$$\hat{\zeta} = \hat{\zeta}\left(\hat{\beta}\right) = 1 + \frac{\sigma\alpha^2 P\left(\hat{\beta}\right)}{1 - \sigma\alpha^2 P\left(\hat{\beta}\right)} = \frac{1}{1 - \sigma\alpha^2 P\left(\hat{\beta}\right)}. \tag{10.3.27}$$

Hansen, Sargent, and Wang (2002) solve (10.3.23), (10.3.25), and (10.3.27) to obtain

$$\hat{\beta}(\sigma) = \frac{1}{R} + \frac{\sigma\alpha^2}{R-1}. \tag{10.3.28}$$

For $\sigma \in [\underline{\sigma}, 0]$, equation (10.3.28) defines a locus of $(\sigma, \hat{\beta})$'s, each point of which is observationally equivalent to $(0, \beta)$ for observations on (c_t, k_t) because each supports the benchmark ($\sigma = 0$) allocation. ∎

This proposition means that with the appropriate adjustments in β given by $\hat{\beta}(\sigma)$, the robust decision maker chooses precisely the same quantities $\{c_t, k_t\}$ as a decision maker without a concern for robustness. Thus, as far as these quantity observations are concerned, the robust $(\sigma < 0, \hat{\beta}(\sigma))$ version of the permanent income model is observationally equivalent to the benchmark $(\sigma = 0, \beta)$ version.[18] However, as we shall see in chapter 13, (σ, β) pairs that imply equivalent allocations because they satisfy (10.3.28) do *not* imply the same asset prices. The reason is that as we alter (σ, β) within this observationally equivalent set, we alter continuation valuations by altering $\mathcal{D}(P)$.

[18] The asset pricing theory developed by HST, which is encoded in (10.3.23), implies that the price of a sure claim on consumption one period ahead is R^{-1} for all t and for all $(\sigma, \hat{\beta})$ in the locus (10.3.18). Therefore, these different parameter pairs are also observationally equivalent with respect to the risk-free rate. In this model, the technology (10.2.3) ties down the risk-free rate. For a version of the model with quadratic costs of adjusting capital, the risk-free rate comes to depend on σ, even though the observations on quantities are approximately independent of σ. See Hansen and Sargent (2008).

10.3.5. Precautionary savings interpretation

The consumer's concern about model misspecification activates a particular kind of precautionary savings motive that underlies our observational equivalence proposition. A concern about robustness inspires the consumer to save *more*. Decreasing his discount factor induces the consumer to save *less*. The observational equivalence proposition asserts that these two effects can be arranged to offset each other.

The following experiment highlights the precautionary motive for savings. Take the base model with $\sigma = 0$ used in our proof of Theorem 10.3.1. Then activate a concern about robustness by setting $\sigma < 0$, but offset its effect on consumption by setting β equal to $\hat{\beta}(\sigma)$. Notice from (10.3.28) that $\hat{\beta}(\sigma)$ depends on the volatility parameter α. Consider a $(\sigma, \hat{\beta}(\sigma))$ pair corresponding to a given $\alpha > 0$. The innovation volatility associated with a positive α means that future endowments are forecast with error. If future endowments and preference shifters *could* be forecast perfectly, then at the value $\beta = \hat{\beta}(\sigma)$, the consumer would choose to make his capital stock, and therefore also his consumption, drift downward because discounting is large relative to the marginal productivity of capital. Investment would be sufficiently unattractive that the optimal linear rule would eventually send both consumption and capital below zero.[19, 20] However, when randomness is activated (i.e., the innovation variances are positive), this downward drift is arrested or even completely offset, as it is in our observational equivalence proposition. Thus, our robust control interpretation of the permanent-income decision rule delivers a form of precautionary savings.

The precautionary savings coming from a concern about robustness differs in structure from another, perhaps more familiar, kind of precautionary savings motive that has attracted much attention in the macroeconomics literature and that emerges when a positive variance of the innovations to the endowment process interacts with a convex derivative of the marginal utility of consumption.[21] In contrast, the precautionary savings induced by a con-

[19] Introducing nonnegativity constraints in capital and/or consumption would induce nonlinearities into the consumption and savings rules, especially near zero capital. But investment would remain unattractive in the presence of those constraints for experiments like the one we are describing here. See Deaton (1992) for a critical survey and quantitative assessment of consumption models with binding borrowing constraints.

[20] As emphasized by Carroll (1992), even when the discount factor is small relative to the interest rate, precautionary savings can emerge when there is a severe utility cost for zero consumption. Such a utility cost is absent in our formulation.

[21] Take the Euler equation $E_t \beta R u'(c_{t+1}) = u'(c_t)$ and assume that $\beta R = 1$ so that $E_t u'(c_{t+1}) = u'(c_t)$. If u' is a convex function, then applying Jensen's inequality implies $E_t c_{t+1} > c_t$, so that consumption is expected to grow when the conditional distribution of

cern about robustness emerges because the consumer wants to protect himself against mistakes in specifying conditional *means* of shocks to the endowment. Thus, a concern for robustness inspires precautionary savings because of how fears of misspecification are expressed in conditional *first* moments of shocks. This type of precautionary saving does not require that the marginal utility of consumption be convex and occurs even in models with quadratic preferences, as we have shown.

A concern about robustness affects consumption by slanting probabilities in the way Fellner described in the passage cited on page 38 of this book. The household saves more for a given β because it makes pessimistic forecasts of future endowments. Precisely how pessimism manifests itself depends on the detailed structure of the permanent income model and the temporal properties of the endowment process, as we shall discuss in the next section.

