

Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making *

Timothy Cogley[†] Thomas J. Sargent[‡]

Revised: March 2005

Abstract

For a Markov decision problem in which unknown transition probabilities serve as hidden state variables, we study the quality of two approximations to the decision rule of a Bayesian who each period updates his subjective distribution over the transition probabilities by Bayes' law. The first is the usual rational expectations approximation that assumes that the decision maker knows the transition probabilities. The second approximation is a version of Kreps' (1998) anticipated utility model in which decision makers update using Bayes' law but optimize in a way that is myopic with respect to their updating of probabilities. For a range of consumption smoothing examples, the anticipated utility approximation outperforms the rational expectations approximation. The anticipated utility and Bayesian models augment market prices of risk relative to the rational expectations approximation.

KEY WORDS: Rational expectations, Bayes' Law, anticipated utility, market price of risk.

1 Introduction

A standard model of intertemporal choice involves dynamic programming with time-invariant beliefs. Let s_t , u_t , and ε_t represent vectors of state variables, controls, and shocks, respectively, and suppose an agent maximizes a discounted sum of returns

$$E_0 \sum_{t=0}^{\infty} \beta^t r(s_t, u_t), \tag{1}$$

*We thank Lars Hansen, Narayana Kocherlakota, and seminar participants at Stanford for their comments and suggestions.

[†]University of California, Davis. Email: twcogley@ucdavis.edu

[‡]New York University and Hoover Institution. Email: ts43@nyu.edu

subject to a transition equation

$$s_{t+1} = g(s_t, u_t, \varepsilon_{t+1}). \quad (2)$$

The return function is assumed to be concave, the constraint set is convex and compact, and the decision maker knows all the parameters of the model. With these assumptions, the maximization problem can be cast as a dynamic program. The Bellman equation is

$$V(s) = \max_u \{r(s, u) + \beta E(V[g(s, u, \varepsilon)|s])\} \quad (3)$$

and the optimal decision rule is

$$u = h(s). \quad (4)$$

The assumption that agents know the parameters of the model means that learning has been completed. Accordingly, the standard model is most useful for studying mature economies in which agents have already acquired enough experience that new observations have a negligible effect on their beliefs. The no-learning assumption is a convenient simplification in cases like this. In other circumstances, however, learning is a more prominent feature of the problem. Examples include transition economies in which equilibria are punctuated by big changes in economic or political institutions, as well as economies in which government policy makers adapt their policy rules as their understanding of the structure evolves. Changing beliefs are likely to be more important in cases like these, and their effect on outcomes would be lost if learning were neglected.

One way to model learning and dynamic choice is to withdraw knowledge of the transition equation (2) and replace it with an estimated transition equation

$$s_{t+1} = g(s_t, u_t, \varepsilon_{t+1}, \hat{\theta}_t), \quad (5)$$

where $\hat{\theta}_t$ represents estimates of the parameters governing (2) conditional on data through date t , updated by recursive least squares or the application of Bayes's theorem. Then one could solve the dynamic program as before using the estimated transition equation (5) instead of (2). This delivers a time-varying decision rule

$$u_t = h(s_t, \hat{\theta}_t), \quad (6)$$

that depends on date- t beliefs about how the state vector evolves. As beliefs are updated, so too are decision rules. Kreps (1998) recommends this approach and refers to it as an 'anticipated-utility' model. This modeling strategy forms the basis of much of the macroeconomics literature on learning (e.g., see Sargent 1993, 1999 and Evans and Honkapohja 2001).

This approach is mildly schizophrenic, however, in the way it treats $\hat{\theta}_t$. Parameters are treated as random variables when agents learn but as constants when they

formulate decisions. Looking backward, agents can see how their beliefs have evolved in the past, but looking forward they act as if future beliefs will remain unchanged forever. Agents are eager to learn at the beginning of each period, but their decisions reflect a pretence that this is the last time they will update beliefs, a pretence that is falsified at the beginning of every subsequent period.

A Bayesian decision maker does not behave this way. A Bayesian would treat θ as a random variable both for learning and decision making and so would recognize that beliefs will continue to evolve going forward in time. He would recognize this source of uncertainty when formulating decision rules.

This is easier said than done. In many interesting applications, especially those in macroeconomics, a full Bayesian procedure is too complicated to be implemented. Macroeconomists might justify anticipated-utility models as an approximation to a correctly formulated Bayesian decision problem. The computational cost of the full Bayesian calculation makes anticipated utility approximation appealing, but the appeal would be more compelling if one could also show that anticipated-utility decisions well approximate Bayesian decisions. As far as we know, no one has assessed the quality of the approximation, mainly because one has to calculate Bayesian choices in order to make the comparison, and that is hard to do.

In this paper, we develop a laboratory for exploring Bayesian and anticipated-utility choices. Our objectives are to study how to implement the Bayesian procedure and to examine how well anticipated-utility choices approximate Bayesian decisions. We show how to find the exact Bayesian solution in a few simple examples rigged for maximum tractability, and then we compare Bayesian and anticipated-utility choices. The Bayesian approach turns out to be simpler than we first imagined, making us hopeful about extensions to more realistic problems. We also find that the anticipated-utility model often provides an excellent approximation, which makes us more confident about its application as well.

The remainder of the discussion is organized as follows. Section 2 discusses a key step in the Bayesian treatment, namely, how to expand the state vector to encompass learning and how to derive the transition equation for the expanded state. Sections 3-5 study a trio of simple examples, solve for the Bayesian equilibrium in each case, and compare the outcomes with those of two approximations, based on anticipated utility and complete-information rational expectations, respectively. We conclude with an assessment of anticipated utility as an approximation strategy.

2 The State Vector and Transition Equation for a Bayesian Model

As before, let s_t represent the ‘natural’ state vector. We assume it is observable, exogenous, and that it evolves according to the transition density

$$f(s_{t+1}|s_t, \theta). \tag{7}$$

Agents know the functional form of $f(\cdot)$, but they do not know θ . They learn about θ by applying Bayes's theorem to observations on s_t .

We make two other assumptions for tractability. First, we limit attention to probability models in which the information in the history s^t can be summarized by a finite-dimensional vector of sufficient statistics ζ_t . These statistics encode the information relevant for the learning problem. A big class of probability models satisfies this condition; for example, all members of the exponential family possess a finite-dimensional vector of sufficient statistics. Second, we also limit attention to models in which the learning problem can be cast in terms of a conjugate prior and likelihood, so that posterior distributions can be expressed analytically. This limitation is more restrictive than the first, but many interesting models can be set up in this way.

Our strategy is to append the vector of sufficient statistics ζ_t to the natural state vector s_t ,

$$S_t = [s'_t, \zeta'_t]', \quad (8)$$

and then to derive a transition density for the expanded state vector S_t ,

$$f(S_{t+1}|S_t). \quad (9)$$

If (9) can be expressed as a time-invariant function of the augmented state vector, then decision rules can be derived by dynamic programming. We first consider a general form for (9) and then give an example to show how it works.

2.1 A General Form of the Transition Density

To begin, factor (9) as

$$\begin{aligned} f(S_{t+1}|S_t) &= f(s_{t+1}, \zeta_{t+1}|s_t, \zeta_t), \\ &= f(\zeta_{t+1}|s_{t+1}, s_t, \zeta_t) f(s_{t+1}|s_t, \zeta_t). \end{aligned} \quad (10)$$

The first term on the right side, $f(\zeta_{t+1}|s_{t+1}, s_t, \zeta_t)$, is the density for the learning statistics ζ_{t+1} conditional on their previous values along with realizations of the natural states. This term describes how the sufficient statistics are updated in light of observations on s_{t+1} and s_t . The updating rules are deterministic given the conditioning information, so we can express the updated value as a function $\zeta(s_{t+1}, S_t)$. That ζ_{t+1} is a deterministic function of (s_{t+1}, S_t) means that this term is a delta function assigning unit probability mass to the updated value,

$$f(\zeta_{t+1}|s_{t+1}, s_t, \zeta_t) = \delta(\zeta(s_{t+1}, S_t)). \quad (11)$$

The particular form of the updating rule $\zeta(s_{t+1}, S_t)$ follows from Bayes's theorem and depends on how the agent's priors and the conditional likelihood for (7) are specified. Analytical expressions for the updating formulas are available as long as we work within a conjugate family.

The second term on the right side of (10) can be regarded as a posterior predictive density. This term can be expanded as

$$\begin{aligned} f(s_{t+1}|s_t, \zeta_t) &= \int f(s_{t+1}, \theta|s_t, \zeta_t) d\theta, \\ &= \int f(s_{t+1}|s_t, \zeta_t, \theta) f(\theta|s_t, \zeta_t) d\theta, \\ &= \int f(s_{t+1}|s_t, \theta) f(\theta|s_t, \zeta_t) d\theta. \end{aligned} \tag{12}$$

On the bottom line, the first term in the integrand, $f(s_{t+1}|s_t, \theta)$, is the natural transition equation with which we started. Here it is conditioned on a particular value of θ , which makes the learning statistics redundant as conditioning variables. This is a conditional likelihood function for s_{t+1} .

Of course, θ is unknown, and the second term in the integrand summarizes beliefs about it. The term $f(\theta|s_t, \zeta_t)$ is the posterior for θ conditioned on information available through date t . Since the vector ζ_t is a sufficient statistic for the history s^t , this can also be expressed as $f(\theta|s_t, \zeta_t) = f(\theta|s^t)$. Again, as long as we work within a conjugate family, expressions for $f(\theta|s_t, \zeta_t)$ will be available in closed form.

Equation (12) represents beliefs about s_{t+1} conditioned on information available through date t . Analytically solving this integral is often hard, but there are a number of cases in which it can be done. Two examples are given below. In other cases, the integral can be solved numerically, e.g. as described by Gelman, et. al. (1995, p. 301),¹ but we do not pursue numerical strategies here.

Equation (12) is the second piece needed for (9). Combining this with the deterministic updating relation (11) produces an expression for (9),

$$f(S_{t+1}|S_t) = \delta(\zeta(s_{t+1}, S_t)) \int f(s_{t+1}|s_t, \theta) f(\theta|s_t, \zeta_t) d\theta. \tag{13}$$

As promised, this is a time-invariant function of S_t , suitable for dynamic programs.

2.2 An Example: A Beta-Binomial Model for a Two-State Markov Process

Suppose the exogenous state s_t takes two values, 0 and 1, and that the transition probabilities are governed by a Markov transition matrix,

$$\Pi = \begin{bmatrix} p & 1-p \\ 1-q & q \end{bmatrix}, \tag{14}$$

¹This involves drawing a θ from its posterior, plugging the draw into the conditional model $f(s_{t+1}|s_t, \theta)$, and then simulating s_{t+1} from that model. A sample generated in this way approximates a sample from the posterior predictive density (12). For an application, see Cogley, Morozov, and Sargent (2003).

where

$$\begin{aligned}
pr(s_{t+1} = 1 | s_t = 1) &= p, \\
pr(s_{t+1} = 0 | s_t = 1) &= 1 - p, \\
pr(s_{t+1} = 0 | s_t = 0) &= q, \\
pr(s_{t+1} = 1 | s_t = 0) &= 1 - q.
\end{aligned} \tag{15}$$

The states are assumed to be observable, but the transition probabilities are unknown. Our agent learns about them using Bayes's theorem. Because the states are discrete and take on two values, a beta-binomial probability model is convenient. We assume the agent has independent beta priors over (p, q) ,

$$f(p, q) = f(p)f(q), \tag{16}$$

where

$$\begin{aligned}
f(p) &\propto p^{n_0^{11}-1}(1-p)^{n_0^{10}-1}, \\
f(q) &\propto q^{n_0^{00}-1}(1-q)^{n_0^{01}-1}.
\end{aligned} \tag{17}$$

The variable n_t^{ij} represents a counter that records the number of transitions from state i to j through date t . The parameters n_0^{ij} represent prior beliefs about the frequency of transitions, which may for example come from a training sample. The likelihood function for a batch of data s^t is proportional to the product of binomial densities,

$$f(s^t | p, q) \propto p^{(n_t^{11}-n_0^{11})}(1-p)^{(n_t^{01}-n_0^{01})} q^{(n_t^{00}-n_0^{00})}(1-q)^{(n_t^{01}-n_0^{01})},$$

where $(n_t^{ij} - n_0^{ij})$ are the number of transitions from state i to j observed in the sample.² Multiplying the likelihood by the prior delivers the posterior kernel,

$$\begin{aligned}
f(p, q | s^t) &\propto p^{n_t^{11}-1}(1-p)^{n_t^{10}-1} q^{n_t^{00}-1}(1-q)^{n_t^{01}-1}, \\
&\propto f(p | s^t) f(q | s^t),
\end{aligned} \tag{18}$$

where

$$\begin{aligned}
f(p | s^t) &= \text{beta}(n_t^{11}, n_t^{10}), \\
f(q | s^t) &= \text{beta}(n_t^{00}, n_t^{01}).
\end{aligned} \tag{19}$$

Given independent beta priors over p and q and a likelihood function that is a product of binomials, it follows that the posteriors are also independent and have the beta form.

In terms of the notation of the last subsection, equation (14) summarizes the natural transition density for this problem, and θ refers to the transition probabilities

²According to this notation, n_t^{ij} represents the sum of prior plus observed counters.

