“Certainty Equivalence” and “Model Uncertainty”

Lars Peter Hansen
University of Chicago

and

Thomas J. Sargent
New York University and
Hoover Institution
Stanford University

25 March 2004

ABSTRACT

Simon’s and Theil’s certainty equivalence property justifies a convenient algorithm for solving dynamic programming problems with quadratic objectives and linear transition laws: first, optimize under perfect foresight, then substitute optimal forecasts for unknown future values. A similar decomposition into separate optimization and forecasting steps prevails when a decision maker wants a decision rule that is robust to model misspecification. Concerns about model misspecification leave the first step of the algorithm intact and affect only the second step of forecasting the future. The decision maker attains robustness by making forecasts with a distorted model that twists probabilities relative to his approximating model. The appropriate twisting emerges from a two-player zero-sum dynamic game.

Note: This paper was prepared for a conference at the Federal Reserve Board on March 26–27, 2004 to honor the work of our friends Dale Henderson, Richard Porter, and Peter Tinsley. An earlier version of this paper was presented at a conference to honor the memory of Henri Theil.

1. Certainty equivalence and rational expectations

Lucas and Sargent (1981) attributed to Simon's (1956) and Theil's (1957) certainty equivalence principle an important role in developing applied dynamic rational expectations models. Two of the three examples in Lucas's Critique (1976) and all but one of the papers in Lucas and Sargent (1981) assumed environments for which certainty equivalence prevails. By sharply delineating the two steps of (1) optimizing for a given set of expectations and (2) forming expectations optimally, these problems formed a perfect environment for extracting the methodological and econometric lessons of rational expectations.[1] Two of the most important were: (a) how rational expectations imposes a set of cross-equation restrictions that link the parameters of an optimal decision rule to laws of motion for variables like prices that influence a decision maker's payoffs, but that are beyond his control;[2] and (b) how the concept of Granger-causality, based as it is on a prediction-error criterion, can guide the empirical specification of variables that belong on the right side of decision rule because they help predict those influential variables. Environments for which certainty equivalence holds are ones for which it is easiest to compute decision rules analytically. That has facilitated formal analysis as well as numerical computation.[3]

for Finally, for some important applications, certainty equivalent environments are benchmarks against which departures have been measured.[4]

[1] The literatures on applied dynamic economics and macroeconomics before Muth were ideally set up for application of the certainty equivalence principle. The standard practice then was to apply nonstochastic optimization problems, then to supplement the solution with a theory about expectations.

[2] This *is* the Lucas critique.

[3] By using linear approximations to decision rules, many papers in the real business cycle literature have approximated the solutions to optimization problems for which certainty equivalence does not obtain with related problems in which it does. Schmitt-Grohe and Uribe (2004) study second order approximations to decision rules and discover that for the second order approximants, constants are affected by volatility terms but the linear and quadratic pieces are not.

[4] A leading example occurs in the precautionary savings literature, which by perturbing a linear quadratic benchmark alters the pure 'permanent income' model of consumption to induce an extra source of saving from the extra curvature of the marginal utility of consumption.

2. Model uncertainty

As in a rational expectations model or the subgame perfect equilibrium of a game, the decision maker envisioned in Simon's and Theil's analysis experiences no uncertainty about the specification of his dynamic *model*. He knows the model up to the realization of a Gaussian random disturbance. But in practice a decision maker could find himself knowing less, like econometricians whose finite data sets expose them to doubts provoked by their statistical specification tests.

Diverse literatures on 'uncertainty aversion' and 'robustness' consider decision makers who do not 'know the model'. They make decisions at best knowing a *set* of models. Such agents are inspired either by the axioms of Gilboa and Schmeidler (1989) or their desire for a decision rule that is robust to misspecified model dynamics (e.g., Zhou, Glover, and Doyle (1996)) to choose a decision rule from a systematic worst-model analysis of alternative rules. These literatures represent model uncertainty by having a decision maker retain a set of models that he refuses to reduce to one by the Bayesian device of assigning probabilities over models in the set. This non-Bayesian decision maker behaves as someone who maximizes expected utility while assuming that, given his decision rule, a perverse nature chooses the worst from within his set of models. The literature on robustness shows how such min-max behavior promotes a decision rule that is robust to model misspecification.

The ignorance that is called uncertainty in these literatures on uncertainty aversion and robustness differs from not knowing realizations from a known probability distribution, the subject of Simon's and Theil's certainty equivalence. Nevertheless, a remarkable and useful version of certainty equivalence prevails when a decision maker expresses his fear of model misspecification by wanting good decisions across a set of models near his approximating model.[5] The structure of this modified certainty equivalence result resembles that for ordinary certainty equivalence. It too sharply divides decision making into separate phases of first optimizing given beliefs, then forming beliefs. The first stage turns out to be identical to that for Simon's and Theil's setting. But in the second stage, the decision maker distorts beliefs relative to his approximating model for the purpose of achieving

[5] This insight began with Jacobson (1973) and was developed by Whittle (1982, 1999), Başar and Bernard (1995) for undiscounted problems. Hansen and Sargent (1995, 2005) formulated discounted problems recursively.

robustness. The resulting form of certainty equivalence preserves many of the analytical conveniences of Simon's and Theil's result. It also sheds light on a kind of precaution that is induced by a concern about model misspecification. In addition, it provides underpinnings for a Bayesian rationalization of the decision made by a robust decision maker. The certainty equivalence representation exposes a distorted law of motion in terms of which the robust decision maker behaves as though he were a Bayesian.

3. Ordinary certainty equivalence

3.1. Notation and setup

For any vector y, let $\mathbf{y}_t^s$ denote the history from t to s. If a subscript is omitted, we take it to be zero. If a superscript is omitted, we take it to be $+\infty$. Thus, $\mathbf{y}^t$ is the history from 0 to t, and $\mathbf{y}_t$ is the future from t to $+\infty$.