10.4. Observational equivalence and distorted expectations

In this section, we use insights from a Stackelberg multiplier game to interpret Theorem 10.3.1. In the Stackelberg multiplier game, decisions for the maximizing player can be computed by solving his Euler equations using a particular distorted law of motion to form conditional expectations of the shocks.[22]

In the benchmark $\sigma = 0, \beta R = 1$ case that is contemplated in Theorem 10.3.1, the solution of the planning problem is determined by equations (10.3.4), (10.2.3), and (10.3.7), where the Ψ_j's satisfy (10.3.8) with $\beta R = 1$. For a $\sigma \in [\underline{\sigma}, 0)$ and a $\hat{\beta} = \hat{\beta}(\sigma)$, the decision rule for the robust planner is characterized by equations (10.3.4), (10.2.3), and the following modified version of (10.3.7):

$$\mu_{st} = \hat{\Psi}_1 k_{t-1} + \hat{\Psi}_2 h_{t-1} + \hat{\Psi}_3 \sum_{j=0}^{\infty} R^{-j} \hat{E}_t b_{t+j} + \hat{\Psi}_4 \sum_{j=0}^{\infty} R^{-j} \hat{E}_t d_{t+j}, \quad (10.4.1)$$

where $\hat{\Psi}_j$ are determined by (10.3.8) with $\beta = \hat{\beta}(\sigma)$; and $\hat{E}_t$ is the conditional expectation operator with respect to the distorted law of motion for the state x_t. The observational equivalence Theorem 10.3.1 implies that (10.4.1)

c_{t+1} is not concentrated at a point. Such consumption growth reflects precautionary savings. See Ljungqvist and Sargent (2004, chapter 16) for an analysis of these precautionary savings models.

[22] While the timing protocol for the Stackelberg multiplier game differs from the Markov perfect timing embedded in game (10.2.1), chapter 7 showed that identical equilibrium outcomes and recursive representations of equilibria prevail under these different timing protocols.

and (10.3.7) are identical solutions for μ_{st}. By substituting for the terms in expected future values, the solutions (10.3.7) and (10.4.1) can also be expressed as $\mu_{st} = M_s x_t$ and $\mu_{st} = \hat{M}_s x_t$. Observational equivalence requires that $M_s = \hat{M}_s$. This requires that the $\hat{\Psi}_j$'s and $\hat{E}$ mutually adjust to keep M_s fixed.[23]

To expand on this point, consider the special case that $\lambda = \delta_h = 0$, so that we need not retain h_{t-1} as a state variable. Also, assume for simplicity that $b_t = b$, so that the preference shock is constant. Shutting down the volatility of b prevents distortions in it from affecting the robust decision rule. Then equating the right sides of (10.3.7) and (10.4.1) gives

$$\begin{aligned} 0 = &\left(\Psi_4 - \hat{\Psi}_4\right) R k_{t-1} + \left(\Psi_3 - \hat{\Psi}_3\right)\left(1 - R^{-1}\right)^{-1} b \\ &+ \Psi_4 \sum_{j=0}^{\infty} R^{-j} E_t d_{t+j} - \hat{\Psi}_4 \sum_{j=0}^{\infty} R^{-j} \hat{E}_t d_{t+j} \end{aligned} \tag{10.4.2}$$

where Ψ_j without hats denotes values of Ψ_j that satisfy (10.3.8) and those with hats satisfy (10.3.8) evaluated at $\beta = \hat{\beta}(\sigma)$. Equation (10.4.2) shows how the observational equivalence result asserts offsetting alterations in the coefficients Ψ_j and the distorted expectations operator $\hat{E}_t$ used to form the expected sum of discounted future endowments that defines human wealth.

The distorted expectations operator is to be interpreted in terms of the recursive formulation of the maximizing player's problem in a Stackelberg multiplier game of chapter 7. The Euler equation approach used to derive (10.3.7) or (10.4.1) presumes the following timing protocol. *After* the minimizing player has committed to an entire path for the w_{t+1} process, the maximizing agent faces the following recursive representation of the motion for the endowment and preference shocks:

$$X_{t+1} = \left(A - BF\left(\sigma, \hat{\beta}\right) + CK\left(\sigma, \hat{\beta}\right)\right) X_t + C\tilde{\epsilon}_{t+1} \tag{10.4.3a}$$

$$\begin{bmatrix} b_t \\ d_t \end{bmatrix} = S X_t \tag{10.4.3b}$$

where $\tilde{\epsilon}_{t+1}$ is an i.i.d. shock identical in distribution to that of ϵ_{t+1}.[24] Because the minimizing player has committed himself to a stochastic process for $\{w_{t+1}\}$ that implies the recursive representation (10.4.3) of the endowment and preference shock processes, the maximizing player takes the X_t

[23] Note from formula (10.3.17) that M_s determines α, a key parameter defining the observational equivalence locus (10.3.18). Thus, because M_s remains fixed, so does α so long as $(\sigma, \hat{\beta})$ obey (10.3.18).

