

Computational Experiments and Reality

John Geweke

Departments of Economics and Statistics, University of Iowa
john-geweke@uiowa.edu

August 25, 2006

Abstract

This study explores three alternative econometric interpretations of dynamic, stochastic general equilibrium (DSGE) models. (1) A strong econometric interpretation takes the model literally and directly produces a likelihood function for observed prices and quantities. It is widely recognized that under this interpretation, most DSGE models are rejected using classical econometrics and assigned zero probability in a Bayesian approach. (2) A weak econometric interpretation commonly made in the calibration literature confines attention to only a few functions of observed prices and interest rates and evaluates a model on its predictive distribution for these functions. This approach is equivalent to a Bayesian prior predictive analysis, developed by Box (1980) and predecessors. This study shows that the weak interpretation retains the implications of the strong interpretation, and therefore DSGE's fare no better under this approach. (3) Under a minimal econometric interpretation, DSGE's provide only prior distributions for specified population moments. When coupled with an econometric model (e.g., a vector autoregression) that includes the same moments, DSGE's may be compared and used for inference using conventional Bayesian methods. This interpretation extends and formalizes an approach suggested by Dejong, Ingram and Whiteman (1996). All three interpretations are illustrated using models of the equity premium, and it is shown that the conclusions from a minimal interpretation differ substantially from those under a weak interpretation.

This revision was prepared for the DYNARE Conference, CEPREMAP, Paris, September 4-5, 2006. It is work in progress. Comments welcome. Please do not cite or quote without the author's permission.

1 Introduction

The dynamic, stochastic, general equilibrium (DSGE) model has become a central analytical tool in studying aspects of economic behavior in which aggregate uncertainty is important. Models in this family abstract sufficiently from measured economic behavior that clarification of the dimensions of reality they are intended to mimic is essential if they are to deepen our understanding of real economies. If the relation between DSGE models and measured economic behavior can be made formal, explicit, and simple, then the analytic power of this approach and our understanding of economic behavior will be enhanced. The objective of this study is to provide such a formal, explicit and simple characterization of the relation between DSGE models and measured economic behavior.

The approach taken here is to examine three alternative interpretations of the relationship. The first, called the strong econometric interpretation, leads to conventional, likelihood based, econometric methods. It is widely understood that DSGE models fare badly under this interpretation, and the DSGE literature consistently denies its appropriateness given the level of abstraction in the models. The second, called the weak econometric interpretation, greatly reduces the dimensions of observed behavior a DSGE model is designed to explain. It is the interpretation advanced by Kydland and Prescott (1996). Its assumptions are in fact no weaker than those that lead to likelihood based econometrics, and so DSGE models fare badly under this interpretation as well, although the failure is not so immediately evident. This study develops a third, minimal econometric interpretation of DSGE models introduced by DeJong, Ingram and Whiteman (1996). The assumptions underlying this interpretation are much weaker, and it is immune to the difficulties encountered in likelihood based econometrics. To be capable of explaining measured aggregate economic behavior, however, DSGE models under this interpretation must be married to econometric models that provide empirically plausible descriptions of measured behavior. This study shows how to do this in a way that is formal, explicit, simple, and easy to implement.

The three econometric interpretations are all presented with reference to a particular substantive application – DSGE models designed to explain the equity premium. The paper begins, in Section 2, by setting forth four such models. The alternative econometric interpretations are taken up in turn in Sections 3, 4, and 5. In each case numerical and graphical methods are used to illustrate the application to equity premium models. The weak econometric interpretation (Section 4) corroborates the findings of the DSGE literature regarding the equity premium puzzle – as it must, for this is the interpretation used there. The minimal econometric interpretation (Section 5) overturns some of the findings widely regarded as established by DSGE models. These findings illustrate some of the returns to a formal, explicit and simple approach to inference in DSGE models.

2 The essential elements of DSGE models

Dynamic, stochastic general equilibrium (DSGE) models have several common elements. They specify preferences of economic agents over alternative paths of consumption, a technology of production, and perhaps a government sector. They assume that all economic agents choose their most preferred path of consumption. They allow stochastic perturbations to the production technology. They use the principle of competition to determine equilibrium paths of quantities and prices, as functions of tastes, technology, and stochastic shocks. Tastes, technology, and the assumption of competition transform the technology shock distribution to a distribution of quantities and prices.