An exogenous component z_t of a state vector $y_t = \begin{bmatrix} x_t \\ z_t \end{bmatrix}$ has transition law

$$z_{t+1} = f\left(z_t, \epsilon_{t+1}\right) \tag{3.1}$$

where ϵ_{t+1} is an i.i.d. sequence of random vectors with cumulative distribution function Φ. The endogenous component x_t is influenced by the decision maker's control u_t via a transition law

$$x_{t+1} = g\left(x_t, z_t, u_t\right). \tag{3.2}$$

The decision maker evaluates stochastic processes $\{x_t, z_t, u_t\}_{t=0}^{\infty}$ according to

$$E\left[\sum_{t=0}^{\infty} \beta^t r\left(y_t, u_t\right) | \mathbf{y}^0\right] \tag{3.3}$$

where $\beta \in (0,1)$, $E(\cdot)|\mathbf{y}^t$ is the mathematical expectation conditioned on $\mathbf{y}^t \equiv (\mathbf{x}^t, \mathbf{z}^t)$, and u_t is required to be a measurable function of $\mathbf{y}^t$. The decision maker chooses $\mathbf{u}_t$ to maximize (3.3) subject to (3.1) and (3.2). The solution is a decision rule

$$u_t = h\left(x_t, z_t\right). \tag{3.4}$$

Along with Simon (1956) and Theil (1957), throughout this paper we use

Assumption 1: The one-period return function r is quadratic, e.g., $r(y,u) = -y'Qy - u'Ru$ where Q, R are both positive semidefinite matrices; f and g are both linear; and Φ is multivariate Gaussian with mean zero and identity covariance matrix.

Simon and Theil showed that under Assumption 1 the solution of the stochastic optimization problem of maximizing (3.3) subject to (3.2) and (3.1) has a special structure. In particular, the problem can be separated into two parts (notice that only (3.3) and (3.2) appear in part 1, while only (3.1) is used in part 2):

1. Solve the nonstochastic or perfect-foresight problem of maximizing (3.3) subject to (3.2), assuming that the future sequence $\mathbf{z}_t$ is known. This leads to a 'feedback-feedforward solution'

$$u_t = h_1\left(x_t, \mathbf{z}_t\right). \tag{3.5}$$

The decision for u_t feeds back on the endogenous state vector x_t and feeds forward on the future of the exogenous component of the state vector $\mathbf{z}_t$. The function h_1 can be computed from knowing only V and g, and without knowing f. In particular, h_1 is obtained by deducing the decision maker's Euler equation for u_t, then solving it forward. The parts of assumption A1 stating that V is quadratic and g is linear make h_1 a linear function.

2. Using the function f and the c.d.f. Φ in (3.1), compute the mathematical expectation of $\mathbf{z}_t$ conditioned on the history $\mathbf{z}^t$. Because f is linear, by iterating on (3.1), the future sequence $\mathbf{z}_t$ can be expressed as a linear function of an initial condition, z_t, and the sequence of future shocks:

$$\mathbf{z}_t = h_2 \cdot z_t + h_3 \cdot \epsilon_{t+1}^{\infty}. \tag{3.6}$$

Then from the assumed properties of the i.i.d. sequence $\{\epsilon_{t+1}\}$, the solution of the forecasting problem takes the form

$$E\mathbf{z}_t|\mathbf{z}^t = h_2 \cdot z_t. \tag{3.7}$$

The *certainty equivalence* or *separation* principle states that the optimal decision rule in pure feedback form can be obtained by replacing $\mathbf{z}_t$ in (3.5) with $E\mathbf{z}_t|\mathbf{z}^t$ from (3.7):

$$u_t = h_1\left(x_t, h_2 \cdot z_t\right) = h\left(x_t, z_t\right). \tag{3.8}$$

Each of h_1, h_2, and h are linear functions. The decision rule (3.8) feeds back on both exogenous and endogenous components of the state.

The original stochastic control problem thus *separates* into a nonstochastic control problem and a statistical estimation problem. An inspiration for the term 'certainty equivalence' is that the control problem in step (1) can be solved by assuming that the future sequence of z's, $\mathbf{z}_t$, is known.

3.2. Certainty equivalence and the value function

Evidently from (3.6),

$$E\left[\left(\mathbf{z}_t - E\mathbf{z}_t|\mathbf{z}^t\right)\left(\mathbf{z}_t - E\mathbf{z}_t|\mathbf{z}^t\right)'|\mathbf{z}^t\right] = h_3 h_3'. \tag{3.9}$$

The optimized value of (3.3) starting from state y_0 is given by a value function

$$V(y_0) = -y_0' P y_0 - p. \tag{3.10}$$

The constant p depends on the 'volatility statistics' $h_3 h_3'$ in (3.9), but the matrix P in the quadratic form in (3.10) does not. In particular, the matrix P is the fixed point of an operator

$$T(P) = T(P; r, g, f_1), \tag{3.11}$$

where we express the linear law of motion for z_t as $z_{t+1} = f_1 z_t + f_2 \epsilon_{t+1}$. The volatility parameters f_2 do not appear in T. The constant p *does* depend on f_2 via the operator

$$p = p(P; f_2, \beta). \tag{3.12}$$

Note that the fixed point $P = T(P)$ can be solved first, then the solution can be put into (3.12) to find p.

This last observation leads to another manifestation of the certainty equivalence principle that comes from noting that the optimal decision rule (3.8) for the stochastic optimization problem also emerges from the *nonstochastic* problem of maximizing

$$\sum_{t=0}^{\infty} \beta^t r(x_t, z_t, u_t) \tag{3.13}$$

subject to (3.2) and the *nonstochastic* law of motion

$$z_{t+1} = f(z_t, 0). \tag{3.14}$$

Equation (3.14) sets all of the shocks impinging on the state to their unconditional means of zero. For this problem, the optimal value function becomes

$$V(y_0) = -y_0' P y_0, \tag{3.15}$$

which differs from (3.10) only by the absence of the constant p. The decision rule depends only on the quadratic forms in the continuation value functions (3.10) and (3.10), not on the constants. Therefore, the presence of uncertainty in the original problem lowers the *value* of the problem by lowering p, but does not affect the decision rule. The decision maker prefers less uncertainty to more, i.e., he prefers a f_2 and therefore a smaller h_3 loading in (3.6), but does not allow his decisions to respond to different values of h_3.

4. Model misspecification and robustness

We now turn to a type of certainty equivalence that prevails when the decision maker does not trust his model of the dynamics for z_t. This certainty equivalence principle lets the decision maker use the type of two-step optimization process described above and to express his doubts about the transition law in the way he forms expectations in the second step.