[24] In (10.4.3), X_t is used to attain a recursive representation of the worst-case endowment and preference shock processes that keeps them exogenous to the maximizer's decisions.

process as exogenous and uses the forecasting rule $\hat{E}_t X_{t+j} = (A - BF(\sigma,\hat{\beta}) + CK(\sigma,\hat{\beta}))^j X_t$ to form forecasts of (b_{t+j}, d_{t+j}) in (10.4.1). These forecasts, together with (10.4.1), (10.3.4), and (10.2.3) can be solved to yield a decision rule $c_t = -\mathcal{F}\begin{bmatrix} x_t \\ X_t \end{bmatrix}$ as in chapter 7. After computing the decision rule as a function of x_t, X_t, we equate $x_t = X_t$; that gives the maximizing agent's decision rule in the form $c_t = -Fx_t$.[25]

10.4.1. Distorted endowment process

Figures 10.4.1 and 10.4.2 illustrate the probability slanting that leads to precautionary savings. The figures assume HST's parameter values that are reported in appendix A and record impulse response functions for the total endowment d_t under the approximating model and a worst-case model associated with $\sigma = -.0001$, where β is adjusted according to (10.3.18) as required under our observational equivalence proposition in order to preserve the same decision rule $F(\sigma,\hat{\beta})$ for different σ's.[26]

For the approximating and the worst-case models with $\sigma = -.0001$, the figures report the response of the total endowment d_t to innovations ϵ_t^* and $\hat{\epsilon}_t$ in the relatively permanent and transitory components of the endowment, $\tilde{d}_t$, $\hat{d}_t$, respectively. Under the distorted model, the impulse response functions diverge and the eigenvalue of $A - BF(\sigma,\hat{\beta}) + CK(\sigma,\hat{\beta})$ that has maximum modulus increases from its value of unity under the approximating model to 1.0016.

The distorted endowment processes respond to innovations with more persistence than they do under the approximating model. With a fixed β, the increased persistence makes the agent save more than under the approximating model, which the observational equivalence proposition offsets by decreasing the household's patience via (10.3.18).

Figures 10.5.1 and 10.5.2 record impulse response functions for the total endowment d_t under the approximating model and a worst-case model associated with $\sigma = -.0001$, *where β is held fixed at HST's benchmark value.* Because these figures do not adjust the discount factor according to (10.3.18) as it was done for figures 10.4.1 and 10.4.2, the distorted impulse response functions deviate from those of the approximating model even more than those of these earlier figures. The reduction in β from (10.3.18) works through two channels to make the $\sigma < 0$ decision rule equal to that for a $\sigma = 0$ rule: (1) it brings the distorted impulse response functions closer to those of the

[25] The procedure of first optimizing, then setting $x_t = X_t$ to eliminate X_t is a common way of formulating rational expectations equilibria in macroeconomics, where it is sometimes called the "Big K, little k" method.

[26] The observational equivalence proposition makes the decision rules equivalent under the approximating model.

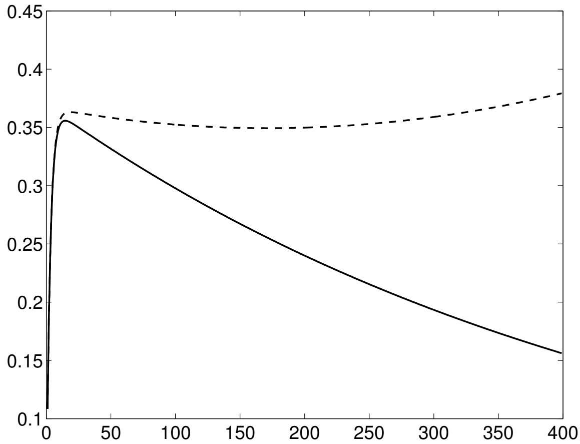

Figure 10.4.1: Response of total endowment d_t to innovation in 'permanent' component $\tilde{d}_t$ under the approximating model (dotted line) and the distorted model associated with the worst-case shock (dashed line) for the $\sigma = -.0001,\ \beta = \beta(\sigma)$ model.

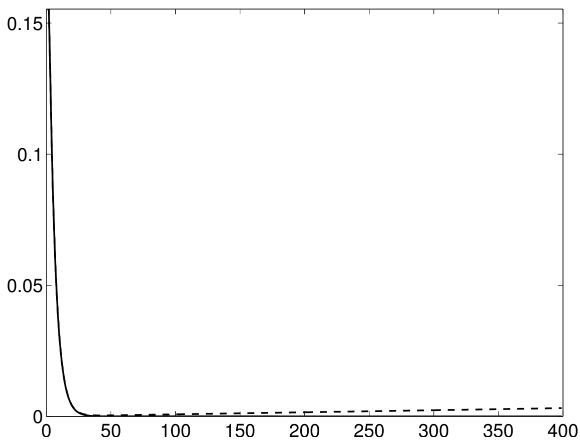

Figure 10.4.2: Response of total endowment d_t to innovation in 'transitory' component $\hat{d}_t$ under the approximating model (solid line) and the distorted model associated with the worst-case shock (dotted line) for the $\sigma = -.0001, \beta = \beta(\sigma)$ model.

approximating model, and (2) more impatience combats the precautionary savings motive.

10.5. Another view of precautionary savings

To interpret the precautionary savings motive inherent in our model, appendix B asserts another observational equivalence proposition. Theorem 10.B.1 takes a baseline case where $\beta R = 1$ and shows that in its effects on (c, i), activating a concern for robustness operates just like an *increase* in the discount factor. This result is useful because the $\beta R = 1$ case forms a benchmark in the permanent income literature (for example, see Hall (1978)). Theorem 10.B.1 shows that the effects of activating concerns about robustness by putting $\sigma < 0$ are replicated by keeping $\sigma = 0$ and raising β so that $\beta R > 1$.