(p, q) . The vector of counters is a sufficient statistic for learning about the transition probabilities, so n_t takes the place of ζ_t . Thus the expanded state S_t consists of the two Markov states s_t along with the counters n_t ,

$$S_t = [s'_t, n'_t]'. \quad (20)$$

For finite-horizon economies, this is a finite-state process, because there are only two values for s_t at any date plus a finite number of permutations of n_t . The nodes for this state are every possible permutation of counters that can be attained along paths of s_t , combined with both elements of s_t . This transforms the two-state process into a large multi-state process. Also notice that the dimension of S_t grows with t , because the number of possible permutations of n_t increases.

Our goal is to derive a matrix of transition probabilities that maps the probability of moving from any element of S_t to any element of S_{t+1} ,

$$P_{t,t+1}^{lm} = pr[S_{t+1} = m | S_t = l], \quad (21)$$

for $l = 1, \dots, \dim(S_t)$ and $m = 1, \dots, \dim(S_{t+1})$. The expanded state is a Markov random variable because the probability that $S_{t+1} = m$ conditional on the past history S^t depends on a single lag S_t . But the process is not homogenous because of the expansion of the dimension of the state. The matrix $P_{t,t+1}$ is rectangular, not square, and the dimensions of $P_{t,t+1}$ change from date to date. It follows that the transition probabilities cannot be independent of t . The elements of $P_{t,t+1}$ are, however, time-invariant functions of S_t .

To derive $P_{t,t+1}^{lm}$, we first deduce the posterior predictive density, $f(s_{t+1}|s_t, n_t)$, and then incorporate how the counters are updated conditional on the passage from s_t to s_{t+1} . For a given s_t and n_t , there are two possible outcomes for s_{t+1} , and the posterior predictive density assigns probabilities to them. In other words, the posterior predictive density consists of two numbers, $pr(s_{t+1} = 1|s_t, n_t)$ and $pr(s_{t+1} = 0|s_t, n_t)$, which of course must sum to one. The first of these can be calculated as

$$\begin{aligned} pr(s_{t+1} = 1|s_t, n_t) &= \iint [ps_t + (1-q)(1-s_t)]f(p, q|s_t, n_t)dpdq, \\ &= s_t \int pf(p|s_t, n_t)dp + (1-s_t) \int (1-q)f(q|s_t, n_t)dq. \end{aligned} \quad (22)$$

This is just the posterior mean of p if $s_t = 1$ and one minus the posterior mean of q if $s_t = 0$. To evaluate the posterior mean, we integrate with respect to the posterior beta density,

$$\begin{aligned}
Ep &= \frac{\Gamma(n_t^{11} + n_t^{10})}{\Gamma(n_t^{11})\Gamma(n_t^{10})} \int p \cdot p^{n_t^{11}-1} (1-p)^{n_t^{10}-1} dp, \\
&= \frac{\Gamma(n_t^{11} + n_t^{10})}{\Gamma(n_t^{11})\Gamma(n_t^{10})} \int p^{n_t^{11}} (1-p)^{n_t^{10}-1} dp, \\
&= \frac{\Gamma(n_t^{11} + n_t^{10})}{\Gamma(n_t^{11})\Gamma(n_t^{10})} \frac{\Gamma(1 + n_t^{11})\Gamma(n_t^{10})}{\Gamma(1 + n_t^{11} + n_t^{10})}, \\
&= \frac{\Gamma(n_t^{11} + n_t^{10})}{\Gamma(1 + n_t^{11} + n_t^{10})} \frac{\Gamma(1 + n_t^{11})}{\Gamma(n_t^{11})}, \\
Ep &= \frac{\Gamma(n_t^{11} + n_t^{10})}{(n_t^{11} + n_t^{10})\Gamma(n_t^{11} + n_t^{10})} \frac{(n_t^{11})\Gamma(n_t^{11})}{\Gamma(n_t^{11})}, \\
&= \frac{n_t^{11}}{n_t^{11} + n_t^{10}} \equiv \hat{p}_t.
\end{aligned} \tag{23}$$

The third equality expresses the integral of a beta density in terms of gamma functions, and the fifth equality follows from a recursive property of the gamma function, viz. that $\Gamma(n+1) = n\Gamma(n)$. The bottom line states the intuitive result that the posterior mean of p is just the fraction of times the system stays in state 1 when it begins there, counting both prior and observed transitions. Similarly, by following the same steps, one can show that

$$Eq = \frac{n_t^{00}}{n_t^{00} + n_t^{01}} \equiv \hat{q}_t, \tag{24}$$

the fraction of times the system remains in state 0 when starting there.

These depend on the counters through date t , so we label them \hat{p}_t and \hat{q}_t , respectively. The posterior predictive probability for state 1 can therefore be expressed as

$$pr(s_{t+1} = 1 | s_t, n_t) = \hat{p}_t s_t + (1 - \hat{q}_t)(1 - s_t). \tag{25}$$

In the same way, the predictive probability for state 0 is

$$pr(s_{t+1} = 0 | s_t, n_t) = (1 - \hat{p}_t)s_t + \hat{q}_t(1 - s_t). \tag{26}$$

To represent the joint transition density for s_{t+1} and n_{t+1} , we need to combine these marginal predictive probabilities with the deterministic conditional relationship that governs how the counters are updated in light of the news in s_{t+1} . For each of

the four possible s -transitions, one element of n_t is incremented by 1, and the others remain constant. The first element n_{t+1}^{11} increases by 1 when passing from $s_t = 1$ to $s_{t+1} = 1$, the second element n_{t+1}^{10} increases by 1 when passing from $s_t = 1$ to $s_{t+1} = 0$, and so on. One can enumerate the probabilities associated with the four possible joint transitions:

$$\begin{aligned} pr(s_{t+1} = 1, n_{t+1} = n_t + [1 \ 0 \ 0 \ 0]' | s_t = 1, n_t) &= \hat{p}_t, \\ pr(s_{t+1} = 1, n_{t+1} = n_t + [0 \ 1 \ 0 \ 0]' | s_t = 0, n_t) &= 1 - \hat{q}_t, \\ pr(s_{t+1} = 0, n_{t+1} = n_t + [0 \ 0 \ 1 \ 0]' | s_t = 1, n_t) &= 1 - \hat{p}_t, \\ pr(s_{t+1} = 0, n_{t+1} = n_t + [0 \ 0 \ 0 \ 1]' | s_t = 0, n_t) &= \hat{q}_t. \end{aligned} \tag{27}$$

The matrix $P_{t,t+1}$ is formed from these outcomes. Imagine iterating through the rows of S_t and S_{t+1} , matching each element of the former with each element of the latter. There are two kinds of matches, admissible and inadmissible. A match is admissible if the updating of n_{t+1} is consistent with the movement from s_t to s_{t+1} . Otherwise it is inadmissible. The four admissible matches are as follows. For $l = 1, \dots, \dim(S_t)$ and $m = 1, \dots, \dim(S_{t+1})$,

- if $s_t^l = 1$, $s_{t+1}^m = 1$, $n_{m,t+1}^{11}$ increases by one, and the other counters remain unchanged, set $P_{t,t+1}^{lm} = n_{lt}^{11} / (n_{lt}^{11} + n_{lt}^{10})$;
- if $s_t^l = 1$, $s_{t+1}^m = 0$, $n_{m,t+1}^{10}$ increases by one, and the other counters remain unchanged, set $P_{t,t+1}^{lm} = n_{lt}^{10} / (n_{lt}^{11} + n_{lt}^{10})$;
- if $s_t^l = 0$, $s_{t+1}^m = 1$, $n_{m,t+1}^{01}$ increases by one, and the other counters remain unchanged, set $P_{t,t+1}^{lm} = n_{lt}^{01} / (n_{lt}^{01} + n_{lt}^{00})$.
- If $s_t^l = 0$, $s_{t+1}^m = 0$, $n_{m,t+1}^{00}$ increases by one, and the other counters remain unchanged, set $P_{t,t+1}^{lm} = n_{lt}^{00} / (n_{lt}^{01} + n_{lt}^{00})$.

All other matches correspond to pairs in which the change in s_{t+1} is inconsistent with the change in the counters. Accordingly, their transition probabilities are zero.

This delivers one-step transition matrices for the expanded state. As promised, the conditional transition probabilities depend on a single lag of S_t , and the elements of $P_{t,t+1}$ are time-invariant functions of the state. Armed with these transition matrices, one can solve finite-horizon dynamic programs in the usual way. The main challenge involves coping with the curse of dimensionality.

This example can be extended without too much trouble to a Markov process with more than two states by adopting a Dirichlet-multinomial probability model. The logic of the argument is the same, only the details differ. This extension is presented in appendix A.

Infinite-horizon problems are more difficult because the counters are unbounded, violating the assumption that the state space is compact. In this example, however,

the Bayesian consistency theorem holds, so eventually \hat{p}_t and \hat{q}_t converge in probability to p and q . In other words, learning is relevant only in finite samples in any case. One way to approximate the solution of an infinite-horizon model is to choose a horizon large enough that further learning can be neglected. Then adopt the finite-horizon approach to approximate outcomes until that date and a no-learning model for outcomes thereafter. Terminal conditions for the finite-horizon calculation can be derived from the value function and decision rules for the no-learning model. Thus, at least in principle, one can approximate an infinite-horizon problem by decomposing it into two segments, a finite-horizon problem in which learning matters plus an infinite-horizon remainder in which beliefs have settled down.³

3 Three Experiments

Having cast the Bayesian problem in a form suitable for dynamic programming, we can proceed to compare anticipated-utility and Bayesian choices. As a laboratory for this comparison, we adopt a finite-horizon version of a permanent income model. We chose the permanent income model because it is a canonical example of dynamic choice, and we assume a finite horizon because this keeps the state space compact. We examine three versions to explore how the quality of the approximation depends on various features of the environment.

3.1 A Finite-Horizon Version of Hall's Model

We begin with a finite-horizon version of Hall's (1978) consumption model. We start here because Hall's model assumes certainty equivalence, which is a common assumption in macroeconomics,⁴ and because certainty equivalence helps us cope with the curse of dimensionality. For Hall's model, consumption decision rules can be expressed in terms of wealth and expectations of future labor income, and the Bayesian dynamic program reduces to a forecasting problem. This is easily solved if labor income evolves exogenously as a finite-state Markov process.

Imagine an economy that operates for T periods. Aggregate income is exogenous and follows a two-state Markov process to be described below. The economy is inhabited by a representative consumer who wants to smooth fluctuations in income. He formulates a consumption plan to maximize lifetime utility,

$$V = E_0 \sum_{t=0}^T \beta^t u(c_t) \quad (28)$$

³We do not pursue this idea further. The calculations reported below focus on finite-horizon economies.

⁴In some cases, certainty equivalence follows from assumptions about preferences; in others, it is invoked as part of an approximation strategy.

where $u(c_t) = -1/2(c_t - c_0)^2$. In the simulations conducted below, the bliss point c_0 is set equal to 5, which corresponds to 5 times mean income in one case and 5 times initial income in another.

Three assets are available for smoothing consumption. One is a linear storage technology,

$$A_{t+1} = R(A_t + i_t), \quad (29)$$

where A_t represents the amount already in storage at the beginning of period t and i_t is the net amount added to or subtracted from it. There are no shocks here, so storage is a risk-free investment paying a constant gross return R , which we assume is the inverse of the subjective discount factor, $R = \beta^{-1}$.

Two Arrow securities also trade at each date. Each pays one unit of consumption at $t + 1$ in the event that a particular income state is realized. They are purely inside assets, however, and are in zero net supply, so they do not alter the aggregate resource constraint. After substituting $i_t = y_t - c_t$, we can re-write the aggregate resource constraint as

$$A_{t+1} = R(A_t + y_t - c_t). \quad (30)$$

To find the equilibrium for this economy, we first solve a planning problem for the consumption allocation and then calculate the Arrow prices that are implied by that allocation. Assuming absence of arbitrage, returns for all investment opportunities must satisfy

$$E_t \beta R_{t+1} \frac{u'(c_{t+1})}{u'(c_t)} = 1. \quad (31)$$

When applied to the storage technology, this condition can be solved along with the aggregate resource constraint to obtain decision rules for consumption. When applied to the Arrow securities, (31) pins down asset prices as a function of consumption.

The consumption allocation can be solved by backward induction. Because this is a finite-horizon economy, it is optimal to set $A_{T+1} = 0$. This makes the terminal decision rule

$$c_T = A_T + y_T. \quad (32)$$

The maximization for the penultimate period then becomes

$$\max_{c_{T-1}} u(c_{T-1}) + \beta E_{T-1} u(c_T), \quad (33)$$

subject to

$$\begin{aligned} c_T &= A_T + y_T, \\ A_T &= R(A_{T-1} + y_{T-1} - c_{T-1}). \end{aligned}$$

The first-order condition is

$$u'(c_{T-1}) = \beta R E_{T-1} u'(c_T), \quad (34)$$

which simplifies to

$$c_{T-1} = E_{T-1}c_T, \quad (35)$$

when $\beta R = 1$ and preferences are quadratic. We substitute the terminal decision rule (32) into (35) to express c_T in terms of A_T and y_T , and then use the aggregate resource constraint (30) to express future wealth in terms of current variables. Eventually we find

$$c_{T-1} = \frac{1}{1 + R^{-1}} [A_{T-1} + y_{T-1} + R^{-1}E_{T-1}y_T]. \quad (36)$$

Next, we go back to period $T - 2$ and do the same thing. Continuing backward in this fashion, the consumption decision rule for period $T - h$ can be expressed as

$$c_{T-h} = \gamma_{T-h} \left[A_{T-h} + E_{T-h} \sum_{j=0}^h R^{-j} y_{T-h+j} \right], \quad (37)$$

where

$$\gamma_{T-h} \equiv \left[\sum_{j=0}^h R^{-j} \right]^{-1}. \quad (38)$$

This converges to the infinite-horizon decision rule as h grows large. In the infinite-horizon version of the model, consumption equals the annuity value of wealth. In the finite-horizon version, consumption equals the annuity value plus a fraction of the principal. That fraction grows as the end draws near, so γ_{T-h} increases as h falls.