The decision maker fears that (3.1) is misspecified because the data might actually be generated by the law of motion

$$z_{t+1} = f(z_t, \epsilon_{t+1} + w_{t+1}) \tag{4.1}$$

where w_{t+1} is a process whose time $t+1$ component is a measurable function ω_t of the history of state at t

$$w_{t+1} = \omega_t\left(\mathbf{x}^t, \mathbf{z}^t\right).$$

The decision maker thinks that his model (3.1) is a good approximation to the data generating mechanism (4.1) in the sense that

$$\hat{E}\left[\sum_{t=0}^{\infty} \beta^t w_{t+1}' w_{t+1} \middle| y_0 \leq \eta_0\right] \tag{4.2}$$

where η_0 measures the size of the maximal specification error and the expectation operator $\hat{E}$ is evaluated with respect to the distribution generated by (4.1).[6]

[6] See Anderson, Hansen, and Sargent (2003) and Hansen and Sargent (2005) for a reasonable way to measure discrepancies between models.

To construct a robust decision rule, the decision maker computes the Markov perfect equilibrium of the following two-player zero-sum game that we call a *multiplier problem*:

$$\min_{\mathbf{w}_1} \max_{\mathbf{u}_0} \hat{E} \left[\sum_{t=0}^{\infty} \beta^t \left\{ r\left(y_t, u_t\right) + \theta \beta w_{t+1}' w_{t+1} \right\} \right] \Bigg| y_0 \tag{4.3}$$

where the components of u_t, w_{t+1} must each be measurable functions of the time t histories $\mathbf{y}^t$. In the Markov perfect equilibrium, the timing protocol is that each player chooses sequentially and simultaneously each period, taking the other player's decision rule as given. The worst case shock mean distortions w_{t+1} can feed back on the endogenous state vector x_t. Allowing this feedback is part of the way that the maximizing agent designs a rule that is robust to the possibility that the dynamics of (3.1) are misspecified.

A robust rule is generated by the u component of a Markov perfect equilibrium of game (4.3):

$$u_t = h\left(x_t, z_t\right) \tag{4.4a}$$

$$w_{t+1} = W\left(x_t, z_t\right). \tag{4.4b}$$

The rule is *robust* in the sense that it promises a lower maximal rate at which the objective $\hat{E}[\sum_{t=0}^{\infty} \beta^t r(y_t, u_t)]|y_0$ can deteriorate with increases in misspecification as measured by the terms $w_{t+1}' w_{t+1}$.

4.1. Stackelberg equilibrium and certainty equivalence

The Markov perfect equilibrium conceals that a form of certainty equivalence prevails despite the decision maker's uncertainty about his approximating model. To reveal the certainty equivalence within the robust decision rule (4.4*b*), it helps to formulate another game with the same players and payoffs, but a different timing protocol. Remarkably, this change in timing protocol leaves intact both the equilibrium outcome and its recursive representation (4.4), a consequence of the special zero-sum feature of the dynamic game (see Başar and Bernhard (1995) and Hansen and Sargent (2005)).

We now impose the following timing protocol on a two-player zero-sum game with transition laws (3.2), (4.1) and the payoff (4.3). At time 0, the minimizing player once-and-for-all chooses a plan $\mathbf{w}_1$, where w_{t+1} is a measurable function of information known to him at time t, which we denote $\mathbf{Y}^t$. Given the random sequence $\mathbf{w}_1$, the maximizing player chooses u_t each period. This timing protocol makes the minimizing player the Stackelberg leader. At time 0, the minimizing player chooses $\mathbf{w}_1$, taking into account the best response of the maximizing player, who chooses sequentially and regards $\{w_{t+1}\}$ as an exogenous process.

4.2. Representing the Stackelberg timing protocol

To reflect the Stackelberg timing protocol within a recursive representation, we add state variables to describe the decisions to which the minimizing player commits at time 0. We use upper case counterparts of x_t and z_t to denote these additional state variables, all of which the maximizing player takes as exogenous. Each of the vectors X, Y, Z has the same dimension as its lower case counterpart. At time 0, the Stackelberg leader makes w_{t+1} follow the law of motion:

$$w_{t+1} = W\left(X_t, Z_t\right) \tag{4.5a}$$

$$X_{t+1} = g\left(X_t, Z_t, h\left(X_t, Z_t\right)\right) \tag{4.5b}$$

$$Z_{t+1} = f\left(Z_t, W\left(X_t, Z_t\right) + \epsilon_{t+1}\right) \tag{4.5c}$$

which we summarize as

$$w_{t+1} = W\left(Y_t\right) \tag{4.6a}$$

$$Y_{t+1} = MY_t + N\epsilon_{t+1}. \tag{4.6b}$$

Note the appearance in (4.5b) of the maximizing player's best response function $h(X, Z)$ from (4.4b). Also W in (4.5a) is the *same* best response function that appears in (4.4b). Thus, equations (4.5b) and (4.5c) are the 'big letter' counterparts to (4.1) and (3.2), where we have substituted 'big letter' analogues of the best response functions (4.4) for the big letter counterparts to u_t and w_{t+1}.

The maximizing agent takes (4.6) as given and forecasts z_t according to the distorted law of motion (4.1), (4.6), where the role of the latter system of equations is to describe the distortion process w_{t+1}. Thus, given the time 0 choice of the minimizing player, the maximizing player assumes the law of model for z_t to be

$$Y_{t+1} = MY_t + N\epsilon_{t+1} \tag{4.7a}$$

$$w_{t+1} = W(Y_t) \tag{4.7b}$$

$$z_{t+1} = f(z_t, \epsilon_{t+1} + w_{t+1}). \tag{4.7c}$$

The maximizing player faces an ordinary dynamic programming problem of maximizing (3.3) subject to (3.2) and (4.7), and so chooses a decision rule of the form

$$u_t = \tilde{H}(x_t, z_t, Y_t). \tag{4.8}$$

Under the Stackelberg timing protocol, we use the following

Definition: An *equilibrium* of the Stackelberg game is a collection of functions $h, W, \tilde{H}$ such that: (a) Given $\tilde{H}$, the decision rule (4.6) solves the minimizing player's problem; (b) Given (4.6), decision rule (4.8) solves the maximizing player's problem; (c) $\tilde{H}(X_t, Z_t, Y_t) = h(X_t, Y_t)$; (d) $Y_0 = y_0$.

Conditions (c) and (d) impose versions of what in macroeconomics are called the 'big K equals little k' conditions.

Başar and Bernhard (1995) and Hansen and Sargent (2004) prove that

$$\tilde{H}(X_t, Z_t, Y_t) = h(Y_t). \tag{4.9}$$

This result asserts that the same decision rule for u_t emerges from the Markov perfect and the Stackelberg games. It rationalizes an interpretation of robust decision rule $h(Y_t)$ in terms of a modified type of certainty equivalence.