To use this result to shed more light on how the precautionary motive manifests itself in the decision rule for consumption, we consider the important special case that $\delta = \lambda = \tilde{\delta} = 0$. Then $\mu_{st} = \mu_{ct} = b - c_t$ and the consumption Euler equation $(10.3.2d)$ without a concern about robustness becomes

$$b - c_t = E_t \left[(\beta R) (b - c_{t+1}) \right].$$

If $\beta R > 1$, this equation implies that $b - c_t > E_t(b - c_{t+1})$, or

$$c_t < E_t c_{t+1}, \tag{10.5.1}$$

so that the optimal policy is to make consumption grow on average.

Theorem 10.B.1 shows that when $\beta R = 1$, a concern about robustness $(\sigma < 0)$ has the same effect on c_t, i_t as setting $\sigma = 0$ and setting a particular β for which $\beta R > 1$. Therefore, when $\beta R = 1$, the precautionary savings that occurs when $\sigma < 0$ follows from $(10.5.1)$. Activating a concern about robustness imparts an upward drift to the expected consumption profile.

We can also use Theorem 10.B.1 to discuss some facts about the decision rule for consumption in our special case that $\lambda = \delta = \tilde{\delta} = 0$. The solution $(10.3.8)$ for $\sigma = 0$ implies the consumption rule

$$c_t = \left(1 - R^{-2}\beta^{-1}\right) \left[R k_{t-1} + E_t \sum_{j=0}^{\infty} R^{-j} d_{t+j} \right] + \left(\frac{(R\beta)^{-1} - 1}{R - 1} \right) b. \tag{10.5.2}$$

Notice that the marginal propensity to consume out of financial wealth Rk_{t-1} equals that out of human wealth $E_t \sum_{j=0}^{\infty} R^{-j} d_{t+j}$.[27] Further, an increase in β *decreases* the constant $\left(\frac{(R\beta)^{-1}-1}{R-1}\right) b$ and *increases* the marginal propensity

[27] This implication of precautionary savings coming from robustness differs from that coming from convex marginal utility functions, where precautionary savings reduces the marginal propensity to consume out of endowment income relative to that from financial wealth. See Wang (2003).

to consume $1 - R^{-2}\beta$. Relative to the baseline $\beta R = 1$ case, raising β raises the marginal propensity to consume out of wealth by $R^{-1}(1 - (R\beta)^{-1})$. This increase in the marginal propensity to consume still allows wealth to have an upward trajectory because of the reduction in the second term $\frac{(R\beta)^{-1}-1}{R-1}b$.

The permanent income model of consumption has an interpretation in terms of the frequency domain that is familiar to macroeconomists. It is that his concave one-period utility function makes the permanent income consumer dislike high-frequency volatility in consumption and therefore adjust his asset holdings in a way that protects his consumption from high-frequency fluctuations in income. The following section views the precautionary savings that are inspired by fears of model misspecification from the vantage point of the frequency domain.

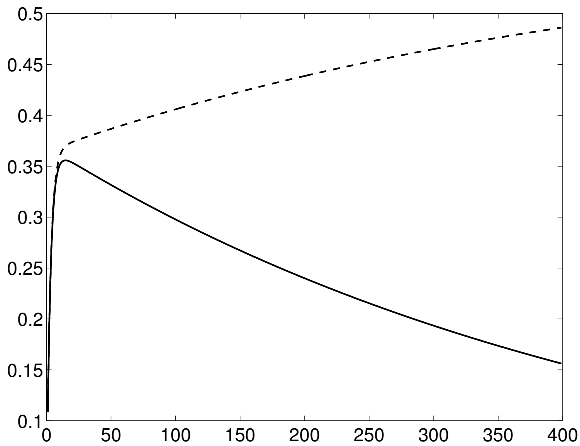

Figure 10.5.1: Response of total endowment d_t to innovation in "permanent" component $\tilde{d}_t$ under the approximating model (solid line) and the distorted model associated with the worst-case shock (dotted line) for $\sigma = -.0001$, with β at benchmark value.

10.6. Frequency domain representation

This section uses HST's estimated permanent income model to illustrate features of the frequency domain decompositions of the consumer's objective function and of the worst-case shocks for different values of σ.

Importing some notation from chapter 8, denote the transfer function from shocks ϵ_t to the "target" $s_t - b_t$ as $G(\zeta)$. For the baseline model with

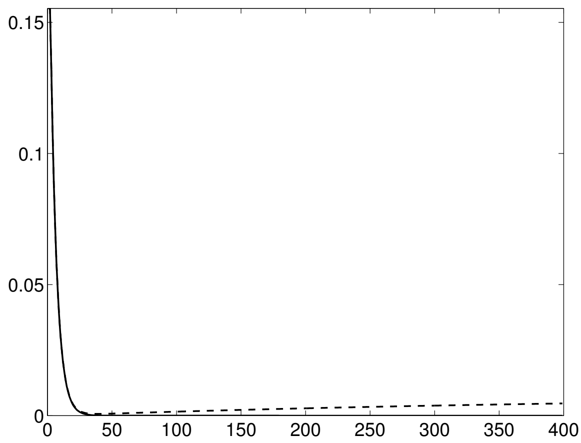

Figure 10.5.2: Response of total endowment d_t to innovation in "permanent" component $\tilde{d}_t$ under the approximating model (solid line) and the distorted model associated with the worst-case shock (dotted line) for $\sigma = -.0001$ with β at benchmark value.