The decision rule for consumption depends on wealth plus the expected present value of future labor income. To solve for consumption, we just need to calculate those present values. How that is done depends on how agents form expectations.

With the consumption allocation in hand, we go on to calculate prices of Arrow securities. The no-arbitrage condition for an Arrow security promising to pay one unit of consumption in state 1 is,

$$\begin{aligned} 1 &= E_t \beta R_{t+1}^1 \frac{u'(c_{t+1})}{u'(c_t)}, \\ &= \beta \sum_{j=1}^2 R_{t+1}^1(s_{t+1} = j) \frac{u'[c_{t+1}(s_{t+1} = j)]}{u'[c_t(s_t = i)]} pr(s_{t+1} = j | s_t = i). \end{aligned} \quad (39)$$

If state 1 occurs, the return is the inverse of the price, $1/Q_{1t}$; otherwise it is zero. After substituting the return into the no-arbitrage condition, we find that the price is

$$Q(s_{t+1} = 1 | s_t = i) = \beta \frac{u'[c_{t+1}(s_{t+1} = 1)]}{u'[c_t(s_t = i)]} pr(s_{t+1} = 1 | s_t = i). \quad (40)$$

The price of the other Arrow security is found in the same way. Finally, notice that a portfolio consisting of one share of both Arrow securities replicates the risk-free investment, so Arrow prices must also satisfy

$$\sum_{j=1}^2 Q(s_{t+1} = j | s_t) = R^{-1}. \quad (41)$$

In a certainty-equivalent setting such as this, differences between anticipated-utility and Bayesian models boil down to how expectations are formed. Bayesians use the appropriate expectations operator; anticipated-utility consumers take a short cut.

3.1.1 Bayesian Forecasts

Suppose that labor income is exogenous and follows a two-state Markov process. Consumers know the number of states and also the value of income in each state, but they do not know the transition probabilities. They use the beta-binomial model of section 2 for learning and forecasting. The natural state is labor income, and the expanded state adds counters that track its transitions.

Consumption decisions involve multi-step forecasts, so we need the m -step transition matrix $P_{t,t+m}^{ij} = pr[S_{t+m} = j | S_t = i]$. We can express $P_{t,t+m}$ in terms of the one-step transition matrices defined above using the Chapman-Kolmogorov equation,

$$P_{t,t+m} = P_{t,t+1}P_{t+1,t+2}\dots P_{t+m-1,t+m}. \quad (42)$$

This relationship is derived in appendix B. Armed with this formula, multi-step forecasts of labor income can be calculated as

$$E(y_{t+m} | S_t = i) = \sum_j P_{t,t+m}(i, j) S_{t+m}(j, 1), \quad (43)$$

with the convention that labor income is recorded in the first column of S_{t+m} .⁵

At each date t , we forecast labor income over horizons $h = t + 1, \dots, T$, discount the forecasts at rate $R^{-(h-t)}$, and then sum to find the expected present value of labor income. Adding this to the predetermined value of assets A_t gives us total wealth. To determine consumption, we multiply by γ_t , the date t marginal propensity to consume out of wealth. Finally, we update next period's assets A_{t+1} by substituting consumption into the aggregate resource constraint. In this way, we recursively compute Bayesian consumption choices.

3.1.2 Anticipated Utility Forecasts

Now consider how an anticipated-utility consumer makes forecasts. In each period, he updates estimates of the transition probabilities for s_t by plugging in the current counters, arriving at an estimate Π_t of the transition matrix Π . Looking forward in time, he simplifies by disregarding how future realizations of labor income will alter the counters and future one-step transition probabilities. In other words, he pretends the chain is homogenous, with constant transition probabilities going forward in time. If the chain were homogenous, the m -step transition density would simplify to

$$\Pi_{t,t+m} = \Pi_t \Pi_t \dots \Pi_t = \Pi_t^m. \quad (44)$$

⁵The other columns store the counters.

Forecasts of labor income are now given by

$$E(y_{t+m}|s_t = i) = \sum_j \Pi_{t,t+m}(i, j) s_{t+m}(j, 1), \quad (45)$$

Optimal consumption is calculated in the same way as before, but with a different expectations operator.

Thus the differences between the two models reduce to how the m -step transition matrix is specified. A Bayesian consumer recognizes that the chain is not homogenous and alters the transition matrix according to date and horizon. An anticipated-utility consumer updates his beliefs each period, but then uses a homogenous-chain approximation for long-horizon forecasts.

3.1.3 Rational Expectations

As a benchmark, we also consider a rational expectations consumer, who is assumed to know the true transition probabilities, Π . In this case, the counters are redundant, and s_t is the full state vector. This evolves according to a homogenous Markov chain, so multi-step transition probabilities are just

$$\Pi_{t,t+m} = \Pi^m, \quad (46)$$

for all t . An anticipated-utility consumer behaves like a Bayesian when updating estimates but acts like a rational-expectations consumer when making forecasts.

3.1.4 A Business-Cycle Experiment

Next we study how much this matters in particular cases. The first example is calibrated to mimic a business cycle. The time period is a year, we set $R = \beta^{-1} = 1.04$, and we study a consumer who lives 50 years. To represent a business cycle, we assume that labor income switches between two values

$$y_h = 1.1, \quad y_l = 0.9,$$

with true but unknown transition probabilities

$$\Pi = \begin{bmatrix} 0.75 & 0.25 \\ 0.25 & 0.75 \end{bmatrix}. \quad (47)$$

Because of the symmetry of y and Π , the economy spends half its time in each state, and average income is 1. Income is 10 percent above average in a boom and 10 percent below in a recession, so these are big business cycles. That the diagonal elements are greater than 0.5 means labor income is persistent. If income is high this year, the odds are 3 to 1 that it will be high again next year, and similarly for low-income states. The mean time until the next switch is 4 years.

Appendix C describes an algorithm for constructing the state space for this problem. Its chief virtue is that it finds a parsimonious representation for $\{S_t\}_{t=1}^T$. This is helpful for managing the curse of dimensionality.

Each of our simulations starts with 1000 draws of $\{y_t\}_{t=1}^{50}$ from this process. We initialize wealth at $A_0 = 0$ and allow households to borrow or lend as much as they want at interest rate R , subject to paying off their debts before they die. Then we compare the consumption choices of Bayesian, anticipated-utility, and rational-expectations consumers who face identical income paths, as well as the Arrow security prices that decentralize those allocations.

For the Bayesian and anticipated-utility consumers, we simulate the model for 4 sets of priors that represent various kinds and degrees of disagreement with the true transition probabilities. We want the priors to disagree with the true transition probabilities so that consumers are actually learning (altering their beliefs) as time passes. The experiment would not be challenging enough for the anticipated-utility approximation if experience merely confirmed correct priors.

The prior counters for each case are displayed in table 1.

Table 1: Prior Counters for the Business Cycle Example

	n_0^{hh}	n_0^{hl}	n_0^{lh}	n_0^{ll}
Case 1: Excess Sensitivity	9	1	1	9
Case 2: Excess Smoothness	5	5	5	5
Case 3: Irrational exuberance	9	1	1	1
Case 4: Depression Generation	1	1	1	9

In the first two cases, consumers begin with mistaken beliefs about the persistence of business cycles. This is important for consumption smoothing because the persistence of labor income determines how much of an innovation is consumed and how much is saved. In case 1, consumers initially over-estimate business-cycle persistence. While the true odds of remaining in the current state are 3 to 1, consumers think they are 9 to 1. Because they over-estimate income persistence, their consumption choices are initially too sensitive to movements in labor income, at least relative to consumers who know the true transition probabilities. Accordingly, case 1 is labeled ‘excess sensitivity.’

Case 2 represents the opposite belief. Here we assume a prior that labor income is distributed identically and independently, so that this year’s realization has no predictive power for next year’s value. Because consumers initially under-estimate income persistence, their consumption choices early in life are too smooth relative to consumers who know the true transition probabilities. Accordingly, case 2 is labeled ‘excess smoothness.’

In cases 3 and 4, consumers hold mistaken priors not only about the persistence of states but also about their relative frequencies. The consumers of case 3 are ‘irrationally exuberant.’ Their prior is that good times are the norm and that recessions

are rare events which terminate quickly when they do occur. They initially save too little for rainy days because they under-estimate how many rainy days they will encounter. The consumers of case 4 are more pessimistic, believing that bad times are the norm and that good times occur infrequently and do not last. Because they begin life with a grim prior, we label them the ‘depression generation.’

3.1.5 A Unit-Root Experiment

Our second example involves more persistent movements in income. The time period is still one year, $R = 1.04$, and the consumer lives 50 years. But instead of transient movements between high and low levels of income, we introduce persistent movements in increments to income. We assume that income evolves as

$$y_{t+1} - y_t = \mu_t, \quad (48)$$

where μ_t is a two-state Markov process that switches between

$$\mu_h = 0.025, \quad \mu_l = -0.025. \quad (49)$$

Income is initialized at $y_0 = 1$, so these values for μ represent increments of roughly 2.5 percent per year. Not only does labor income have a unit root, but we also assume that the Markov states are more persistent than in the last example. Here we set the true transition probabilities to

$$\Pi = \begin{bmatrix} 0.95 & 0.05 \\ 0.05 & 0.95 \end{bmatrix}. \quad (50)$$

The economy still spends half its time in each state, but the expected time between switches is 20 years. Thus, the states represent periods of sustained growth and contraction, respectively. Although we abstract from a deterministic trend in y_t , one can interpret the two states as ‘new economies’ and ‘productivity slowdowns.’ The magnitude of the shifts is probably larger than in actual data, but the duration is about right.

In this example, the natural state variable is μ_t , and the learning statistics are the counters that track its transitions. After some algebra, consumption decision rules can be expressed in terms of current income, financial wealth, and forecasts of μ_t ,

$$c_{T-s} = y_{T-s} + \frac{1 - R^{-1}}{1 - R^{-s-1}} A_{T-s} + E_{T-s} \sum_{j=1}^s \frac{R^{-j} - R^{-s-1}}{1 - R^{-s-1}} \mu_{T-s+j}. \quad (51)$$

This follows from the generic form of the decision rule (equation 37) along with the specific form of expectations implied by the unit-root specification for income (equation 48). As before, the differences between anticipated-utility and Bayesian choices involve how multi-step forecasts of the Markov variate μ_t are formed.

The priors for this example are listed in table 2. In the first two scenarios, consumers recognize that the chain is symmetric and that the states occur equally often in the long run, but they are initially mistaken about the degree of persistence. Their priors involve too much persistence in one case and too little in the other. The true mean time between switches is 20 years. In the first scenario, consumers think the current situation is much more durable than it really is, with a prior mean time of 40 years. In the second, they think that states represent business cycles, with a mean switching time of 5 years, when in fact they are much longer lasting.

Table 2: Prior Counters for Unit-Root Example

	n_0^{hh}	n_0^{hl}	n_0^{lh}	n_0^{ll}
Case 1: Too Persistent	39	1	1	39
Case 2: Not Persistent Enough	8	2	2	8
Case 3: Over-Estimate Positive	19	1	1	1
Case 4: Over-Estimate Negative	1	1	1	19

In scenarios 3 and 4, consumers focus too much on one state and are hardly aware of the other. In this instance, they underestimate how often the other state occurs and how long it will last when it does occur. In case 4, a ‘new economy’ really is new, and in case 3 a productivity slowdown takes consumers by surprise. Experience gradually teaches them about the properties of these ‘hidden’ states.

Once again, we simulate 1000 trajectories for y_t and calculate Bayesian, anticipated-utility, and rational-expectations choices for each. We set initial wealth at zero and allow households to borrow or lend as much as they want subject to the constraint that debts are repaid before they die. The state space and transition matrices for μ_t are constructed in the same way as before; only the details differ.

3.2 Departures from Certainty Equivalence: A *CRRA* Experiment

Many macroeconomic models exploit certainty equivalence – either by specifying quadratic preferences or by adopting a certainty-equivalent solution algorithm – but some do not. Departures from certainty equivalence are especially important for asset pricing and welfare comparisons. Accordingly, we also consider an example in which consumption depends on higher-order moments of future labor income.

This example has the same structure as the previous ones, except that we replace quadratic preferences with a period utility function that has constant relative risk aversion,

$$u(c_t) = \frac{c_t^{1-\alpha}}{1-\alpha}. \quad (52)$$

The parameter α is the coefficient of relative risk aversion; in the simulations reported below, we set α equal to 2, 5, 10, or 20. Values around 2 are probably most relevant for macroeconomics, but higher values are sometimes adopted in asset-pricing models.

With *CRRA* preferences, the model becomes computationally more demanding because decision rules cannot be derived analytically. We compensate by altering the length of the decision period. Here we set $\beta = 1.04^3$, so that each period corresponds to 3 years, and $T = 20$, so that consumers live for 20 periods (60 years).