4.3. The Stackelberg game and certainty equivalence

With the Stackelberg timing protocol, we can characterize the maximizing player's choice with an Euler equation that we can solve forward to get his time t decision as a function of future actions of the minimizing player. This legitimizes the following two-step solution procedure.

1. Solve the same non-stochastic problem as in step (1) above, assuming again that the sequence $\mathbf{z}_t$ is known, leading to the same solution (3.5):

$$u_t = h_1(x_t, \mathbf{z}_t).$$

2. Using the distorted law of motion (4.7), compute the mathematical expectation of $\mathbf{z}_t$ conditioned on the history $\mathbf{z}^t$. The Markov property of (4.7), the linearity of f, and the Gaussian distribution of ϵ_{t+1}^{∞} make the solution of this forecasting problem take the form

$$\hat{E}\mathbf{z}_t|\mathbf{z}^t = \hat{h}_2 \cdot \begin{bmatrix} z_t \\ Y_t \end{bmatrix} \tag{4.10}$$

where $\hat{E}$ denotes the conditional expectation using model (4.7).

A modified separation principle states that the robust decision rule in pure feedback form can be obtained by replacing $\mathbf{z}_t$ in (3.5) with $\hat{E}[\mathbf{z}_t|\mathbf{z}^t, \mathbf{Y}^t]$ from (4.10), and then setting $Y_t = y_t$.[7] This yields the robust decision rule

$$u_t = h_1 \left(x_t, \hat{h}_2 \cdot y_t \right). \tag{4.11}$$

Our certainty equivalence result under robustness asserts that (4.11) is equivalent with the robust rule (4.4a), the u component of the Markov perfect equilibrium of game (4.3), so that $u_t = h_1(x_t, \hat{h}_2 \cdot y_t) = h(x_t, z_t)$.

Another view of this certainty equivalence result come from noting that there is a counterpart to our finding without robustness that the same decision rule (but a different optimal value function) emerges had we just set $\{\epsilon_t\}_{t=1}^{\infty}$ to zero. That is, the same decision rule $h(x_t, z_t)$ would emerge from our Stackelberg problem had we solved the *nonstochastic* control problem that emerges upon replacing (4.7) with the nonstochastic law of motion

$$\begin{aligned} Y_{t+1} &= MY_t \\ w_{t+1} &= W\left(Y_t\right) \\ z_{t+1} &= f\left(z_t, w_{t+1}\right), \end{aligned}$$

which we form by setting $\epsilon_{t+1} \equiv 0$ in (4.7).

[7] This is the stage at which we set 'big K' equal to 'little k'.

5. Other uses of the Stackelberg timing protocol

In addition to clarifying the certainty equivalence embedded in the robust decision rule, the 'big letter' law of motion associated with the Stackelberg timing protocol has other important uses. These include:

- Supporting a recursive representation of state-contingent prices for a decentralized version of an economy with a robust representative agent. These prices can be expressed as linear functions of a state Y_t that is described by a version of (4.5). See Anderson, Hansen, and Sargent (2003) and Hansen and Sargent (2005) for details.
- Supplying a Bayesian interpretation of a robust decision maker. Equation (4.7) provides a unique model that rationalizes the robust decision maker's choices. See Blackwell and Girshik (1954) for related material about interpreting the outcomes of zero sum games.

6. Certainty equivalence and the value function

Under a preference for robustness, the optimized value of (4.3) starting from state y_0 is given by a value function

$$V(y_0) = -y_0' P y_0 - p \tag{6.1}$$

but now the constant p and the matrix P both depend on the volatility parameters f_2 in the representation

$$z_{t+1} = f_1 z_t + f_2 (\epsilon_{t+1} + w_{t+1}). \tag{6.2}$$

In particular, the matrix P is now the fixed point of a composite operator $T \circ \mathcal{D}$ where T is the same operator defined in (3.11) and $\mathcal{D}$ is an operator

$$\mathcal{D}(P) = \mathcal{D}(P; f_2, \theta), \tag{6.3}$$

that depends on f_2 and θ. Given a fixed point P of $T \circ \mathcal{D}$, p is then a function

$$p = p(P; f_2, \beta, \theta). \tag{6.4}$$

The $\mathcal{D}$ operator represents the distortion through w imposed by the minimizing player, while the T operator represents the maximizing player's choice of u. Because P is a fixed point of $T \circ \mathcal{D}$, the volatility parameters f_2 now influence P and therefore the decision rule h. Nevertheless, a form of certain equivalence prevails.

Thus, paralleling our earlier discussion without robustness, another manifestation of the certainty equivalence principle under robustness comes from noting that identical decision rules (4.4) emerge from the *nonstochastic* zero-sum two-player game associated with extremizing (i.e., jointly minimizing and maximizing) (4.3) subject to (3.2) and the *nonstochastic* law of motion

$$z_{t+1} = f(z_t, w_{t+1}) = f_1 z_t + f_2 w_{t+1}. \tag{6.5}$$

Equation (6.5) sets all of the shocks impinging on the state to their unconditional means of w_{t+1}. For this problem, the equilibrium value function becomes

$$V(y_0) = -y_0' P y_0, \tag{6.6}$$

which differs from (6.1) only by the absence of the constant p. The decision rule depends only on the quadratic forms in the continuation value functions (3.10) and (3.10), not on the constants. Therefore, the presence of randomness lowers the *value* of the game by lowering p, but does not affect the decision rule.

7. Risk sensitive preferences

Building on results of Jacobson (1973) and Whittle (1990), Hansen and Sargent (1995) showed that another way to interpret the decision rules that we obtain under a preference for robustness is to regard them as reflecting a particular version of Epstein and Zin's (1989) specification of recursive preferences. Thus, suppose now that the decision maker believes his model and so has no concern about model misspecification. Let the model be

$$y_{t+1} = Ay_t + Bu_t + C\epsilon_{t+1} \tag{7.1}$$

where ϵ_{t+1} is an i.i.d. Gaussian process with contemporaneous mean zero and identity covariance matrix. Hansen and Sargent suppose that the decision maker evaluates processes $\{y_t, u_t\}_{t=0}^{\infty}$ according to a utility function U_0 that is defined by the fixed point of recursions on

$$U_t = r(y_t, u_t) + \beta \mathcal{R}_t (U_{t+1}) \tag{7.2}$$

where

$$\mathcal{R}_t\left(U_{t+1}\right) = \left(\frac{2}{\sigma}\right) \log E\left[\exp\left(\frac{\sigma U_{t+1}}{2}\right) \Big| \mathbf{y^t}\right] \tag{7.3}$$

and $\sigma \leq 0$ is the risk sensitivity parameter. When $\sigma = 0$, an application of l'hospital's rule shows that the right side of (7.2) becomes $r(y_t, u_t) + \beta E U_{t+1} | \mathbf{y^t}$. However, when $\sigma < 0$, $\mathcal{R}_t$ departs from the ordinary conditional expectation operator and puts an additional correction for risk into the evaluation of continuation utility U_{t+1}. The risk-sensitive control problem is to choose a decision rule $u_t = -F x_t$ that maximizes U_0 defined by recursions on (7.2), subject to (7.1).