habit persistence, recall formula (8.4.3) for the frequency decomposition of H_2:

$$H_2 = -\frac{1}{2\pi}\int_{-\pi}^{\pi} \text{trace}\left[G\left(\sqrt{\beta}\exp\left(i\omega\right)\right)' G\left(\sqrt{\beta}\exp\left(i\omega\right)\right)\right] d\omega.$$

A reinterpretation of formula (8.3.5) also gives us the frequency domain representation

$$E\sum_{t=0}^{\infty}\beta^t w_t' w_t = \frac{1}{2\pi}\int_{-\pi}^{\pi} W\left(\sqrt{\beta}\exp\left(i\omega\right)\right)' W\left(\sqrt{\beta}\exp\left(i\omega\right)\right) d\omega.$$

Figure 10.6.1 shows $G(\sqrt{\beta}\exp(i\omega))'G(\sqrt{\beta}\exp(i\omega))$ for the baseline ($\sigma = 0$) line as a function of frequency ω; $G'G$ is larger at lower frequencies. Remember that $G(\zeta) = (I-(A_o - BF)\zeta)^{-1}C$ embodies the consumer's optimal decision rule F. The noise process ϵ_t upon which $G(\zeta)$ operates is i.i.d. under the approximating model, so that the spectral density matrix of ϵ_t is constant across frequencies. But seeing that the consumer's policy makes him most vulnerable to the low-frequency components of ϵ_t, the minimizing player makes the conditional mean of the worst-case shock w_{t+1} highly serially correlated. For two values of σ, figure 10.6.2 shows frequency decompositions of trace $W(\zeta)'W(\zeta)$ for $\zeta = \sqrt{\beta}\exp(i\omega)$. Notice how most of the power is at the lowest frequencies. As we varied σ from zero to the two values in figure

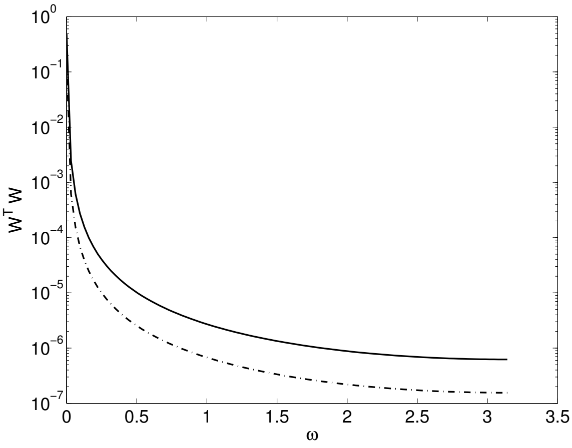

Figure 10.6.2: Frequency decomposition of volatility of worst-case shocks for $-\theta^{-1} = \sigma = -.0001$ (solid line) and $\sigma = -.00005$ (dotted line); trace$[W(\zeta)'W(\zeta)]$ plotted as a function of ω where $\zeta = \sqrt{\beta}\exp(i\omega)$.

10.6.2, we adjusted $\beta = \hat{\beta}$ according to (10.3.18), which keeps the robust $\sigma < 0$ decision rule for consumption equal to that for the baseline no robustness ($\sigma = 0$) model. Notice that [trace $W(\zeta)'W(\zeta)$] varies directly with the absolute value of σ.

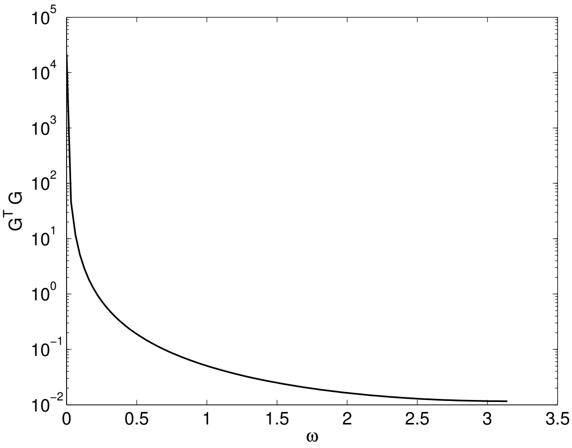

Figure 10.6.1: Frequency decomposition of criterion function; $G(\zeta)'G(\zeta)$ plotted as a function of ω where $\zeta = \sqrt{\beta}\exp(i\omega)$.

10.7. Detection error probabilities

For HST's parameter values, figure 10.7.1 reports detection error probabilities associated with various values of σ, adjusting β according to (10.3.18) so as to keep the decision rule fixed. These detection error probabilities were calculated by the method of chapter 9 for a sample of the same length that HST used to estimate their model and for HST's initial conditions. To calculate the detection error probabilities, all other parameter values were frozen at the values from table table 10.A.1. Then the formula for the worst-case distortions $w_{t+1} = K(\sigma, \hat{\beta})x_t$ was used to compute an alternative law of motion for the endowment process.

For different values of σ, figure 10.7.1 records the detection error probabilities for distinguishing an approximating model from a worst-case model associated with that value of σ. The approximating model is

$$x_{t+1} = (A - BF(0, \beta))\, x_t + C\epsilon_{t+1}$$

while the distorted model associated with σ is

$$x_{t+1} = \left(A - BF(0, \beta) + CK\left(\sigma, \hat{\beta}\right)\right) x_t + C\tilde{\epsilon}_{t+1}$$

where both ϵ_t and $\tilde{\epsilon}_t$ are i.i.d. processes with mean zero and identity covariance matrix, and where $F(0, \beta) = F(\sigma, \hat{\beta})$ by the observational equivalence proposition.