To isolate the econometrically relevant implications of these models, let “ A ” denote the assumptions of a particular model. For example, this could include the assumptions that preferences are time separable with constant relative risk aversion in each period, production is Cobb-Douglas, shocks to technology are log-normal and first order autoregressive, and equilibrium is competitive. Let θ_A denote the vector of parameters that provide quantitative content for the model – for example, the specification that labor’s share is .70, the coefficient of relative risk aversion is 2.0, and so on. Finally, let “ \mathbf{y} ” denote an observable, finite sequence of quantities and prices whose equilibrium values the model describes, for example 90 years of annual asset returns and output growth.

If the model has a unique equilibrium then it implies a distribution of \mathbf{y} , given the values of the parameters. A generic expression for this distribution is $p(\mathbf{y} \mid \theta_A, A)$. In most DSGE models, $p(\mathbf{y} \mid \theta_A, A)$ cannot be derived in closed form. However, it is typically not difficult to learn about $p(\mathbf{y} \mid \theta_A, A)$ by means of forward simulations: given a value of θ_A , pseudo-random vectors $\tilde{\mathbf{y}}$ can be drawn independently and repeatedly from $p(\mathbf{y} \mid \theta_A, A)$. In many cases, this ability to simulate is sufficient to draw formal conclusions about the model and use it to study the substantive questions it was designed to address.

2.1 An example: General equilibrium models of the equity premium

Average annual real returns on relatively riskless short-term securities in the U.S. have been about one percent during the past one hundred years. Annual real returns on equities over the same period have averaged above six percent. The equity premium—the difference between the return to equities and the return to relatively riskless short-term securities—has therefore exceeded five percent during the past century in the U.S. Many simple general equilibrium models predict average returns on riskless assets that are much higher than the observed average value, and average equity premia that are much lower, given parameter values generally regarded as reasonable. This predictive failure has become known as the equity premium puzzle. Kocherlakota

(1996) provides a review of the literature.

In the simplest general equilibrium model of the equity premium there is a single perishable good produced and consumed each period. Let period t production of the good be y_t , and denote the period-to-period gross growth rate of output by $x_t = y_t/y_{t-1}$. The representative agent orders preferences over random paths of consumption $\{y_t\}$ by

$$E_t \left[\sum_{s=0}^{\infty} \delta^s U(y_{t+s}) \right]. \quad (1)$$

In this expression $\delta \in (0, 1)$ is the subjective discount factor, and E_t denotes expectation conditional on time t information. The instantaneous utility function is the constant relative risk aversion (CRRA) utility function

$$U(c_t) = c_t^{1-\alpha} / (1-\alpha), \quad (2)$$

it being understood that $U(c_t) = \log(c_t)$ when $\alpha = 1$. The parameter α is the coefficient of relative risk aversion in the instantaneous utility function (2), and is also proportional to the marginal rate of intertemporal substitution in the preference ordering (1).

Define a riskless asset to be a claim to one unit of consumption in the next period. If such an asset is held in this economy, its period t price must be

$$p_t = \delta E [U'(y_{t+1}) / U'(y_t)] = \delta E_t (x_{t+1}^{-\alpha}). \quad (3)$$

Define one share of equity to be a claim to the fraction f of output in all future periods. If this asset is held in this economy, its period t price must be

$$\begin{aligned} q_t &= f \cdot E_t \left[\sum_{s=1}^{\infty} \delta^s U'(y_{t+s}) y_{t+s} / U'(y_t) \right] \\ &= f y_t E_t \left[\sum_{s=1}^{\infty} \delta^s U'(y_{t+s}) y_{t+s} / U'(y_t) y_t \right] = f y_t E_t \left[\sum_{s=1}^{\infty} \delta^s \prod_{j=1}^s x_{t+j}^{1-\alpha} \right] \end{aligned} \quad (4)$$

from which

$$q_t = \delta E_t [U'(y_{t+1}) (f y_{t+1} + q_{t+1}) / U'(y_t)] = \delta E_t [x_{t+1}^{-\alpha} (f y_{t+1} + q_{t+1})]. \quad (5)$$

From (4), the share price is proportional to output. If the growth rate x_t is stationary, then q_t/y_t is also stationary even though output y_t is not. This is a consequence of the assumption that instantaneous utility is of the CRRA form (2) – in fact, (2) is the unique instantaneous utility function with this property in (1) (King, Plosser and Rebelo, 1990).