Despite their different motivations, a risk-sensitive control problem yields precisely the same decision rule as a corresponding robust control problem with $\theta = -\sigma^{-1}$. This useful fact can be established by following Jacobson (1973) and evaluating $\mathcal{R}_t(U^e_{t+1})$ for a candidate quadratic continuation value function

$$U^e_{t+1} = -y'_{t+1}\Omega y_{t+1} - \rho \tag{7.4}$$

and a time t decision rule $u_t = -F_t x_t$, so that (7.1) becomes

$$y_{t+1} = \hat{A}_t y_t + C\epsilon_{t+1} \tag{7.5}$$

where $\hat{A}_t = A - BF_t$. Using the properties of the Gaussian distribution, one obtains

$$\mathcal{R}_t U^e_{t+1} = -y_t \hat{A}'_t \mathcal{D}\left(\Omega\right) \hat{A}_t y_t - \hat{\rho}, \tag{7.6}$$

where $\mathcal{D}$ is the *same* operator defined in (6.3) with $\theta = -\sigma^{-1}$ and[8]

$$\hat{\rho} = \rho - \sigma^{-1} \log\det\left(I + \sigma C'\Omega C\right). \tag{7.7}$$

If we guess that the value function is quadratic, the Bellman equation for a risk-sensitive control problem is[9]

$$-y'Py - p = \max_u \left[r\left(y,u\right) + \beta\mathcal{R}_t\left(-y^{*\prime}Py^* - p\right)\right] \tag{7.8}$$

[8] In particular,

$$\mathcal{D}\left(\Omega\right) = \left[I - \sigma\Omega C\left(I + \sigma C'\Omega C\right)^{-1} C'\right]\Omega.$$

In the shorthand notation used in (6.3), f_2 plays the role of C.

[9] The P that solves (7.7) matches the P in the value function (6.1) for the robust control problem, but the constants p differ. The decision rules depend on P but not on the constants.

where the maximization is subject to $y^* = Ay + Bu + C\epsilon$. It follows that P is the fixed point of iterations on $T \circ \mathcal{D}$, where T is the operator defined in (3.11) that is associated with the ordinary quadratic intertermporal optimization problem without risk-sensitivity and without a preference for robustness. Since the matrices in the quadratic forms in the value functions are the same, both being fixed points of $T \circ \mathcal{D}$, the decision rule for u_t is the same under risk sensitivity or a preference for robustness. Under risk-sensitive preferences, the $\mathcal{D}$ operator comes from evaluating the $\mathcal{R}$ operator, while the T operator comes one-step on the ordinary Bellman equation, taking $\mathcal{R}_t(-y^{*\prime} P y^* - p)$ as the continuation value function.

It immediately follows from these results that with quadratic $r(y, u)$, a form of certainty equivalence holds for the risk-sensitive preference specification.

8. Example: robustness and discounting in a permanent income model

This section illustrates certainty equivalence in the context of a linear-quadratic version of a simple permanent income model.[10] In the basic permanent income model, a consumer applies a single marginal propensity to consume to the sum of his financial wealth and his 'human wealth', defined as the expected present value of his labor (or endowment) income discounted at the same risk-free rate of return that he earns on his financial assets. In the usual permanent income model without a preference for robustness, the consumer has no doubts about the probability model used to form the conditional expectation of discounted future labor income. But with a preference for robustness, the consumer doubts that model and therefore forms forecasts of future income by using a probability distribution that is twisted or slanted relative to his approximating model for his endowment process. Except for this slanting, the consumer behaves as an ordinary permanent income consumer, i.e., uses the same component function h_1 from (3.5), a reflection of our certainty equivalence principle under robustness.

[10] See Sargent (1987) and Hansen, Roberds, and Sargent (1991) for accounts of the connection between the permanent income consumer and Barro's (1979) model of tax smoothing. See Aiyagari, Marcet, Sargent, and Seppälä (2002) for a more extensive exploration of the connections.

This slanting of probabilities leads the consumer to engage in a form of *precautionary savings* that under the approximating model tilts his consumption profile toward the future relative to what it would be without a preference for robustness. Indeed, so far as his consumption and saving program is concerned, activating a preference for robustness is equivalent with making the consumer more patient. However, that is not the end of the story because elsewhere (Hansen, Sargent, and Tallarini (1999) and Hansen, Sargent, and Wang (2002)) we have shown that attributing a preference for robustness to a representative consumer has different effects on asset prices than does varying his discount factor, even though it can leave a consumption-savings program unaltered.

8.1. The LQ permanent income model

In the linear-quadratic permanent income model (e.g., Hall (1978)), a consumer receives an exogenous endowment process $\{z_{2t}\}$ and wants to allocate it between consumption c_t and savings x_t to maximize

$$-E_0 \sum_{t=0}^{\infty} \beta^t \left(c_t - z_1\right)^2, \beta \in (0,1) \tag{8.1}$$

where E_0 denotes the mathematical expectation conditioned on time 0 information, and the constant z_1 is a bliss level of consumption. We simplify the problem by assuming that the endowment process is a first-order autoregression. Thus, the household faces the state transition laws

$$x_{t+1} + c_t = Rx_t + z_{2t} \tag{8.2a}$$

$$z_{2,t+1} = \mu_d \left(1 - \rho\right) + \rho z_{2t} + c_d \left(\epsilon_{t+1} + w_{t+1}\right), \tag{8.2b}$$

where $R > 1$ is a time-invariant gross rate of return on financial assets x_t held at the beginning of period t, and $|\rho| < 1$ describes the persistence of his endowment process. In (8.2b), w_{t+1} is a distortion to the mean of the endowment process that represents possible model misspecification. For convenience, we follow Whittle (1990) and use the *risk-sensitivity* parameter $\sigma = -\theta^{-1}$ to measure the consumer's preference for robustness. We begin by setting $\sigma = 0$ and solving the problem with no preference for robustness, then activate a preference for robustness by setting $\sigma < 0$ (i.e., $\theta < +\infty$). It is useful to let the consumer's choice variable be $\mu_{ct} = z_1 - c_t$, the marginal utility of consumption and to express (8.2a) as