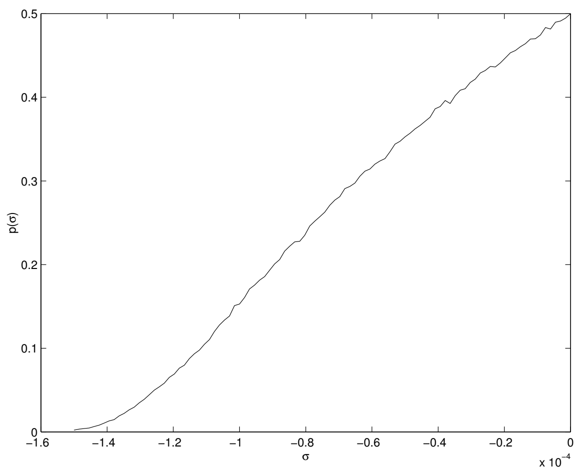

Figure 10.7.1: Detection error probabilities as a function of σ.

The detection error probability equals .5 for $\sigma = 0$ because then the models are identical and, hence, cannot be distinguished. The detection error probability falls with σ because the two models differ more from one another. In the following section, we use figure 10.7.1 to guide a choice of σ as measuring the size of a set of models against which it is plausible for the consumer to seek robustness.

10.8. Robustness of decision rules

For $\sigma = -\theta^{-1}$, express the equilibrium decision rules of game (10.2.1) as

$$c_t = -F(\sigma) x_t \tag{10.8.1a}$$

$$w_{t+1} = K(\sigma) x_t \tag{10.8.1b}$$

and express $s_t - b$ as $H(\sigma)x_t$. For possibly different values σ_1, σ_2, consider the law of motion of the state under the consumption plan $F(\sigma_2)x_t$ and the worst-case shock process $K(\sigma_1)x_t$:

$$x_{t+1} = (A - BF(\sigma_2) + CK(\sigma_1)) x_t + C\epsilon_{t+1}. \tag{10.8.2}$$

For x_0 given, we evaluate the expected payoff

$$\pi(\sigma_1; \sigma_2) = -E_{0,\sigma_1} \sum_{t=0}^{\infty} \beta^t x_t' H(\sigma_2)' H(\sigma_2) x_t \tag{10.8.3}$$

under the law of motion (10.8.2). That is, we want to evaluate the performance of the rule designed by setting σ_2 when the data are generated by the distorted model associated with σ_1. For three values of σ_2, figure 10.8.1 plots $\pi(\sigma_1; \sigma_2)$ as a function of the parameter σ_1 that indexes the magnitude of the distortion in the model generating the data. By construction, the $\sigma_2 = 0$ decision rule does better than the other rules when $\sigma_1 = 0$. But its performance deteriorates faster with decreases in σ_1 below zero than do the more robust $\sigma_1 = -.00004, \sigma_1 = -.00008$ rules.

From figure 10.8.1, $\sigma = -.00004$ is associated with a detection error probability of over .3, and $\sigma = -.00008$ with a detection error probability about .2. It is plausible for the consumer to want decisions that are robust against alternative models that are as close as the worst-case models associated with those values of σ.

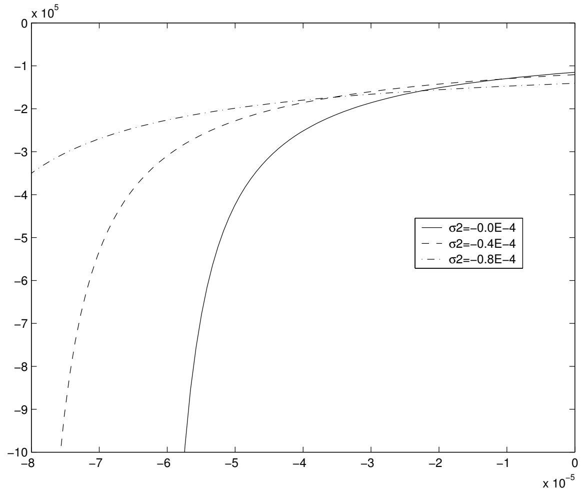

Figure 10.8.1: Payoff

$$\pi(\sigma_1;\sigma_2) = -E_{0,\sigma_1} \sum_{t=0}^{\infty} \beta^t x_t' H(\sigma_2)' H(\sigma_2) x_t$$

as a function of σ_1 on the ordinate axis for decision rules $F(\sigma_2)$ associated with three values of σ_2.

10.9. Concluding remarks

Different observationally equivalent (σ, β) pairs identified by Theorem 10.3.1 have different implications concerning (1) pricing risky assets; (2) the amounts required to compensate the planner for confronting different amounts of risk; (3) the amount of model misspecification used to justify the planner's decisions if risk sensitivity is reinterpreted as reflecting concerns about model misspecification. Hansen, Sargent, and Tallarini (1999) and Hansen, Sargent, and Wang (2002) have analyzed the asset pricing implications of the model in this chapter. They show that although movements along the observational equivalence locus described by (10.3.18) do not affect consumption and investment, they put an adjustment for fear of model misspecification into asset prices and boost what macroeconomists typically measure as market prices of risk. In chapter 13, we shall describe how standard asset pricing formulas are altered when a representative consumer is concerned about robustness. There we shall describe an asset pricing theory under a concern about robustness in the context of a class of general equilibrium models. The model from this chapter can be viewed as a special case of this class of models.