The primary difference between quadratic and *CRRA* utility is not that one involves risk aversion and the other does not,⁶ but rather that *CRRA* preferences involve precautionary motives while quadratic utility does not. To give consumers ample reason to take precautions, we recalibrate the income process to introduce a large but rare ‘crash state.’ Here we assume that labor income switches between two values

$$y_h = 1, \quad y_l = 0.75, \quad (53)$$

with transition probabilities,

$$\Pi = \begin{bmatrix} 0.95 & 0.05 \\ 0.5 & 0.5 \end{bmatrix}. \quad (54)$$

The high-income state represents ‘normal times.’ It occurs about 91 percent of the time and is persistent. The low-income state represents a sharp but infrequent drop in income, and it is not persistent. Conditional on being in the high-income state, the mean switching time is 20 periods, which means that a consumer might go a whole lifetime without experiencing a crash. Nevertheless, the presence of a low-income state activates a motive for precautionary saving.

To activate learning, we once again endow agents with a prior that differs from the true transition probabilities. Here we assume that consumers are pessimistic, as in Cecchetti, Lam, and Mark (2000) and Abel (2002). The prior shown in table 3 exaggerates the frequency and persistence of the crash state and underestimates the frequency and persistence of the normal state. Thus, consumers initially believe that crashes occur too often and last too long when they do occur. This magnifies their interest in precautionary saving.

Table 3: Prior Counters for the *CRRA* Example

	n_0^{hh}	n_0^{hl}	n_0^{lh}	n_0^{ll}
Pessimism	3	1	1	3

The Bayesian consumption allocation is found by dynamic programming. Once again, it is optimal to consume all of one’s resources at the terminal date, implying a

⁶Consumers with quadratic preferences are also risk averse.

terminal value function $V_{T+1} = 0$ and decision rule $A_{T+1} = 0$. For $t \leq T$, the Bellman equation takes the form

$$V_t^B(y_t, A_t, n_t) = \max_{A_{t+1}} [u(A_t + y_t - R^{-1}A_{t+1}) + \beta E_t^B V_{t+1}^B(y_{t+1}, A_{t+1}, n_{t+1})]. \quad (55)$$

The Bayesian expectations operator has the same form as in the earlier examples, but now it applies to continuation values rather than future labor income. The state variables y_t and n_t are already discrete, and it is convenient to approximate the solution by discretizing A_t as well. Accordingly, we specify a grid for A_t in increments of 0.01 and set a range that is broad enough so that the end points are never reached in the simulations.⁷ Then value functions and decision rules are calculated by solving the finite-horizon, finite-state dynamic program, using backward-induction methods described in Judd (1998).

The anticipated-utility allocation is also found by backward induction, but the problem and solution algorithm differ in two respects. First, an anticipated-utility planner neglects to take into account how future counters influence continuation values, so his value function depends only on y_t and A_t ,

$$V_t^{AU}(y_t, A_t) = \max_{A_{t+1}} [u(A_t + y_t - R^{-1}A_{t+1}) + \beta E_t^{AU} V_{t+1}^{AU}(y_{t+1}, A_{t+1})]. \quad (56)$$

To compensate, an anticipated-utility planner re-optimizes each period rather than once and for all, revising value functions and decision rules in accordance with updates of the transition probabilities. Like the Bayesian planner, he solves the problem by backward induction, using finite-horizon, finite-state methods, but he does so repeatedly.

A rational-expectations planner knows the true transition probabilities and need not keep track of the counters. Thus, the rational-expectations state vector is (y_t, A_t) , and the Bellman equation is

$$V_t^{RE}(y_t, A_t) = \max_{A_{t+1}} [u(A_t + y_t - R^{-1}A_{t+1}) + \beta E_t^{RE} V_{t+1}^{RE}(y_{t+1}, A_{t+1})]. \quad (57)$$

But because there is no news about transition probabilities, there is no need to update decision rules, so a rational-expectations planner solves the dynamic program only once.

For this example, our simulations work as follows. First, we solve for the Bayesian and rational-expectations decisions rules and store them. Then we simulate 1000 paths for the exogenous income process. Having stored the decision rules, we can simply read off the associated Bayesian and RE consumption paths for each income path. Solving for anticipated-utility consumption involves an extra step because consumers must also re-optimize at each date. So at each date on each path, they re-solve the AU dynamic program using updated transition probabilities, then read off the consumption choice from the revised decision rule. Finally, with consumption allocations in hand, Arrow security prices are calculated from the Euler equations.

⁷The range is case specific and depends on α , which controls the strength of the precautionary motive.

4 Anticipated Utility v. Bayesian Outcomes

The first set of results compares Bayesian and anticipated-utility outcomes. Figures 1-3 portray consumption in the middle year of each simulation.⁸ Each panel represents a cross section of consumption choices over 1000 draws for labor income, with the Bayesian choice shown on the horizontal axis and the anticipated-utility choice on the vertical. Because the two decision makers face identical income streams within each experiment, a good approximation should result in a tight scatterplot along the 45 degree line. The figures confirm that Bayesian and anticipated-utility consumption choices are highly correlated and tightly arrayed along the 45 degree line.




Figure 1: Consumption in Period 25 of the Business-Cycle Simulation




Figure 2: Consumption in Period 25 of the Unit-Root Simulation

⁸The middle year was chosen to avoid end-point problems; apart from that, the results are typical of those in other years.




Figure 3: Consumption in Period 10 of the CRRA Simulation

The next two tables provide more detail about the quality of the approximation, reporting the relative mean square approximation error (RMSAE) for various years.⁹ If we regard the anticipated-utility model as an approximation to the more complex Bayesian decision problem, the approximation error is $u = c^B - c^{AU}$. The RMSAE is defined as the ratio of mean-square error of u to the variance of c^B , and it is analogous to $1 - R^2$ in a regression. Thus, a small number signifies a good approximation.

Table 4: RMSAE for Consumption in the Certainty-Equivalent Examples

Period	5	15	25	35	45
<i>Business-Cycle Example</i>					
Excess Sensitivity	0.0063	0.0044	0.0015	0.0010	0.0006
Excess Smoothness	0.0002	0.0001	0.0001	0.0001	0.0001
Irrational Exuberance	0.0182	0.0046	0.0056	0.0067	0.0054
Depression Generation	0.0196	0.0042	0.0060	0.0066	0.0051
<i>Unit-Root Example</i>					
Too Persistent	0.0011	0.0004	0.0005	0.0005	0.0005
Not Persistent Enough	0.0005	0.0005	0.0002	0.0002	0.0004
Over-Estimate Positive State	0.0240	0.0027	0.0055	0.0081	0.0100
Over-Estimate Negative State	0.0213	0.0025	0.0051	0.0074	0.0098

Note: $\text{RMSAE} = \text{MSE}(c^B - c^{AU}) / \text{var}(c^B)$.

⁹Watson (1993) used a measure like this to assess approximation errors in DSGE models.

Table 5: RMSAE for Consumption in the *CRRA* Example

Period	5	10	15
$\alpha = 2$	0.0395	0.0064	0.0101
$\alpha = 5$	0.0197	0.0109	0.0402
$\alpha = 10$	0.0085	0.0261	0.0706
$\alpha = 20$	0.0933	0.0601	0.1235

Note: $\text{RMSAE} = \text{MSE}(c^B - c^{AU}) / \text{var}(c^B)$.

Measured on this scale, the approximation errors are often quite small, confirming the visual impression made by figures 1-3. In the certainty-equivalent simulations, all of the entries are less than 2.5 percent, and many are less than 1 percent. In the *CRRA* example, the quality of the approximation is also quite good when α is small, but it deteriorates a bit as α increases. For α equal to 2 or 5, the RMSAE is around 1 to 4 percent, but it increases to 6 to 12 percent when $\alpha = 20$. Even so, the bottom right panel of figure 3 suggests that a 6 percent RMSAE is still a pretty good approximation. These numbers suggest that an anticipated-utility approximation for consumption is excellent for typical calibrations in macroeconomics.

Not only is Bayesian consumption well approximated by anticipated-utility models, so too are security prices, provided that consumers are not too risk averse. Figures 4-6 depict outcomes for the Arrow security that pays off in the high-income state. In both of the certainty-equivalent simulations, the scatterplots indicate that Bayesian and anticipated-utility prices are tightly arrayed along the 45 degree line. Pricing errors are also quite small in the *CRRA* example when α is 2 or 5, but they grow in magnitude as α increases. Notice in particular that when $\alpha = 20$ the scatterplot is steeper than the 45 degree line, which means that the mean pricing error is negative. In this instance, an anticipated-utility model systematically overstates the price of an Arrow security in a Bayesian economy.

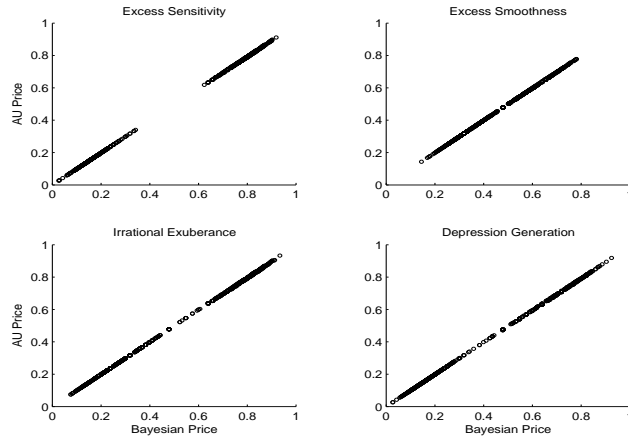


Figure 4: Arrow Prices in Period 25 of the Business-Cycle Example

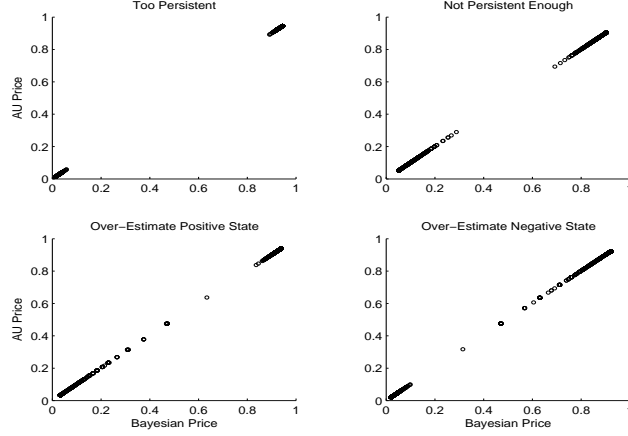


Figure 5: Arrow Prices in Period 25 of the Unit-Root Example

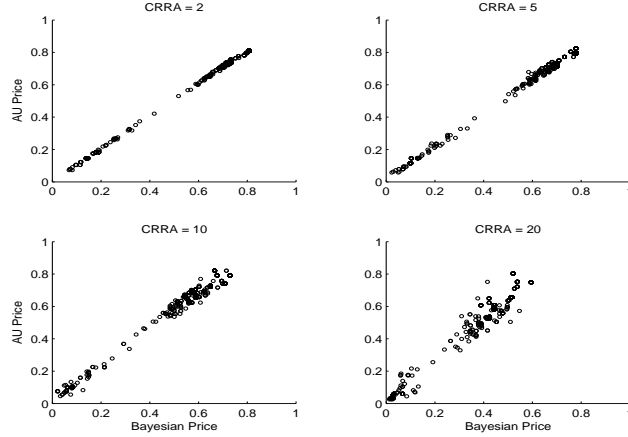


Figure 6: Arrow Prices in Period 10 of the CRRA Example

Tables 6 and 7 record the RMSAE for prices in various years. As before, the quality of the approximation is excellent for the certainty-equivalent simulations, indeed the RMSAEs for security prices are even lower than those for consumption. In these examples, the mean-square approximation error amounts to less than one-twentieth of one percent of the variance of Bayesian prices. Pricing errors are also quite small in *CRRA* simulations with small α . When α is 2 or 5, the RMSAEs are only slightly larger than in the certainty equivalent examples, but the quality of the approximation deteriorates as α increases. The RMSAEs grow as α increases to 10, and they become very large – sometimes more than 100 percent – when $\alpha = 20$. The large RMSAEs that occur in this case reflect the upward bias in AU prices that was mentioned above.

Table 6: RMSAE for Prices in the Certainty-Equivalent Examples

Period	5	15	25	35	45
<i>Business-Cycle Example</i>					
Excess Sensitivity	0.0003	0.0004	0.0004	0.0003	0.0001
Excess Smoothness	0.0001	0.0002	0.0002	0.0003	0.0001
Irrational Exuberance	0.0001	0.0002	0.0002	0.0002	0.0001
Depression Generation	0.0005	0.0005	0.0005	0.0004	0.0001
<i>Unit-Root Example</i>					
Too Persistent	<0.0001	<0.0001	<0.0001	<0.0001	<0.0001
Not Persistent Enough	<0.0001	<0.0001	<0.0001	<0.0001	<0.0001
Over-Estimate Positive State	0.0001	<0.0001	<0.0001	<0.0001	<0.0001
Over-Estimate Negative State	0.0001	<0.0001	<0.0001	<0.0001	<0.0001

Note: $\text{RMSAE} = \text{MSE}(Q^B - Q^{AU}) / \text{var}(Q^B)$.

Table 7: RMSAE for Prices in the *CRRA* Example

Period	5	10	15
$\alpha = 2$	0.0013	0.0020	0.0029
$\alpha = 5$	0.0126	0.0215	0.0332
$\alpha = 10$	0.0641	0.1818	0.3248
$\alpha = 20$	0.4038	1.0787	2.9262

Note: $\text{RMSAE} = \text{MSE}(Q^B - Q^{AU}) / \text{var}(Q^B)$.