$$x_{t+1} = Rx_t + (z_{2t} - z_1) + \mu_{ct}. \tag{8.3}$$

8.2. Solution when $\sigma = 0$

As promised, first we solve the household's problem *without* a preference for robustness, so that $\sigma = 0$. The household's Euler equation is

$$E_t \mu_{c,t+1} = (\beta R)^{-1} \mu_{ct}. \tag{8.4}$$

Treating (8.2a) as a difference equation in x_t, solving it forward in time, and taking conditional expectations on both sides gives

$$x_t = \sum_{j=0}^{\infty} R^{-(j+1)} E_t \left(c_{t+j} - z_{2,t+j} \right). \tag{8.5}$$

This equation imposes expected present value budget balance on the household. Solving (8.4) and (8.5) and using $\mu_{ct} = z_1 - c_t$ gives the following representation for μ_{ct}:

$$\mu_{ct} = -\left(1 - R^{-2}\beta^{-1}\right) \left(Rx_t + E_t \sum_{j=0}^{\infty} R^{-j} \left(z_{2,t+j} - z_1 \right) \right) \tag{8.6}$$

Note that because $\mu_{ct} = z_1 - c_t$, this equation is a version of (3.5), i.e., it is the household's Euler equation "solved forward", so that the right side of (8.6) portrays the function h_1 after $E\mathbf{z}_t|\mathbf{z}^t$ has replaced $\mathbf{z}_t$. Equations (8.4) and (8.6) can be used to deduce the following representation for μ_{ct}

$$\mu_{c,t+1} = (\beta R)^{-1} \mu_{c,t} + \nu \epsilon_{t+1} \tag{8.7}$$

where ν is a scalar.[11]

Given an initial condition $\mu_{c,0}$, equation (8.7) describes the consumer's optimal behavior. This initial condition is determined by solving (8.6) at $t = 0$. It is easy to use (8.6) to deduce an optimal consumption rule of the form

$$\mu_{ct} = h \begin{bmatrix} z_t \\ x_t \end{bmatrix} = hy_t, \tag{8.8}$$

which is a particular version of (3.4). In the case $\beta R = 1$ that was analyzed by Hall (1978), (8.7) implies that the marginal utility of consumption μ_{ct} is a martingale, which, because $\mu_{ct} = z_1 - c_t$, also implies that consumption is a martingale.

---

[11] See Hansen and Sargent (2005) for a way to compute ν from the linear regulator.

8.3. A robust consumption rule $\sigma < 0$

Under a preference for robustness, the decision rule for consumption is summarized by the following solved-forward Euler equation for μ_{ct}:

$$\mu_{ct} = -\left(1 - R^{-2}\beta^{-1}\right)\left(Rx_t + \hat{E}_t \sum_{j=0}^{\infty} R^{-j}\left(z_{2,t+j} - z_1\right)\right) \tag{8.9}$$

where $\hat{E}_t \equiv \hat{E}[\cdot|\mathbf{z}^t, \mathbf{Y}^t]$ is the distorted expectations operator associated with the law of motion chosen by the malevolent agent in a Stackelberg equilibrium with $\sigma < 0$. Note how (8.9) assumes the certainty-equivalent form $u_t = h_1(x_t, \hat{E}\mathbf{z}_t|\mathbf{z}^t, \mathbf{Y}^t)$ that underlies (4.11), where h_1 is the *same* function that pertains to the no-robustness $\sigma = 0$ problem. By substituting explicit formulas for the forecasts $\hat{E}[\mathbf{z_t}|\mathbf{z}^t, \mathbf{Y^t}]$ that appear in (8.9), we obtain a robust consumption rule in the feedback form (4.11).

To compute the distorted measure with respect to which $\hat{E}$ in (8.9) is to be computed, we pose the problem of the minimizing player in the Stackelberg game. It can be represented as

$$\min_{\{w_{t+1}\}} - \sum_{t=0}^{\infty} \hat{\beta}^t \left\{\mu_{ct}^2 + \hat{\beta}\sigma^{-1} w_{t+1}^2\right\} \tag{8.10}$$

subject to

$$\mu_{c,t+1} = \left(\tilde{\beta}R\right)^{-1} \mu_{c,t} + \nu w_{t+1}. \tag{8.11}$$

Equation (8.11), the consumption Euler equation of the maximizing player (the household), encodes the maximizing player's best response to the minimizing player's choice of a sequence $\{w_{t+1}\}$. Under the Stackelberg timing, the minimizing player commits to a sequence $\{w_{t+1}\}_{t=0}^{\infty}$ that the maximizing player takes as given. The minimizing player determines that sequence by solving (8.10), (8.11). The worst case shock that emerges from this problem can be portrayed as a feedback rule $w_{t+1} = k\mu_{ct}$. Since from (8.8) $\mu_{ct} = hy_t$ in the Markov perfect equilibrium, we can also represent the shock mean distortion as

$$w_{t+1} = khy_t \equiv Wy_t, \tag{8.12}$$

a version of (4.4b).

The Stackelberg timing protocol confronts the consumer with an exogenous w_{t+1} sequence of the form (4.6). We can use (8.12) in the manner described above to create such an exogenous representation for w_{t+1}. The household can then be regarded as solving an ordinary control problem subject to (8.2), (4.6), and (8.12).

8.4. Effects on consumption of a preference for robustness

To understand the effects on consumption of a preference for robustness, we use as a benchmark Hall's case of $\beta R = 1$ and no preference for robustness ($\sigma = 0$). In that case, we have seen that μ_{ct} and consumption are both driftless random walks. To be concrete, we set the parameters of our example to be consistent with ones calibrated from post-World War U.S. time series by Hansen, Sargent, and Tallarini (1999) for a more general permanent income model. They set $\beta = .9971$ and fit a two-factor model for the endowment process where each factor is a second order autoregression. To simplify their specification, we replace their estimated two-factor endowment process with the population first-order autoregression one would obtain if that two factor model actually generated the data. Thus, fitting the first-order autoregressive process (8.2b) with $w_{t+1} \equiv 0$ using the population moments implied by Hansen, Sargent, and Tallarini's (HST's) estimated endowment process we obtain the endowment process $z_{t+1} = .9992 z_t + 5.5819 \epsilon_{t+1}$, where ϵ_{t+1} is an i.i.d. scalar process with mean zero and unit variance.[12] We use $\hat{\beta}$ to denote HST's $\beta = .9971 \equiv \hat{\beta}$. Throughout, we suppose that $R = \hat{\beta}^{-1}$.