Table 10.A.1: HST's parameter estimates

Object	Habit Persistence	No Habit Persistence
Risk Free Rate	.025	.025
β	.997	.997
δ_h	.682	
λ	2.443	0
α_1	.813	.900
α_2	.189	.241
ϕ_1	.998	.995
ϕ_2	.704	.450
μ_d	13.710	13.594
$c_{\hat{d}}$	.155	.173
$c_{\tilde{d}}$	.108	.098
$2 \times$ LogLikel	779.05	762.55

A. Parameter values

HST calibrated a $\sigma = 0$ version of their permanent income model by maximizing a likelihood function conditioned only on U.S. quarterly consumption and investment data. They used U.S. quarterly data on consumption and investment for the period 1970I–1996III. They measured consumption by nondurables plus services and investment by the sum of durable consumption and gross private investment.[28] They estimated the model from data on (c_t, i_t), setting $\sigma = 0$, then deduced pairs (σ, β) that are observationally equivalent, using formula (10.3.18).

The forcing processes are governed by seven free parameters: $(\alpha_1, \alpha_2, c_{\hat{d}}, \phi_1, \phi_2, c_{\tilde{d}}, \mu_d)$. The parameter μ_b sets a bliss point. While μ_b alters the marginal utilities, it does not influence the decision rules for consumption and investment. HST fixed μ_b at an arbitrary number, namely 32, for estimation.

Four parameters govern the endogenous dynamics: $(\gamma, \delta_h, \beta, \lambda)$. HST set $\delta_k = .975$, and imposed the permanent-income restriction, $\beta R = 1$. The restrictions that $\beta R = 1, \delta_k = .975$ pin down γ once β is estimated. HST imposed $\beta = .9971$, which after adjustment for the effects of the geometric growth factor of 1.0033 implies an annual real interest rate of 2.5%.

Table 10.A.1 reports HST's estimates for the parameters governing the endogenous and exogenous dynamics. Figures 10.A.1 and 10.A.2 report impulse response functions for consumption and investment to innovations in both components of the endowment process. For comparison, table 10.A.1 reports estimates from a no habit persistence ($\lambda = 0$) model as well.

Notice that the persistent endowment shock process contributes much more to consumption and investment fluctuations than does the transitory endowment shock process.

[28] They estimated the model from data that had been scaled through multiplication by 1.0033^{-t}.

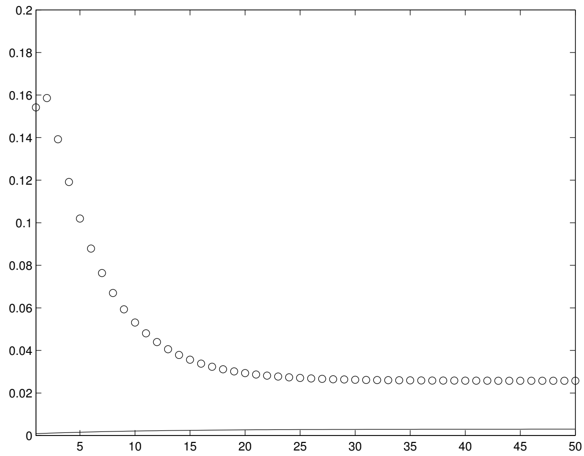

Figure 10.A.1: Impulse response functions of investment (circles) and consumption (solid line) to innovation in transitory endowment process ($\hat{d}$), at maximum likelihood estimate of habit persistence.

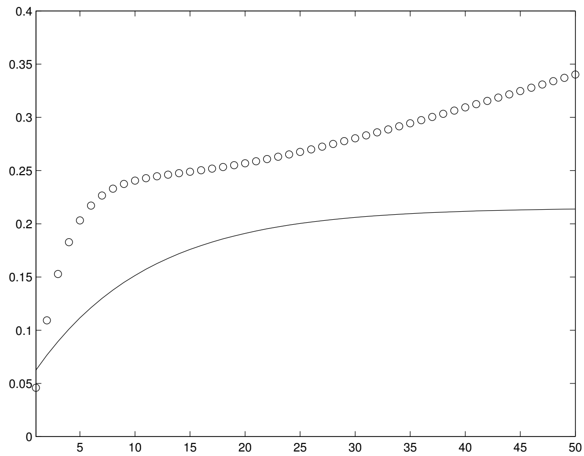

Figure 10.A.2: Impulse response functions of investment (circles) and consumption (solid line) to innovation in persistent shock ($\tilde{d}$), at maximum likelihood estimate of habit persistence.

B. Another observational equivalence result

To shed more light on the form of precautionary savings, we state another observational equivalence result that takes as its benchmark an initial allocation associated with parameter settings $\beta R = 1$ and $\sigma < 0$. Then we find another value of β that implies the same decisions for c_t, i_t as the base model when $\sigma = 0$, so that the decision maker fears model misspecification. This entails working backwards from the worst-case model that is reflected in the $\sigma < 0$ decision rule to the associated approximating model.