2.2 The Mehra-Prescott and Rietz models

Mehra and Prescott (1985) assume that the growth rate x_t is a first order Markov chain with n discrete states. The growth rate is λ_j in state j . Assume that the time t information set includes the history of growth rates, and let P_t denote probability conditional on time t information. Then the Mehra-Prescott assumption can be expressed

$$P(x_{t+1} = \lambda_j \mid x_t = \lambda_i) = \phi_{ij}. \quad (6)$$

Suppose that this economy is in state i at time t . Then from (3) and (6), the price of the riskless asset is

$$p_t = p^{(i)} = \delta \sum_{j=1}^n \phi_{ij} \lambda_j^{-\alpha},$$

and the return to the riskless asset held from period t to period $t + 1$ is

$$r_{t+1} = r^{(i)} = 1/p^{(i)} - 1.$$

From (4) the share price q_t is proportional to output y_t . Hence for $x_t = \lambda_i$, denote the share price $q_t = w_i y_t$. Substituting in (5),

$$w_i = \delta \sum_{j=1}^n \phi_{ij} \lambda_j^{-\alpha} (f \lambda_j + w_j \lambda_j) = \delta \sum_{j=1}^n \phi_{ij} \lambda_j^{(1-\alpha)} (f + w_j) \quad (i = 1, \dots, n).$$

Solving this system of n linear equations for (w_1, \dots, w_n) yields the share prices $w_i y_t$. If $x_{t-1} = \lambda_i$ and $x_t = \lambda_j$, then the net return to equity holding from period $t - 1$ to period t is

$$s_t = s^{(i,j)} = \frac{q_t + f y_t - q_{t-1}}{q_{t-1}} = \frac{w_j y_t + f y_{t-1} \lambda_j - w_i y_{t-1}}{w_i y_{t-1}} = \frac{\lambda_j (w_j + f)}{w_i} - 1.$$

Mehra and Prescott (1985) take up the case $n = 2$, and restrict $\phi_{11} = \phi_{22} = \phi$. They choose $\lambda_1 = 1.054$, $\lambda_2 = 0.982$, and $\phi = 0.43$ to match the mean, standard deviation, and first order autocorrelation in the annual growth rate of per capita U.S. real consumption between 1889 and 1978. They then examine whether there are values of α less than 10 and any values of $\delta \in (0, 1)$ consistent with the observed average annual real returns of 0.0080 for short-term relatively riskless assets, and 0.0698 for the Standard and Poor's Composite Stock Price Index, over the same period. Their conclusion is negative.

Rietz (1988) uses the same model but adds a third state for output growth ($n = 3$). The third state occurs with low probability, the growth rate in this state is quite negative, and return to one of the two normal growth states occurs with certainty in the next period. Rietz concludes that this model is consistent with the observed average returns to riskless assets and the Stock Price Index, for some combinations of the parameter values — for example, α in the range of 5 to 7, δ above .98, and a probability of about 0.1% of a growth rate in which half of output is lost.

2.3 The Labadie and mixture models

Labadie (1989) takes

$$\log x_t = \beta_0 + \beta_1 \log x_{t-1} + \varepsilon_t, \quad \varepsilon_t \stackrel{iid}{\sim} N(0, \sigma^2).$$

Tsionas (1994) generalizes this to

$$\log x_t = \beta_0 + \beta_1 \log x_{t-1} + \omega_t^{1/2} \varepsilon_t, \quad \varepsilon_t \stackrel{iid}{\sim} N(0, 1), \omega_t \stackrel{iid}{\sim} p_\omega(\cdot),$$

with $\{\omega_t\}$ and $\{\varepsilon_t\}$ mutually independent. The riskless asset price follows from (3):

$$\begin{aligned} p_t &= \delta E_t(x_{t+1}^{-\alpha}) = \delta \exp[-\alpha(\beta_0 + \beta_1 \log x_t)] E_t \left[\exp \left(-\alpha \omega_{t+1}^{1/2} \varepsilon_{t+1} \right) \right] \\ &= \delta \exp(-\alpha \beta_0) x_t^{-\alpha \beta_1} E_t(\alpha^2 \omega_{t+1} / 2) = \delta \exp(-\alpha \beta_0) x_t^{-\alpha \beta_1} M(\alpha^2 / 2), \end{aligned}$$

where $M(\cdot)$ denotes the moment generating function of ω , $M(t) = E[\exp(\omega t)]$. The net return on holding the riskless asset is thus

$$r_t = \exp(\alpha \beta_0) x_t^{\alpha \beta_1} / \delta M(\alpha^2 / 2) - 1. \quad (7)$$

Direct substitution in (4) leads to

$$q_t / f y_t = \sum_{s=1}^{\infty} c_s m_s x_t^{\alpha_s}, \quad (8)$$

with

$$\begin{aligned} \alpha_s &= \rho \beta_1 (1 - \beta_1^s) / (1 - \beta_1), \\ c_s &= \exp \left[\frac{\rho s \beta_0}{1 - \beta_1} - \frac{\rho \beta_0 \beta_1 (1 - \beta_1^s)}{(1 - \beta_1)^2} \right], \\ m_s &= \delta^s \prod_{j=1}^s M \left[\rho^2 (1 - \beta_1^j)^2 / 2 (1 - \beta_1)^2 \right] \end{aligned}$$

where $\rho = 1 - \alpha$. The expression (8) converges, and equilibrium with finite equity prices exists, if and only if $M[\rho^2 / 2 (1 - \beta_1^2)] < \delta^{-1}$. Defining the left side of (8) to be h_t , the return to equity is then $s_t = x_t (h_t + 1) / h_{t-1} - 1$.