We now consider three cases.

- The $\beta R = 1, \sigma = 0$ case studied by Hall (1978). With $\beta = \hat{\beta} \equiv .9971$, we compute that the marginal utility of consumption follows the law of motion

$$\mu_{c,t+1} = \mu_{c,t} + 4.3825 \epsilon_{t+1} \tag{8.13}$$

  where we computed the coefficient 4.3825 on ϵ_{t+1} using a formula from Hansen and Sargent (2005).
- A version of Hall's $\beta R = 1$ specification with a preference for robustness. Retaining $\hat{\beta} R = 1$, we activate a preference for robustness by setting $\sigma = \hat{\sigma} - 2E - 7 < 0$. Under the approximating model, we now compute that

$$\mu_{c,t+1} = .9976 \mu_{c,t} + 8.0473 \epsilon_{t+1}. \tag{8.14}$$

[12] We computed ρ, c_d by calculating autocovariances implied by HST's specification, then used them to calculate the implied population first-order autoregressive representation.

When $b - c_t > 0$, this equation implies that $E_t(b - c_{t+1}) = .9976(b - c_t) < (b - c_t)$ which in turn implies that $E_t c_{t+1} > c_t$. Thus, the effect of activating a preference for robustness is to put upward drift into the consumption profile, a manifestation of a kind of 'precautionary savings'.

- A case that raises the discount factor relative to the $\beta R = 1$ benchmark prevailing in Hall's model, but withholds a preference for robustness. In particular, while we set $\sigma = 0$ we increase β to $\tilde{\beta} = .9995$. Remarkably, with $(\sigma, \beta) = (0, \tilde{\beta})$, we compute that $\mu_{c,t+1}$ continues to obey exactly (8.14).

The second and third bullets suggest an observational equivalence result about offsetting changes in σ and β. Thus, starting from $(\sigma, \beta) = (0, \hat{\beta})$, the effect on consumption and saving of activating a preference for robustness by lowering σ so that $\sigma < 0$ while keeping β constant are evidently equivalent with keeping $\sigma = 0$ but *increasing* the discount factor to $\tilde{\beta} > \hat{\beta}$.

These numerical examples illustrate what can be confirmed more generally, that in the permanent income model an increased preference for robustness operates exactly like an increase in the discount factor β. In particular, let $\alpha^2 = \nu'\nu$ and suppose that instead of the particular pair $(\hat{\sigma}, \hat{\beta})$, where $(\hat{\sigma} < 0)$, we use the pair $(0, \tilde{\beta})$, where $\tilde{\beta}$ satisfies:

$$\tilde{\beta}(\sigma) = \frac{\hat{\beta}\left(1 + \hat{\beta}\right)}{2\left(1 + \sigma\alpha^2\right)} \left[1 + \sqrt{1 - 4\hat{\beta}\frac{1 + \sigma\alpha^2}{\left(1 + \hat{\beta}\right)^2}}\right]. \tag{8.15}$$

Then the laws of motion for $\mu_{c,t}$, and therefore the decision rules for c_t, are identical across these two preference specifications. Hansen and Sargent (2005) establish formula (8.15).

8.5. Observational equivalence and distorted expectations

8.6. Distorted endowment process

When $\sigma = 0$, the consumer faces a law of motion for the state vector y_t that we can represent as

$$y_{t+1} = Ay_t + B\mu_{ct} + C\epsilon_{t+1}. \tag{8.16}$$

When $\sigma < 0$, the minimizing agent adds Cw_{t+1} to the right side of this transition law, where $w_{t+1} = Wy_t$. Under the Stackelberg timing protocol, the maximizing agent faces an exogenous $\{w_{t+1}\}$ process that evolves according to

$$Y_{t+1} = (A + Bh + CW)\, Y_t + C\epsilon_{t+1} \tag{8.17a}$$

$$w_{t+1} = WY_t. \tag{8.17b}$$

Under the Stackelberg timing protocol, the maximizing player thus faces the following transition law for y:

$$\begin{bmatrix} y_{t+1} \\ Y_{t+1} \end{bmatrix} = \begin{bmatrix} A & CW \\ 0 & A + Bh + CW \end{bmatrix} \begin{bmatrix} y_t \\ Y_t \end{bmatrix} + \begin{bmatrix} C \\ C \end{bmatrix} \epsilon_{t+1}. \tag{8.18}$$

If the decision maker solves an ordinary dynamic programming program without a preference for robustness but substitutes the distorted transition law for the one given by his approximating model, he attains robust decision rules. Thus, when $\sigma < 0$, instead of facing the transition law (8.16) that prevails under the approximating model, the household would use the distorted transition law (8.18).

It is useful to consider our observational equivalence result in light of the distorted law of motion (8.18). Let $\hat{E}_t$ denote a conditional expectation taken with respect to the distorted transition law (8.18) for the endowment shock and let E_t denote the expectation taken with respect to the approximating model. Then the observational equivalence of the pairs $(\hat{\sigma}, \hat{\beta})$ and $(0, \tilde{\beta})$ means that the following two versions of (8.6) imply the same μ_{ct} processes:

$$\mu_{ct} = -\left(1 - R^{-2}\hat{\beta}^{-1}\right)\left(Rx_{t-1} + \hat{E}_t \sum_{j=0}^{\infty} R^{-j}\left(z_{t+j} - b\right)\right)$$

and

$$\mu_{ct} = -\left(1 - R^{-2}\tilde{\beta}^{-1}\right)\left(Rx_{t-1} + E_t \sum_{j=0}^{\infty} R^{-j}\left(z_{t+j} - b\right)\right).$$

For both of these expressions to be true, the effect on $\hat{E}$ of setting σ less than zero must be just offset by the effect of raising β from $\hat{\beta}$ to $\tilde{\beta}$.