Theorem 10.B.1. *(Observational Equivalence, II) Fix all parameters except* (σ, β). *Consider a consumption-investment allocation for* $(\hat{\sigma}, \hat{\beta})$ *where* $\hat{\beta}$ *satisfies* $\hat{\beta}R = 1$ *and* $\hat{\sigma} < 0$ *and* $\underline{\hat{\sigma}} < \hat{\sigma}$. *Then there exists a* $\tilde{\beta} > \hat{\beta}$ *such that the* $(\hat{\sigma}, \hat{\beta})$ *allocation also solves the* $(0, \tilde{\beta})$ *problem.*

Proof. We suppose that $\hat{\sigma} < 0$, so that the worst-case model differs from the approximating model. We want to find the approximating model and a value $\tilde{\beta}$ of β for which a $\sigma = 0$ decision maker would choose the $(\hat{\sigma}, \hat{\beta})$ allocation. Under the model with $\hat{\sigma} < 0$, where $\hat{E}_t$ denotes a conditional expectation under the worst-case model, we have

$$\hat{E}_t \mu_{c,t+1} = \mu_{c,t} \tag{10.B.1}$$

because $\hat{\beta}R = 1$. Let

$$\hat{E}_t \mu_{s,t+1} = \xi\left(\tilde{\beta}\right) \mu_{s,t}. \tag{10.B.2}$$

Equation (10.B.1) implies that we want

$$1 = \xi\left(\tilde{\beta}\right) \tag{10.B.3}$$

where the projection coefficient $\xi(\tilde{\beta})$ emerges from the multiplier problem for the evil agent for $\hat{\sigma} < 0$, which can be cast as

$$\min_{\{w_{t+1}\}} \left[-\sum_{t=0}^{\infty} \hat{\beta}^t \{\mu_{st}^2 + \hat{\beta}\frac{1}{\hat{\sigma}} w_{t+1}^2\} \right]$$

subject to the law of motion

$$\mu_{st} = \delta\left(\tilde{\beta}\right) \mu_{s,t-1} + \alpha w_t \tag{10.B.4}$$

where $\delta(\tilde{\beta}) = \frac{1}{\tilde{\beta}R}$ and α is given by (10.3.17), (10.3.14), (10.3.15) under the $(\hat{\sigma}, \hat{\beta})$ model. (Remember that the decision rule for c_t and therefore the law for μ_{st} will be the same under our two observationally equivalent (σ, β) pairs, so we can use the benchmark case to compute α.) We freeze all parameters except σ, β. The approximating model would be $\mu_{st} = \delta\mu_{s,t-1} + \alpha\epsilon_t$, so that (10.$B$.4) adds a perturbation αw_t to the law of motion of μ_{st} under a deterministic version of the approximating model. The Bellman equation for the minimizing agent is evidently

$$-P\mu_s^2 = -\mu_s^2 + \hat{\beta}\min_w \left[-\frac{1}{\hat{\sigma}} w^2 - P\left(\delta\mu_s + \alpha w\right)^2 \right]. \tag{10.B.5}$$

Notice the presence of both $\hat{\beta}$ and $\tilde{\beta}$, via δ and α. The first-order condition is

$$w = K\mu_s,$$

where

$$K = -\frac{\alpha\delta\hat{\sigma}P}{1 + \alpha^2\hat{\sigma}P}.$$

Notice that

$$\xi\left(\tilde{\beta}\right) = A + KC = \delta + K\alpha = 1,$$

which implies that

$$1 = \xi\left(\tilde{\beta}\right) = \delta + K\alpha = \frac{\delta}{1 + \alpha^2 \hat{\sigma} P}.$$

Therefore,

$$\delta = 1 + \hat{\sigma}\alpha^2 P < 1. \tag{10.B.6}$$

Equation $(10.B.5)$ implies that

$$-P = -1 + \hat{\beta}\left[-\frac{1}{\hat{\sigma}}K^2 - P\left(\delta + K\alpha\right)^2\right].$$

Simplifying the above identity leaves

$$P = \frac{1}{1 - \hat{\beta}}\left[1 + \frac{\hat{\beta}}{\hat{\sigma}}\left(\frac{1 - \delta}{\alpha}\right)^2\right]. \tag{10.B.7}$$

Equations $(10.B.6)$ and $(10.B.7)$ together imply that

$$0 = \hat{\beta}\left(1 - \delta\left(\tilde{\beta}\right)\right)^2 + \left(1 - \hat{\beta}\right)\left(1 - \delta\left(\tilde{\beta}\right)\right) + \alpha\left(\tilde{\beta}\right)^2 \hat{\sigma}.$$

A solution of this equation determines $\tilde{\beta}$. The solution of this quadratic equation is

$$\delta = 1 - \frac{-\left(1 - \hat{\beta}\right) \pm \sqrt{\left(1 - \hat{\beta}\right)^2 - 4\hat{\beta}\sigma\alpha^2}}{2\hat{\beta}}.$$

If $\sigma = 0$, this equation implies $\delta = 1$. When $\sigma < 0$, the appropriate root is

$$\delta = 1 - \frac{-\left(1 - \hat{\beta}\right) + \sqrt{\left(1 - \hat{\beta}\right)^2 - 4\hat{\beta}\sigma\alpha^2}}{2\hat{\beta}}.$$

Using $\hat{\beta}R = 1$, this is equivalent to

$$\tilde{\beta}\left(\sigma\right) = \frac{\hat{\beta}\left(1 + \hat{\beta}\right)}{2\left(1 + \sigma\alpha^2\right)}\left[1 + \sqrt{1 - 4\hat{\beta}\frac{1 + \sigma\alpha^2}{\left(1 + \hat{\beta}\right)^2}}\right]. \tag{10.B.8}$$

$\blacksquare$