The close correspondence of anticipated-utility and Bayesian prices in the certainty-equivalent and small- α cases follows from two facts, that consumption allocations are very similar across models and that one-step-ahead transition probabilities are identical. In each case, there are two possible outcomes for s_{t+1} given (s_t, n_t) , and along every sample path the one-step-ahead transition probabilities to these outcomes depend on the same counters n_t in the same way. Thus, the one-step-ahead transition probabilities always agree. Multi-step transition probabilities differ because Bayesian consumers update counters across potential future paths, while anticipated-utility consumers do not, but multi-step transition probabilities matter only indirectly for Arrow prices, affecting $Q(s_{t+1}, s_t)$ only through c_{t+1} . In the certainty-equivalent and small- α examples, the disparities in consumption are negligible, so Arrow prices are also in close agreement.

The consumption disparities are a bit larger in the large- α simulations, however, and they are magnified because inverse-consumption growth is raised to a large exponent when calculating security prices. Thus, seemingly minor discrepancies in approximating consumption can matter a lot for asset prices when consumers are highly risk averse.

Once again, the lesson seems to be that an anticipated-utility approximation is fine as long as consumers are not too risk averse. But it can be problematic for models with high degrees of risk aversion, especially when modeling security prices.

5 Rational Expectations v. Bayesian Outcomes

For the sake of comparison, we also consider a rational-expectations approximation to the Bayesian economy. In a rational-expectations approximating model, consumers are endowed with knowledge of the true transition probabilities from the outset, so a rational-expectations approximation abstracts from learning.

Results for consumption are reported in figures 7-9 and in tables 8 and 9. In many cases, the quality of the approximation is still quite good. RE consumption is still highly correlated with outcomes in the Bayesian economies, and in many instances the RMSAEs are also quite low, sometimes around 15 percent or less. But, not surprisingly, the correspondence is not as close as for the anticipated-utility approximation. In addition, the rational-expectations approximation is not uniformly reliable. Sometimes there are substantial approximation errors, accounting for as much as three-quarters of the total variation in Bayesian consumption in the certainty-equivalent and small- α simulations and for more than 100 percent of the variance when α is larger.




Figure 7: Consumption in Period 25 of the Business-Cycle Simulation




Figure 8: Consumption in Period 25 of the Unit-Root Simulation




Figure 9: Consumption in Period 10 of the *CRRA* Simulation

Table 8: RMSAE for Consumption in the Certainty-Equivalent Examples

Period	5	15	25	35	45
<i>Business-Cycle Example</i>					
Excess Sensitivity	0.4179	0.1876	0.1194	0.1643	0.2232
Excess Smoothness	0.1423	0.0673	0.0368	0.0794	0.1770
Irrational Exuberance	0.8388	0.1474	0.2365	0.5363	0.7287
Depression Generation	0.8266	0.1508	0.2748	0.5452	0.7786
<i>Unit-Root Example</i>					
Too Persistent	0.0780	0.0333	0.0391	0.0671	0.1015
Not Persistent Enough	0.0272	0.0261	0.0116	0.0099	0.0206
Over-Estimate Positive State	0.2608	0.0296	0.0668	0.1598	0.2720
Over-Estimate Negative State	0.2557	0.0266	0.0639	0.1541	0.2763

Note: $\text{RMSAE} = \text{MSE}(c^B - c^{RE}) / \text{var}(c^B)$.

Table 9: RMSAE for Consumption in the *CRRA* Example

Period	5	10	15
$\alpha = 2$	0.3051	0.7271	2.2731
$\alpha = 5$	0.2742	0.7891	2.6728
$\alpha = 10$	0.2949	1.0671	3.1396
$\alpha = 20$	0.7945	1.0179	3.3879

Note: $\text{RMSAE} = \text{MSE}(c^B - c^{RE}) / \text{var}(c^B)$.

The reason why the RMSAE can be so high despite the high correlation is that the scatterplots are not always arrayed along the 45 degree line. This means that there are systematic discrepancies between RE and Bayesian outcomes. For example, in two of the business cycle cases (Irrational Exuberance and Depression Generation), the RMSAE of consumption exceeds 50 percent in years 35 and 45. This reflects the consequences for wealth of expectations errors made early in life by Bayesian agents. Those who were irrationally exuberant saved too little early on and had to cut back later when their hopes were not realized. In the rational expectations model, consumers were more realistic and saved more early in life, so they could consume more later on. If we regard the Bayesian model as true and the rational expectations model as an approximation to it, this means the rational expectations model systematically overstates consumption later in life. Rational expectations outcomes are still highly correlated with Bayesian outcomes, but the mean approximation error contributes to a higher RMSAE. The same phenomenon occurs in the Depression-Generation example, except that Bayesian agents are too pessimistic in that case, and the mean approximation error has the opposite sign.

This bias is especially pronounced in the *CRRA* simulation because of how model uncertainty and pessimism affect the slope of the life-cycle consumption path. In the Bayesian economy, three factors contribute to precautionary saving, viz. *CRRA* preferences, model uncertainty, and initial pessimism. Consumers in the rational-expectations approximating model have the same preference, but they are neither pessimistic nor uncertain about the transition probabilities, so they engage in less precautionary saving. Figure 10 illustrates the magnitude of the problem. Solid lines portray Bayesian outcomes, and dashed lines illustrate the rational-expectations approximation. Because the rational-expectations approximation abstracts from model uncertainty and pessimism, it understates the amount of precautionary savings early in life and therefore also understates the amount of consumption later on. This flattens the life-cycle consumption profile relative to that of Bayesian consumers and biases the approximation to consumption, first upward and then downward. These biases contribute to the high RMSAEs recorded in table 9. The anticipated-utility approximation – depicted by solid-dotted lines – does a better job fitting these features of the Bayesian economy.

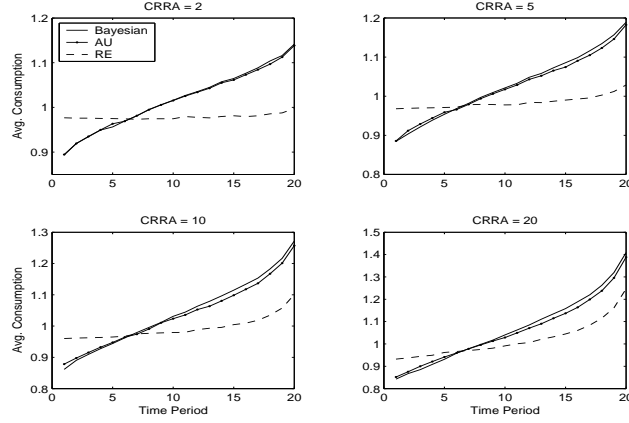


Figure 10: Average Life-Cycle Consumption Profile in the CRRA Simulation

The RE approximation also gives a somewhat misleading impression about Arrow security prices. Figures 11-13 compare Arrow prices from Bayesian economies with simulations of RE approximating models. Bayesian and RE outcomes are highly correlated (i.e., RE prices are systematically above average when Bayesian prices are), but notice how the RE models predict essentially only two values for prices, while the Bayesian economy predicts many. This feature is especially sharp in the certainty-equivalent and small- α simulations, though somewhat more diffuse in the large- α case.

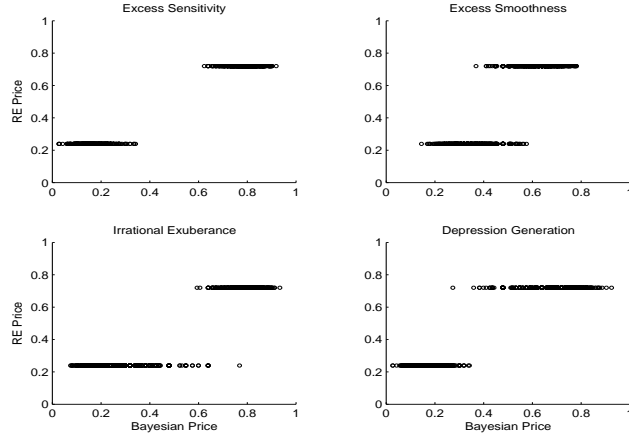


Figure 11: Arrow Prices in Period 25 of the Business-Cycle Simulation




Figure 12: Arrow Prices in Period 25 of the Unit-Root Simulation




Figure 13: Arrow Prices in Period 10 of the *CRRA* Simulation

In the certainty-equivalent and small- α cases, the existence of essentially two prices under rational expectations follows from three facts, that labor income evolves as a two-state process, that the true transition probabilities are known, and that consumers' intertemporal marginal rate of substitution does not vary much. Consider the Arrow security that pays off in the high-income state. The upper and lower branches in the scatterplots represent $Q(s_{t+1} = s_h, s_t = s_h)$ and $Q(s_{t+1} = s_h, s_t = s_l)$, respectively. From the household's first-order condition, these prices are given by

$$Q(s_{t+1} = s_h, s_t = i) = \beta \frac{u'[c_{t+1}(s_{t+1} = s_h)]}{u'[c_t(s_{t+1} = i)]} pr(s_{t+1} = s_h | s_t = i). \quad (58)$$

The linear storage technology provides a powerful tool for consumption smoothing, so the IMRS does not vary much across states, always staying fairly close to $1/R = \beta$

in the certainty-equivalent models and not straying too far from there in the small- α case. It follows that Arrow prices are well approximated by

$$Q(s_{t+1} = s_h, s_t = i) \doteq \beta pr(s_{t+1} = s_h | s_t = i). \quad (59)$$

Because the two transition probabilities are constant under rational expectations, it follows that there are essentially just two prices. The price is higher when $s_t = s_h$, reflecting a high probability that $s_{t+1} = s_h$ when it starts there, and the price is lower when $s_t = s_l$, reflecting a low probability of making a transition to the high-income state.

This two-value representation mirrors in a rough way what goes on in a Bayesian economy. There it is also the case that $Q(s_{t+1} = s_h, s_t)$ tends to be higher when $s_t = s_h$, again reflecting that prices are higher for securities promising a payoff with higher probability. Much of the variation in prices in the Bayesian economy represents movements across branches, and the rational expectations model captures this feature of the data. Cross-branch movements in prices account for the high correlation between the Bayesian and rational-expectations economies.

But there is an additional source of variation in the Bayesian economy that the rational expectations model neglects. In addition to variation across branches, there is also price variation within branches arising from the updating of transition probabilities. A version of equation (59) holds for the Bayesian economy as well,¹⁰ but there are many more nodes in the expanded Bayesian state space, hence many more possible values for the transition probabilities. Accordingly, there is price uncertainty in the Bayesian model even conditional on knowing the state of income today and tomorrow; prices depend on the counters as well. Because the rational expectations approximation abstracts from this source of uncertainty, the RMSAE for within-branch variation exceeds 100 percent.¹¹

This two-value characterization of prices begins to break down as α increases because higher α magnifies variation in the IMRS. In that case, we get many values for prices under rational expectations as well, and two clouds emerge in the scatterplots instead of two branches. But it remains true that within each cloud there is more variation in Bayesian prices than in RE prices, again reflecting that variation in transitions probabilities in Bayesian economies magnifies asset price volatility.

Whether abstracting from learning is critical for modeling security prices depends on the relative importance of variation within and across branches in the Bayesian economy. In many instances, the overall quality of the rational-expectations approximation is not too bad. For example, many of the RMSAEs reported in tables 10 and 11 are 15 percent or less, including all of those in the unit-root example. But the neglected source of variation contributes to higher values of the RMSAE, and in

¹⁰The IMRS varies more in the Bayesian economies, but not a lot more.

¹¹The RMSAE equals 1 if mean Bayesian prices within each branch are the same as the rational expectations value, and it exceeds 1 if the mean error is nonzero.

a few cases the RMSAE is more than 100 percent. Thus, the quality of the rational-expectations approximation for prices is mixed: often it is quite good, but sometimes it is very wide of the mark.

Table 10: RMSAE for Prices in the Certainty-Equivalent Examples

Period	5	15	25	35	45
<i>Business-Cycle Example</i>					
Excess Sensitivity	0.1426	0.1018	0.0807	0.0681	0.0602
Excess Smoothness	3.4148	1.1667	0.7040	0.4218	0.3022
Irrational Exuberance	0.2808	0.1734	0.1432	0.1126	0.0933
Depression Generation	0.2660	0.2114	0.1378	0.1223	0.0889
<i>Unit-Root Example</i>					
Too Persistent	0.0033	0.0035	0.0033	0.0031	0.0029
Not Persistent Enough	0.0723	0.0368	0.0275	0.0260	0.0242
Over-Estimate Positive State	0.0786	0.0481	0.0308	0.0233	0.0155
Over-Estimate Negative State	0.0890	0.0497	0.0418	0.0261	0.0237

Note: $\text{RMSAE} = \text{MSE}(Q^B - Q^{RE}) / \text{var}(Q^B)$.

Table 11: RMSAE for Prices in the *CRRA* Example

Period	5	10	15
$\alpha = 2$	0.5875	0.3621	0.2774
$\alpha = 5$	0.8297	0.4661	0.3027
$\alpha = 10$	1.2778	0.7602	0.4253
$\alpha = 20$	8.5083	1.5293	1.0708

Note: $\text{RMSAE} = \text{MSE}(Q^B - Q^{RE}) / \text{var}(Q^B)$.