8.7. Equivalence of quantities but not continuation values

Our numerical examples illustrate that, holding other parameters constant, there exists a locus of (σ, β) pairs that imply the same consumption, saving programs. Furthermore, it can be verified that the P matrices appearing in the quadratic forms in the value function are identical for the $(\hat{\sigma}, \hat{\beta})$ and $(0, \tilde{\beta})$ problems. However, in terms of their implications for pricing claims on risky future payoffs, it is significant that the $\mathcal{D}(P)$ matrices differ across such (σ, β) pairs. For the $(0, \tilde{\beta})$ pair, $P = \mathcal{D}(P)$. However, when $\sigma < 0$, $\mathcal{D}(P)$ differs from P. HST show that $\mathcal{D}(P)$ encodes the shadow prices that are relevant for pricing uncertain claims on future consumption. Thus, although the $(\hat{\sigma}, \hat{\beta})$ and $(0, \tilde{\beta})$ parameter settings imply identical savings and consumption plans, they imply different valuations of risky future consumption payoffs. HST and Hansen, Sargent, and Wang (2002) used this fact to study how a preference for robustness influences the equity premium.

9. Concluding remarks

In delineating optimization and expectations formation, Simon and Theil's certainty equivalence principle takes for granted that the decision maker has 'rational expectations' about the exogenous variables that impinge on his one-period return function. For single-agent problems, the assumption of rational expectations was so natural for Simon and Theil and their readers that it was received without comment or controversy. The hypothesis of rational expectations became more technically challenging and controversial when Simon's co-author[13] Muth (1961) applied it to forecasts about *endogenous* variables, like prices, whose laws of motion were to be determined by an equilibrium shaped by a representative agent's decision rules. Muth's analysis thus required a fixed point argument that Simon and Theil did not need.

Our introduction reminisced about how extensively the certainty equivalence principle has served as a laboratory for developing applied rational expectations models. This paper has told how the certainty equivalence principle also pertains to settings where the decision

[13] See Holt, Modigliani, Muth, and Simon (1960).

maker distrusts his model, unlike the decision maker inside rational expectations models. Our decision maker's fear of model misspecification makes him appear to have distorted or "irrational" expectations relative to his approximating model. The decision maker achieves robustness by distorting his expectations. Anderson, Hansen, and Sargent (2002), Hansen, Sargent, and Tallarini (1999) and Hansen and Sargent (2005) describe the relevant counterpart to a rational expectations equilibrium for a context where a representative agent fears model misspecification.

References

Aiyagari, Rao, Albert Marcet, Thomas Sargent, and Juha Seppala (2002), "Optimal Taxation without State-Contingent Debt" *Journal-of-Political-Economy.* Vol. 110, No. 6: pp. 1220–1254.

Anderson, E.W. (1998), "Uncertainty and the Dynamics of Pareto Optimal Allocations," University of Chicago Dissertation.

Anderson, E.W., L.P. Hansen, and Thomas Sargent (2003), "A Quartet of Semi-Groups for Model Specification, Robustness, Prices of Risk, and Model Detection," *Journal of the European Economic Association*, Vol.1, No.1, pp. 68–123.

Başar, T. and P. Bernhard (1995), *H^∞-Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*, Boston-Basel-Berlin: Birkhäuser

Blackwell, D. and and M. Girshick (1954), *Theory of Games and Statistical Decisions*, Wiley: New York.

Gilboa, I. and D. Schmeidler (1989), "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics* 18, 141-153.

Hall, Robert E. (1978), "Stochastic Implications of the LIfe-Cycle Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, Vol. 86, No. 6, pp. 971–988.

Hansen, Lars P., William Roberds, and Thomas J. Sargent (1991), "Time Series Implications of Present Value Budget Constraint and of Martingale Models of Consumption and

Taxes," in Lars Peter Hansen and Thomas J. Sargent, *Rational Expectations Econometrics*, Westview Press, Boulder, Colorado

Hansen, L.P. and T.J. Sargent (1995), "Discounted Linear Exponential Quadratic Gaussian Control," *IEEE Transactions on Automatic Control* 40:5, 968-971.

Hansen, L.P. and T.J. Sargent (2005), *Robust Control and Economic Model Uncertainty*, Princeton: Princeton University Press, forthcoming

Hansen, L.P., T.J. Sargent and T.D. Tallarini, Jr. (1999), "Robust Permanent Income and Pricing," *Review of Economic Studies*, 66, 873-907.

Hansen, L.P., T.J. Sargent, and N.E. Wang (2002), "Robust Permanent Income and Pricing with Filtering", *Macroeconomic Dynamics*, Vol. 6, February, pp. 40–84.

Hansen, L.P., T.J. Sargent, G.A. Turmuhambetova, and N. Williams (2004), "Robustness and Uncertainty Aversion," manuscript, University of Chicago and Hoover Institution.

Holt, C.C., F. Modigliani, J.F. Muth, and H.A. Simon (1960), *Planning Production, Inventories, and Work Force*, Englewood Cliffs, N.J.

Jacobson, D. J. (1973), "Optimal Linear Systems with Exponential Performance Criteria and their Relation to Differential Games," *IEEE Transactions on Automatic Control*, Vol.18, pp. 124–131

Lucas, R. E., Jr. and Thomas Sargent (1981), "Introduction" in *Rational Expectations and Econometric Practice*, Minneapolis, University of Minnesota Press.

Lucas, R.E. , Jr (1976), "Econometric Policy Evaluation: A Critique," *Journal of Monetary Economics*, 1:2, 19-46.

Muth, J. F. (1961), "Rational Expectations and the Theory of Price Movements," *Econometrica*, Vol. 29, pp. .

Sargent, Thomas J. (1987), *Macroeconomic Theory*, Academic Press

Schmitt-Grohe, Stephanie and Martin Uribe (2004), "Solving dynamic general equilibrium models using a second-order approximation to the policy function," *Journal of Economic Dynamics and Control*, Vol. 28, pp. 755-775.

Simon, Herbert A. (1956), " Dynamic Programming Under Uncertainty with a Quadratic Criterion Function," *Econometrica*, Vol. 24, pp. 74-81.

Tallarini, Jr., T.D. (2000), "Risk-Sensitive Real Business Cycles," forthcoming in *Journal of Monetary Economics*, vol. 45, pp. 507–532.

Theil, Henri (1957), "A Note on Certainty Equivalence in Dynamic Planning," *Econometrica*, Vol. 25, April, pp. 346–349.

Whittle, P (1990), *Risk-Sensitve Optimal Control*, New York: Wiley.

Whittle, P. (1996), *Optimal Control: Basics and Beyond*, New York: Wiley

Zhou, K., J. C. Doyle and K. Glover (1996), *Robust and Optimal Control*, London: Prentice Hall