5.1 The market price of risk

The rational expectations approximation is also problematic for evaluating the market price of risk, at least if the estimates are interpreted narrowly as signifying risk aversion per se. Following Hansen and Jagannathan (1991 and 1997), we define the conditional price of risk as

$$\rho_t(m_{t+1}) = \frac{\sigma_t(m_{t+1})}{\mu_t(m_{t+1})}, \quad (60)$$

where m_{t+1} represents a stochastic discount factor, $\mu_t(\cdot)$ is a conditional mean, and $\sigma_t(\cdot)$ is a conditional standard deviation. Thus the market price of risk is the conditional coefficient of variation of the stochastic discount factor.

At least three prices of risk are relevant for our economies. If one were to survey the agents in our model and ask them about the price of risk, they would report calculations based on their own preferences and beliefs. Within a Bayesian equilibrium, asset returns and consumption satisfy

$$E_t^B(m_{t+1}R_{t+1}) = 1, \quad (61)$$

where $E_t^B(\cdot)$ represents an expectation taken with respect to Bayesian transition probabilities, $pr_B(s_{t+1} = j|S_t)$, and $m_{t+1} = u'(c_{t+1})/u'(c_t)$ is the consumers' intertemporal marginal rate of substitution. Their subjective price of risk is accordingly

$$\rho_t^B(m_{t+1}) = \frac{\sigma_t^B(m_{t+1})}{\mu_t^B(m_{t+1})}, \quad (62)$$

where a superscript B indicates that conditional moments are evaluated with respect to Bayesian transition probabilities. In a Markov economy such as ours, the subjective price of risk can be expressed in terms of Arrow security prices and Bayesian probabilities. Appendix D shows how to do this and displays a formula for $\rho_t^B(m_{t+1})$.

Now imagine a rational-expectations modeler confronting data on prices and consumption from the Bayesian equilibrium. He has a model discount factor m_{t+1} , which in this case happens to be correctly specified, but he follows the precept of equating the perceived and actual laws of motion, and so he tries to make sense of observed prices in terms of the actual transition probabilities $pr_A(s_{t+1} = j|s_t)$.¹² This combination will not correctly price securities, for his model implies an Euler equation

$$E_t^A(m_{t+1}R_{t+1}) = 1 \quad (63)$$

that involves expectations with respect to different probabilities than those used by consumers. Since $E_t^A(\cdot) \neq E_t^B(\cdot)$, it follows that equilibrium consumption and asset returns will not conform to (63).

Following Hansen and Jagannathan (1991 and 1997), a rational-expectations modeler might summarize the discrepancy by calculating two other prices of risk, one associated with his model stochastic discount factor and another that characterizes the properties a stochastic discount factor must have in order to rationalize observed security prices with his probability model. We label the first $\rho_t^A(m_{t+1})$ and the second $\rho_t^A(\tilde{m}_{t+1})$. Both are calculated with respect to the rational-expectations probabilities $pr_A(s_{t+1} = j|s_t)$, hence the A superscript.

The RE model price of risk, $\rho_t^A(m_{t+1})$, is easy to calculate. One just substitutes observations on consumption into the model discount factor and then computes conditional moments under the assumed probability law. Appendix D displays an alternative calculation that is less straightforward but easier to implement using our simulation output.

¹²A rational expectations modeler can consistently estimate these transition probabilities ex post.

The RE required price of risk, $\rho_t^A(\tilde{m}_{t+1})$, is calculated from securities market data alone, without reference to consumption data or a model discount factor. Assuming absence of arbitrage opportunities, in a Markov economy with transition probabilities $pr_A(s_{t+1} = j|s_t)$, Arrow security prices must satisfy

$$Q_t(s_{t+1} = j|s_t) = \tilde{m}_{t+1}(s_{t+1} = j|s_t)pr_A(s_{t+1} = j|s_t), \quad (64)$$

for some discount factor \tilde{m}_{t+1} . Conditional moments for the unknown discount factor \tilde{m}_{t+1} can therefore be calculated from deflated security prices $Q_t(\cdot)/pr_A(\cdot)$. The prices $Q_t(\cdot)$ are observed in our economies, and the modeler specifies the probabilities $pr_A(s_{t+1} = j|s_t)$, so the properties of \tilde{m}_{t+1} can indeed be inferred from security market data. Once again, see appendix D for the details and a formula for $\rho_t^A(\tilde{m}_{t+1})$.¹³

The literature typically finds that the model price of risk $\rho_t^A(m_{t+1})$ is small if consumers are not too risk averse and that the required price of risk $\rho_t^A(\tilde{m}_{t+1})$ is larger. This is also what we find in our Bayesian economies. For example, figures 14-16 portray RE prices of risk for the middle periods of our simulations. The required price of risk $\rho_t^A(\tilde{m}_{t+1})$ is shown on the horizontal axis, and the model price of risk $\rho_t^A(m_{t+1})$ is on the vertical. The scatterplots are always far from the 45 degree line, and the required price of risk is almost always larger than the model price of risk. Differences between the two are especially pronounced in the certainty-equivalent and small- α simulations. In those cases, the required price of risk is often one or two orders of magnitude larger than the model price of risk. In the CRRA example, the model price of risk grows as α increases, but almost all the points remain below the 45 degree line, so the model price of risk still falls short of the required price of risk.

This is usually interpreted as a sign that the model discount factor is misspecified, but it could also signify that the transition probabilities are off the mark. Indeed, in this instance, we know that the model discount factor is correctly specified, so the discrepancies must be due to how the transition probabilities are modeled.

¹³Hansen and Jagannathan actually estimate a lower bound on the market price of risk. In our economies we can do a little better and estimate the price of risk itself. But this is just a detail. Following their example, we still compare a model price of risk with a required price of risk estimated from security market data alone.

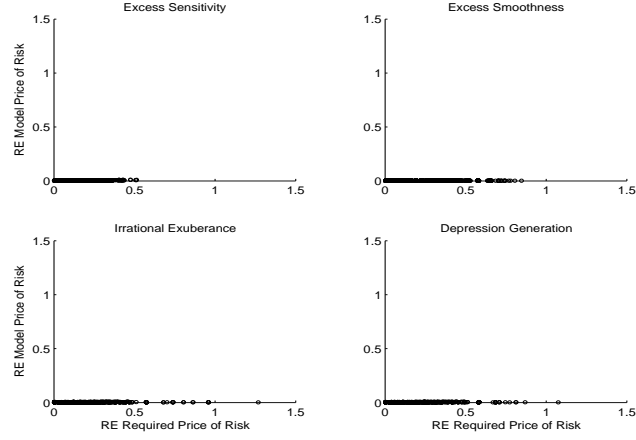


Figure 14: RE Prices of Risk in Period 25 of the Business-Cycle Example

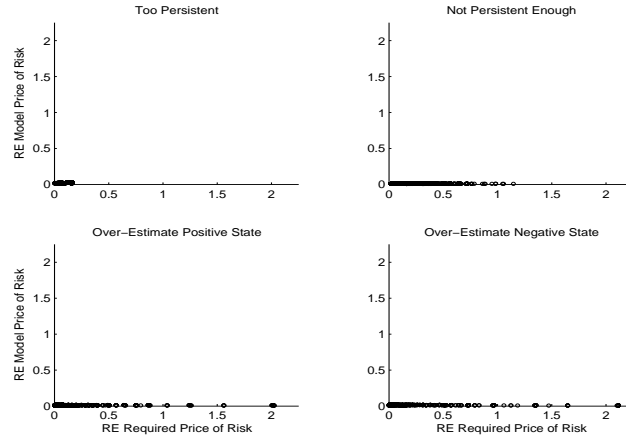


Figure 15: RE Prices of Risk in Period 25 of the Unit-Root Example

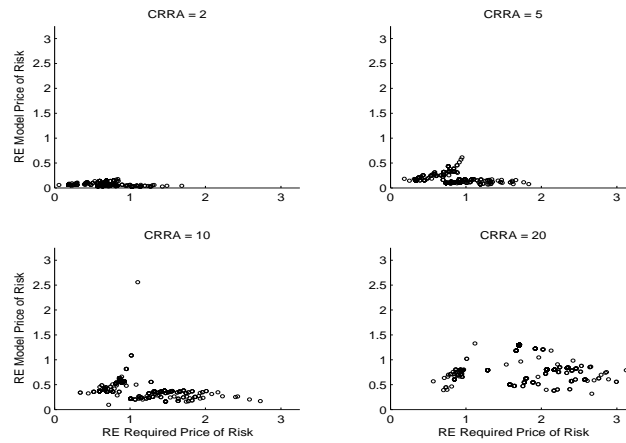


Figure 16: RE Prices of Risk in Period 10 of the *CRRA* Example

The next set of figures compares the RE required price of risk with the subjective measure $\rho_t^B(m_{t+1})$ that appraises the price of risk from the vantage of Bayesian consumers. The RE required price of risk $\rho_t^A(\tilde{m}_{t+1})$ is again plotted on the horizontal axis, and the consumers' price of risk $\rho_t^B(m_{t+1})$ is now shown on the vertical. Once again, the RE required price of risk is often higher by one or two orders of magnitude. If one regards the consumers' price of risk as the true value and the RE calculation as an estimate of it, then the figures can be interpreted as an assessment of the RE approximation. Many of the RMSAE statistics for this comparison are off the chart, on the order of $1.0\text{E}+02$ to $1.0\text{E}+06$. Notice also that the RE model price of risk, shown in the previous figures, is actually closer to subjective evaluations. The irony in this example is that the RE model price of risk, which is usually judged to be too low, is nearer to the truth. It is the required price of risk that is exaggerated.

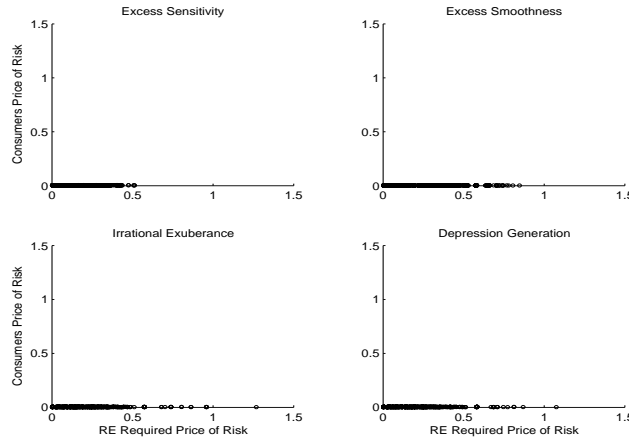


Figure 17: Consumers' and RE Prices of Risk in Period 25 of the Business-Cycle Example

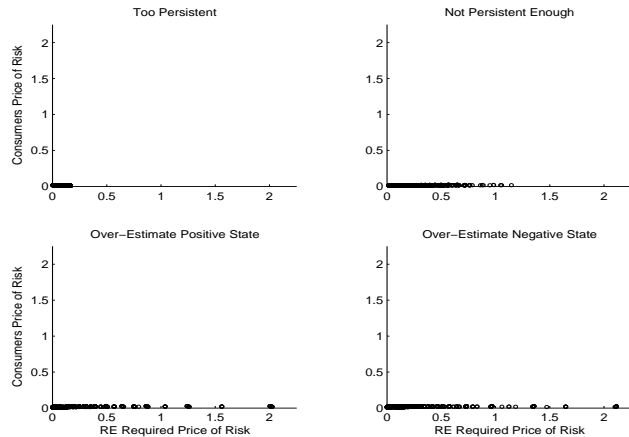


Figure 18: Consumers' and RE Prices of Risk in Period 25 of the Unit-Root Example

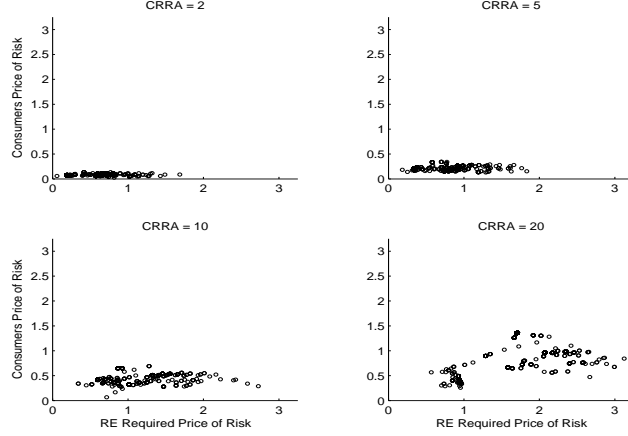


Figure 19: Consumers' and RE Prices of Risk in Period 10 of the *CRRA* Example

The reason why the RE estimate $\rho_t^A(\tilde{m}_{t+1})$ is so much larger is that it encompasses both risk aversion per se and model uncertainty. To see why, re-write the Bayesian first-order condition (61) in terms of rational-expectations probabilities,

$$\begin{aligned}
 1 &= E_t^B(m_{t+1}R_{t+1}), \\
 &= \sum_j m_{t+1}(s_{t+1} = j|S_t)R_{t+1}(s_{t+1} = j|S_t)pr_B(s_{t+1} = j|S_t), \\
 &= \sum_j \left[m_{t+1}(s_{t+1} = j|S_t) \frac{pr_B(s_{t+1} = j|S_t)}{pr_A(s_{t+1} = j|s_t)} \right] R_{t+1}(s_{t+1} = j|S_t)pr_A(s_{t+1} = j|s_t).
 \end{aligned} \tag{65}$$

The term in brackets is the modified discount factor \tilde{m}_{t+1} that reconciles Bayesian prices with RE transition probabilities. It involves the consumers' IMRS along with the Radon-Nikodým derivative of the Bayesian transition density with respect to the actual transition density. Variation in the IMRS reflects how averse consumers are to risk, while variation in the Radon-Nikodým derivative reflects their uncertainty about the law of motion for income. If the perceived and actual laws of motion were the same, the latter term would always equal 1, and its variance would be zero. But because consumers are uncertain about the law of motion and try to learn about it, their beliefs change over time, giving rise to variation in this probability ratio. Hence variation in $pr_B(\cdot)/pr_A(\cdot)$ reflects consumers' uncertainty about the right model for income.

Risk aversion and model uncertainty are both in play in the Bayesian economy, but our consumers are actually very risk tolerant. Their IMRS varies, but unless α is large it varies only a little. It follows that most of the variation in \tilde{m}_{t+1} recorded in $\rho_t^A(\tilde{m}_{t+1})$ arises from variation in the Radon-Nikodým derivative. Thus, the high price of risk required under rational expectations mostly reflects model uncertainty and changing beliefs; risk aversion makes only a small contribution.

Finally, notice that an anticipated-utility modeler trying to rationalize prices and quantities would draw the same conclusions about the conditional market price of

risk as a Bayesian modeler. This follows from the fact that one-step ahead transition probabilities in the anticipated-utility model coincide with those of a Bayesian model. Transition probabilities disagree over longer horizons but are the same for one-period forecasts. Since the one-step transition probabilities agree, the implicit stochastic discount factor \tilde{m}_{t+1}^{AU} is the same as the Bayesian discount factor m_{t+1} . In addition, conditional moments evaluated with respect to one set of probabilities are equivalent to those evaluated with respect to the other, so $\rho_t^B(\cdot) = \tilde{\rho}_t^{AU}(\cdot)$. It follows that the AU price of risk is identical to Bayesian price of risk,¹⁴

$$\rho_t^B(m_{t+1}) = \rho_t^{AU}(\tilde{m}_{t+1}^{AU}). \quad (66)$$

The AU approximation is a shortcut for calculating multi-step transition probabilities. Since multi-step transitions are not in play for calculating the conditional market price of risk, the AU approximation coincides with Bayesian calculations.

6 Conclusion

This paper represents a progress report on our research on anticipated utility. Our results are encouraging in three respects. First, the anticipated-utility approximation is excellent provided that precautionary motives are not too strong. This can be understood in terms of how an anticipated-utility model approximates a Bayesian predictive density. In our examples, the anticipated-utility model well-approximates the mean of the predictive density, but it neglects parameter uncertainty and therefore has tails that are too thin. Whether the quality of the approximation in the tails matters for decisions depends on the strength of precautionary motives. When decisions depend mostly on the mean, as in certainty-equivalent or small- α models, errors in approximating the tails matter hardly at all for decisions. The range of α over which we obtain good results covers values typically used in macroeconomics, so the approximation is likely to be very reliable for macroeconomic applications. The quality of the approximation deteriorates as α increases, however, so anticipated-utility modeling may be problematic for applications in finance when high risk aversion is assumed. But the full Bayesian analysis also turns out to be more tractable than we expected, making us optimistic that the methods can be generalized to more realistic applications when one has doubts about the anticipated-utility approach. Finally, the results on market prices of risk are tantalizing. In Cogley and Sargent (2004), we explore further how learning drives a wedge between subjective and RE prices of risk in the context of a more realistic asset pricing model.

¹⁴One can confirm this by inspecting the formulas in Appendix D. Re-interpret $pr_A(\cdot)$ as an anticipated-utility transition probability and then equate $pr_B(\cdot) = pr_A(\cdot)$. The equality of the prices of risk follows directly.

A A Dirichlet-Multinomial Model for a Multi-State Markov Process

In this example, we suppose that s_t can take on any of k distinct values. The transition probabilities are governed by a Markov transition matrix,

$$\Pi = \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1k} \\ p_{21} & p_{22} & \cdots & p_{2k} \\ \cdots & \cdots & \cdots & \cdots \\ p_{k1} & p_{k2} & \cdots & p_{kk} \end{bmatrix}, \quad (67)$$

where $p_{ij} = pr(s_{t+1} = j | s_t = i)$, the probability of moving from state i to state j . Once again, we assume the states are observable but that the transition probabilities are unknown.

The multi-state extension of the beta-binomial model is a Dirichlet prior and a multinomial likelihood. As before, we assume the agent has independent priors over the rows of Π ,

$$f(p_1, p_2, \dots, p_k) = \prod_{i=1}^k f(p_i), \quad (68)$$

where $p_i = [p_{i1}, p_{i2}, \dots, p_{ik}]$. The prior on each row is assumed to have the Dirichlet form,

$$f(p_i) \propto \prod_{j=1}^k p_{ij}^{n_0^{ij}-1}, \quad (69)$$

where n_0^{ij} records the prior number of transitions from state i to state j . The likelihood function for a batch of data s^t is proportional to the product of multinomial densities,

$$f(s^t | p_1, \dots, p_k) \propto \prod_{i=1}^k p_{i1}^{(n_t^{i1}-n_0^{i1})} p_{i2}^{(n_t^{i2}-n_0^{i2})} \cdots p_{ik}^{(n_t^{ik}-n_0^{ik})}, \quad (70)$$

where n_t^{ij} is the total number of transitions, prior plus observed, from state i to state j through date t . Multiplying the likelihood by the prior delivers the posterior kernel,

$$f(p_1, \dots, p_k | s^t) = \prod_{i=1}^k f(p_i | s^t), \quad (71)$$

where

$$f(p_i | s^t) \sim p_{i1}^{n_t^{i1}-1} p_{i2}^{n_t^{i2}-1} \cdots p_{ik}^{n_t^{ik}-1}. \quad (72)$$

The expression on the right side is a Dirichlet kernel, and its normalizing constant can be calculated as

$$\int \cdots \int p_{i1}^{n_t^{i1}-1} p_{i2}^{n_t^{i2}-1} \cdots p_{ik}^{n_t^{ik}-1} dp_{i1} \cdots dp_{ik} = \frac{\prod_{j=1}^k \Gamma(n_t^{ij})}{\Gamma(\sum_{j=1}^k n_t^{ij})}, \quad (73)$$

(see Sobel, et. al. 1977). Therefore, each term of the posterior density is

$$f(p_i|s^t) = \frac{\Gamma(\sum_{j=1}^k n_t^{ij})}{\prod_{j=1}^k \Gamma(n_t^{ij})} \prod_{j=1}^k p_{i1}^{n_t^{ij}-1}. \quad (74)$$

A Dirichlet prior and a multinomial likelihood form a conjugate pair. With independent Dirichlet priors over the rows of Π , Bayes's law implies that the posterior is also independent across the rows of Π and that each row is distributed as a Dirichlet random vector.

Next we derive the posterior predictive density, $f(s_{t+1}|s_t, n_t)$. This function assigns a probability to each of the k possible values of s_{t+1} for given values of s_t and n_t . There are k^2 possibilities,

$$pr(s_{t+1} = j|s_t = i, n_t), \quad i = 1, \dots, k, \quad j = 1, \dots, k. \quad (75)$$

As before, one can express each predictive probability by demarginalizing with respect to the true but unknown transition probabilities and then integrating across plausible values implied by the posteriors:

$$\begin{aligned} pr(s_{t+1} = j|s_t = i, n_t) &= \int pr(s_{t+1} = j, p_{ij}|s_t = i, n_t) dp_{ij}, \\ &= \int pr(s_{t+1} = j|s_t = i, n_t, p_{ij}) f(p_{ij}|s_t, n_t) dp_{ij}, \\ &= \int pr(s_{t+1} = j|s_t = i, p_{ij}) f(p_{ij}|s_t, n_t) dp_{ij}, \\ &= \int p_{ij} f(p_{ij}|s_t, n_t) dp_{ij}. \end{aligned} \quad (76)$$

The third equality follows from the fact that the counters are redundant given knowledge of the true transition probability, and the fourth reflects the definition of p_{ij} , i.e. $pr(s_{t+1} = j|s_t = i, p_{ij}) = p_{ij}$.

Recall that each row of Π is a Dirichlet random vector. Here the relevant beliefs concern the marginal distribution over a single element of row i , not the joint distribution over the entire row. But we can exploit the fact that the marginal distribution over a single element of a Dirichlet vector is a beta random variable (e.g., see Gelman, et. al. p. 482). Thus, the posterior for p_{ij} can be expressed as

$$f(p_{ij}|s_t = i, n_t) = \frac{\Gamma(\sum_{h=1}^k n_t^{ih})}{\Gamma(n_t^{ij})\Gamma(\sum_{h \neq i} n_t^{ih})} p_{ij}^{n_t^{ij}-1} (1 - p_{ij})^{(\sum_{h \neq i} n_t^{ih})-1}. \quad (77)$$

The transition probability we are seeking can therefore be written as

$$\begin{aligned} pr(s_{t+1} = j|s_t = i, n_t) &= \\ &= \frac{\Gamma(\sum_{h=1}^k n_t^{ih})}{\Gamma(n_t^{ij})\Gamma(\sum_{h \neq i} n_t^{ih})} \int p_{ij} p_{ij}^{n_t^{ij}-1} (1 - p_{ij})^{(\sum_{h \neq i} n_t^{ih})-1} dp_{ij}. \end{aligned} \quad (78)$$

This is the posterior mean of a beta density. By following the same steps as in example 1, we can write this as

$$pr(s_{t+1} = j | s_t = i, n_t) = \frac{n_t^{ij}}{\sum_{h=1}^k n_t^{ih}} \equiv \hat{p}_t^{ij}. \quad (79)$$

The posterior predictive probability of moving from state i to state j is simply the number of transitions from i to j divided by the total number of times the system has visited state i .

This is the probability of moving across the natural states. We actually want the transition probability for the expanded state, $f(s_{t+1}, n_{t+1} | s_t, n_t)$, which also reflects how the counters are updated. In each period, one element of n_t increases by 1, the others remaining constant. Which one is updated depends in a deterministic fashion on the s -transition. It follows that

$$pr(s_{t+1} = j, n_{t+1} = n_t + e_j | s_t = i, n_t) = \frac{n_t^{ij}}{\sum_{h=1}^k n_t^{ih}} \equiv \hat{p}_t^{ij}, \quad (80)$$

where e_j is a row vector with a 1 in column j and zeros everywhere else.

The complete list of \hat{p}_t^{ij} for $i = 1, \dots, k$ and $j = 1, \dots, k$ provides the building blocks for the transition matrix $P_{t,t+1}$. Again, imagine iterating through the rows of S_t and S_{t+1} . The transition probabilities can be found as follows.

- If $s_{t+1}^l = i$, $s_{t+1}^m = j$, and the n -update is consistent with the s -transition, set $P_{t,t+1}^{lm} = n_{lt}^{ij} / \sum_{h=1}^k n_{lt}^{ih}$,
- Otherwise, set $P_{t,t+1}^{lm} = 0$.

Having computed the sequence of transition matrices $\{P_{t,t+1}\}_{t=1}^{T-1}$ in this way, one can proceed to the dynamic program. The curse of dimensionality lurks here as well.

B Deriving the Chapman-Kolmogorov Equation

We want to express the m -step transition matrix

$$P_{t,t+m}(x, y) = pr[S_{t+m} = y | S_t = x]. \quad (81)$$

in terms of one-step transition matrices. Define the joint transition density

$$P(x, x_1, \dots, x_{m-1}, y) = pr[S_{t+m} = y, S_{t+m-1} = x_{m-1}, \dots, S_{t+1} = x_1 | S_t = x]. \quad (82)$$

This density spells out all the intermediate steps that take us from $S_t = x$ to $S_{t+m} = y$. Next, factor this into a conditional density for S_{t+m} and a marginal density for all the previous steps,

$$P(x, x_1, \dots, x_{m-1}, y) = pr[S_{t+m} = y | S_{t+m-1} = x_{m-1}, \dots, S_{t+1} = x_1, S_t = x] \quad (83) \\ \times pr[S_{t+m-1} = x_{m-1}, \dots, S_{t+1} = x_1 | S_t = x].$$

Because the states are Markov, the conditional density in the first line simplifies to

$$pr[S_{t+m} = y | S_{t+m-1} = x_{m-1}, \dots, S_{t+1} = x_1, S_t = x] = pr[S_{t+m} = y | S_{t+m-1} = x_{m-1}], \quad (84)$$

so that the entire expression becomes

$$P(x, x_1, \dots, x_{m-1}, y) = pr[S_{t+m} = y | S_{t+m-1} = x_{m-1}] \quad (85) \\ \times pr[S_{t+m-1} = x_{m-1}, \dots, S_{t+1} = x_1 | S_t = x].$$

Factor the joint density on the second line in the same way, and continue back to period $t + 1$. Eventually we find that the joint transition density (equation (82)) is the product of one-step transition densities,

$$P(x, x_1, \dots, x_{m-1}, y) = pr[S_{t+m} = y | S_{t+m-1} = x_{m-1}] \dots pr[S_{t+1} = x_1 | S_t = x]. \quad (86)$$

The m -step transition density (equation 81) can be found by marginalizing with respect to all the intermediate steps. Because this is a discrete-state model, marginalization is done by summation:

$$P_{t,t+m}(x, y) = \sum_{x_1} \dots \sum_{x_{m-1}} P(x, x_1, \dots, x_{m-1}, y), \quad (87) \\ = \sum_{x_1} \dots \sum_{x_{m-1}} pr[S_{t+m} = y | S_{t+m-1} = x_{m-1}] \dots pr[S_{t+1} = x_1 | S_t = x], \\ = \sum_{x_1} \dots \sum_{x_{m-1}} P_{t,t+1}(x, x_1) P_{t+1,t+2}(x_1, x_2) \dots P_{t+m-1,t+m}(x_m, y).$$

The notation $P_{h,h+1}(z_1, z_2)$ denotes the one-step transition probability associated with moving from state z_1 in period h to state z_2 in period $h + 1$. We need the h subscripts because the Markov chain is not homogenous; the transition matrix changes from period to period as the state space expands.

$P_{t,t+m}(x, y)$ is a particular element of the m -step transition matrix. The full set of transition probabilities can be expressed in matrix form as

$$P_{t,t+m} = P_{t,t+1} P_{t+1,t+2} \dots P_{t+m-1,t+m}, \quad (88)$$

where $P_{h,h+1}$ are the one-step transition matrices derived above. The i, j element of $P_{t,t+m}$ is the probability of moving from state i at time t to state j at time $t + m$. Equation (88) is known as the Chapman-Kolmogorov equation.

C How the State Space is Constructed

For the two-state models used in the examples, there are 4 counters, one for each possible transition of s_t . The expanded state space consists of all possible permutations of counters as well as the two possible values for the natural states. To construct this state space, we start with a large $k \times 4$ matrix in which each row represents a potential configuration of the counters. Initially, we include all combinations of integers from 1 to T , the time horizon of the model. In T periods, the economy can remain in the same state the entire time, so the counters for s_h to s_h and s_l to s_l transitions can in principle reach T . But the state cannot switch from s_h to s_l or s_l to s_h that many times. The maximum number of switches in T periods is $(T + 1)/2$, so we prune the original matrix by eliminating rows in which the counters for s_h to s_l or s_l to s_h exceed this limit. Call this matrix N_0 .

Next, we rearrange the rows of N_0 so that counters are grouped by date. Admissible counters for date t should sum to t , and the elements representing switches should not exceed $(t + 1)/2$. Once again, we prune elements of N_0 that violate this constraint. We label the resulting matrices N_{1t} .

Next we add the prior counters to N_{1t} to get a new set of counters N_{2t} . We also append s_h and s_l to N_{2t} , obtaining a state array

$$S_{0t} = \begin{bmatrix} s_h \cdot \iota & N_{2t} \\ s_l \cdot \iota & N_{2t} \end{bmatrix}, \quad (89)$$

where ι is a column vector of ones conformable with N_{2t} . Thus S_{0t} consists of every permutation of the counters and levels of income.

The matrix S_{0t} is a profligate representation of the state, however, containing many elements that the Markov process cannot actually reach. To identify redundant elements, we compute transition probabilities

$$P_{t,t+1}^0(i, j) = pr(S_{0t+1} = x_j | S_{0t} = x_i). \quad (90)$$

This is done by iterating through the rows and columns of S_{0t} and S_{0t+1} and looking for admissible matches. For $i = 1, 2$ and $j = 1, 2$, a match is admissible if $s_t = s_i$, $y_{t+1} = s_j$, the counter for i to j transitions increases by one, and the other counters remain unchanged. In that case, the transition probability is $n_t^{ij} / (n_t^{ii} + n_t^{ij})$. All other matches between S_{0t} and S_{0t+1} correspond to pairs in which the change in s is inconsistent with the change in n . Accordingly, their transition probabilities are zero.

After cycling through all the rows and columns of S_{0t} and S_{0t+1} , we have a transition matrix $P_{t,t+1}^0$. We identify redundant elements of S_{0t+1} by inspecting the columns of this matrix. If an element of S_{0t+1} cannot be reached from S_{0t} , the corresponding column of $P_{t,t+1}^0$ is zero. Accordingly, we eliminate rows of S_{0t+1} that correspond to columns of $P_{t,t+1}^0$ that sum to zero. The reduced state array is denoted S_{t+1} , and

the trimmed transition matrix connecting the reduced states S_t and S_{t+1} is denoted $P_{t,t+1}$.

These tree-trimming operations substantially reduce the dimension of the state space. For example, for a 50-period economy the matrix N_0 has more than 1.7 million rows. We can eventually reduce the dimension of the state to approximately 44,200.

D Arrow Prices and the Market Price of Risk in a Finite-State Markov Economy

The market price of risk is defined as the ratio of the conditional standard deviation of a stochastic discount factor to its conditional mean. Here we show how to express the market price of risk in terms of Arrow security prices.

Consider first the subjective market price of risk, which is associated with the preferences and beliefs of the consumers who inhabit in the models. In our Markov economies, the conditional mean is

$$\begin{aligned}\mu_t^B(m_{t+1}) &= \sum_{j=1}^k m_{t+1}(s_{t+1} = j|S_t) pr_B(s_{t+1} = j|S_t), \\ &= \sum_{j=1}^k Q_t(s_{t+1} = j|S_t) = \beta.\end{aligned}\tag{91}$$

Similarly, the conditional second moment is

$$\begin{aligned}E_t^B(m_{t+1}^2) &= \sum_{j=1}^k m_{t+1}^2(s_{t+1} = j|S_t) pr_B(s_{t+1} = j|S_t), \\ &= \sum_{j=1}^k \frac{Q_t^2(s_{t+1} = j|S_t)}{pr_B(s_{t+1} = j|S_t)}.\end{aligned}\tag{92}$$

It follows that the price of risk is

$$\begin{aligned}\rho_t^B(m_{t+1}) &= \frac{[E_t^B(m_{t+1}^2) - \mu_t^B(m_{t+1})^2]^{1/2}}{\mu_t^B(m_{t+1})}, \\ &= \frac{\left[\sum_{j=1}^k \frac{Q_t^2(s_{t+1}=j|S_t)}{pr_B(s_{t+1}=j|S_t)} - \left(\sum_{j=1}^k Q_t(s_{t+1} = j|S_t) \right)^2 \right]^{1/2}}{\sum_{j=1}^k Q_t(s_{t+1} = j|S_t)}.\end{aligned}\tag{93}$$

Because the risk-free bond price is constant and equal to β in our models, this expression further simplifies to

$$\rho_t^B(m_{t+1}) = \left[\sum_{j=1}^k \frac{Q_t^2(s_{t+1} = j|S_t)}{\beta^2 pr_B(s_{t+1} = j|S_t)} - 1 \right]^{1/2}.\tag{94}$$

Next, consider the model price of risk under rational expectations, $\rho_t^A(m_{t+1}) = \sigma_t^A(m_{t+1})/\mu_t^A(m_{t+1})$. Under the actual transition law, the conditional first and second moments are

$$\begin{aligned}\mu_t^A(m_{t+1}) &= \sum_{j=1}^k m_{t+1}(s_{t+1} = j|S_t) pr_A(s_{t+1} = j|s_t), \\ &= \sum_{j=1}^k m_{t+1}(s_{t+1} = j|S_t) \frac{pr_A(s_{t+1} = j|s_t)}{pr_B(s_{t+1} = j|S_t)} pr_B(s_{t+1} = j|S_t) \\ &= \sum_{j=1}^k Q_t(s_{t+1} = j|S_t) \frac{pr_A(s_{t+1} = j|s_t)}{pr_B(s_{t+1} = j|S_t)},\end{aligned}\tag{95}$$

and

$$\begin{aligned}E_t^A(m_{t+1}^2) &= \sum_{j=1}^k m_{t+1}^2(s_{t+1} = j|S_t) pr_A(s_{t+1} = j|s_t), \\ &= \sum_{j=1}^k [m_{t+1}(s_{t+1} = j|S_t) pr_B(s_{t+1} = j|S_t)]^2 \frac{pr_A(s_{t+1} = j|s_t)}{pr_B^2(s_{t+1} = j|S_t)}, \\ &= \sum_{j=1}^k \frac{Q_t^2(s_{t+1} = j|S_t)}{pr_B(s_{t+1} = j|S_t)} \frac{pr_A(s_{t+1} = j|s_t)}{pr_B(s_{t+1} = j|S_t)},\end{aligned}\tag{96}$$

respectively. It follows that the RE model price of risk is

$$\rho_t^A(m_{t+1}) = \frac{\left[\sum_{j=1}^k \frac{Q_t^2(s_{t+1}=j|S_t)}{pr_B(s_{t+1}=j|S_t)} \frac{pr_A(s_{t+1}=j|s_t)}{pr_B(s_{t+1}=j|S_t)} - \left(\sum_{j=1}^k Q_t(s_{t+1} = j|S_t) \frac{pr_A(s_{t+1}=j|s_t)}{pr_B(s_{t+1}=j|S_t)} \right)^2 \right]^{1/2}}{\sum_{j=1}^k Q_t(s_{t+1} = j|S_t) \frac{pr_A(s_{t+1}=j|s_t)}{pr_B(s_{t+1}=j|S_t)}}.\tag{97}$$

This is not the most straightforward way to calculate the model price of risk, and a rational expectations modeler would not adopt it because he would not know $pr_B(s_{t+1} = j|S_t)$. But we do know the Bayesian probabilities, and this formula happens to simplify our computations given what is available from our simulation output.

Finally, consider the rational-expectations calculation of the required price of risk, $\rho_t^A(\tilde{m}_{t+1}) = \sigma_t^A(\tilde{m}_{t+1})/\mu_t^A(\tilde{m}_{t+1})$. The conditional mean is

$$\begin{aligned}\mu_t^A(\tilde{m}_{t+1}) &= \sum_{j=1}^k \frac{Q_t(s_{t+1} = j|S_t)}{pr_A(s_{t+1} = j|s_t)} pr_A(s_{t+1} = j|s_t), \\ &= \sum_{j=1}^k Q_t(s_{t+1} = j|S_t) = \beta.\end{aligned}\tag{98}$$

Similarly, the conditional second moment is

$$\begin{aligned}E_t^A(\tilde{m}_{t+1}^2) &= \sum_{j=1}^k \left[\frac{Q_t(s_{t+1} = j|S_t)}{pr_A(s_{t+1} = j|s_t)} \right]^2 pr_A(s_{t+1} = j|s_t), \\ &= \sum_{j=1}^k \frac{Q_t^2(s_{t+1} = j|S_t)}{pr_A(s_{t+1} = j|s_t)},\end{aligned}\tag{99}$$

It follows that the market price of risk is

$$\begin{aligned}\rho_t^A(\tilde{m}_{t+1}) &= \frac{\left[\sum_{j=1}^k \frac{Q_t^2(s_{t+1}=j|S_t)}{pr_A(s_{t+1}=j|s_t)} - \left(\sum_{j=1}^k Q_t(s_{t+1}=j|S_t) \right)^2 \right]^{1/2}}{\sum_{j=1}^k Q_t(s_{t+1}=j|S_t)}, \\ &= \left[\sum_{j=1}^k \frac{Q_t^2(s_{t+1}=j|S_t)}{\beta^2 pr_A(s_{t+1}=j|s_t)} - 1 \right]^{1/2}.\end{aligned}\tag{100}$$

This depends on Arrow security prices, which are observed, and on actual transition probabilities, which a rational expectations modeler can consistently estimate. Hence the RE required price of risk can be calculated from asset price data without making assumptions about the form of the stochastic discount factor.

References

- Abel, A.B., 2002, "An Exploration of the Effects of Pessimism and Doubt on Asset Returns," *Journal of Economic Dynamics and Control* 26, 1075-1092.
- Cecchetti, S.G., P.S. Lam, and N.C. Mark, 2000, "Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to Be True?" *American Economic Review* 90, 787-805.
- Cogley, T., S. Morozov, and T.J. Sargent, 2003, "Bayesian Fan Charts for U.K. Inflation: Forecasting and Sources of Uncertainty in an Evolving Monetary System," forthcoming *Journal of Economic Dynamics and Control*.
- Cogley, T., and T.J. Sargent, 2004, "The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?" unpublished manuscript.
- Evans, G.W. and S. Honkapohja, 2001, *Learning and Expectations in Macroeconomics* (Princeton University Press: Princeton).
- Gelman, A., J.B. Carlin, H.S. Stern, and D.B. Rubin, 1995, *Bayesian Data Analysis* (Chapman and Hall: London).
- Hall, R.J., 1978, "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy* 86, pp. 971-988.
- Hansen, L.P. and R. Jagannathan, 1991, "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy* 99(2), pp. 225-262.
- and —, 1997, "Assessing Specification Errors in Stochastic Discount Factor Models," *Journal of Finance* 52, pp. 557-590.
- Judd, Kenneth, 1998, *Numerical Methods in Economics* (MIT Press: Cambridge MA).

Kreps, D., 1998, "Anticipated Utility and Dynamic Choice," 1997 Schwartz Lecture, in *Frontiers of Research in Economic Theory*, Edited by D.P. Jacobs, E. Kalai, and M. Kamien, Cambridge University Press, Cambridge, England.

Sargent, T.J., 1993, *Bounded Rationality in Macroeconomics*, (Clarendon Press: Oxford).

—, 1999, *The Conquest of American Inflation* (Princeton University Press: Princeton).

Sobel, M., R.R. Uppuluri, and K. Frankowski, 1977, *Selected Tables in Mathematical Statistics, Vol. 4: Dirichlet Distribution Type I* (American Mathematical Society: Providence RI).

Watson, M.W., 1993, "Measures of Fit for Calibrated Models," *Journal of Political Economy* 101, pp. 1011-1041.