Tsionas (1994) thus extends Labadie (1989) by permitting the growth shock to be a scale mixture of normals. The best known scale mixture of normals is the Student- t distribution, corresponding to an inverted gamma mixing distribution $p_\omega(\cdot)$. However, the inverted gamma distribution has no moment generating function: the implicit integral on the left side of (7) diverges, and there is no equilibrium with finite asset prices (see Geweke (2001)). An attractive flexible family of symmetric distributions is the finite scale mixture of normals, for which the moment generating

function is trivial and always exists. A distribution in this family has n components, with component i assigned probability p_i . Conditional on component i , $\omega_t = \omega_{(i)}$ ($i = 1, \dots, n$). Thus, the p.d.f. of $u_t = \omega_t^{1/2} \varepsilon_t$ is

$$p(u_t) = (2\pi)^{-1/2} \sum_{i=1}^n p_i \omega_{(i)}^{-1/2} \exp(-u_t^2/2\omega_{(i)}).$$

To extend Labadie’s model in much the same way that Rietz extended Mehra and Prescott, let $n = 2$, let $i = 1$ denote the “normal” state, and let $i = 2$ denote the “high variance” state. We refer to this subsequently as the mixture model.

The Labadie and mixture models can be calibrated in the same way as the Mehra-Prescott and Rietz models. In the Labadie model choose β_0 , β_1 , and σ^2 to match the same three moments used by Mehra and Prescott: the mean, standard deviation, and first-order autocorrelation coefficient of U.S. consumption growth for 1889 through 1978. For the mixture model do the same thing, except to substitute $\omega_{(1)}$ for σ^2 . To parallel the treatment of Rietz, let $p_1 = 0.99$ and $p_2 = 0.01$ in the mixture model.

3 Strong econometric interpretation

The strong econometric interpretation of a dynamic, stochastic general equilibrium (DSGE) model is that the model provides a predictive distribution for an observable sequence of quantities and/or prices \mathbf{y} . Given the parameter values, $p(\mathbf{y} \mid \boldsymbol{\theta}_A, A)$ is the *ex ante*, predictive distribution for the observables \mathbf{y} . Then, letting \mathbf{y}^o denote the observed value of \mathbf{y} , $L(\boldsymbol{\theta}_A; \mathbf{y}^o, A) = p(\mathbf{y}^o \mid \boldsymbol{\theta}_A, A)$ is the likelihood function *ex post*. To provide such a predictive distribution the model must specify values of the parameters $\boldsymbol{\theta}_A$, or indicate a reasonable range for parameter values and a distribution over that range, not just a distribution conditional on an unknown parameter vector $\boldsymbol{\theta}_A$. Many studies that construct and calibrate DSGE models provide this sort of information about parameter values, at least informally, by means of reference to values in the literature and through their choice of calibrated values used in simulations.

Given a strong econometric interpretation, there is potentially a rich agenda of positive and normative economics. On the positive side, predictive distributions of observables can be used to compare alternative models, discard some, and improve others. On the normative side, DSGE models provide predictive distributions of the welfare implications of the policy interventions they are designed to accommodate.

3.1 Classical econometrics

Formal, classical approaches to model evaluation and comparison are squarely in the tradition of Neyman and Pearson, and have become known as specification testing. Let $t = t(\mathbf{y})$ be a scalar function of the observables, let $p(t \mid \boldsymbol{\theta}_A, A)$ be the density of t implied by $p(\mathbf{y} \mid \boldsymbol{\theta}_A, A)$, and let Θ_A denote the parameter space. If t is pivotal –

i.e., the distribution of t is the same for all $\theta_A \in \Theta_A$ – then one may unambiguously write the probability density function of t as $p(t | A)$ and the cumulative distribution function as $P(t | A)$. In this case, the model A implies that $1 - P(t | A)$ is uniformly distributed on the unit interval. Then $t^\circ = t(\mathbf{y}^\circ)$ is, formally, a test statistic, and $1 - P(t^\circ | A)$ is the marginal significance level, or p -value, of t° . It is then a short step to state that the model A is rejected at marginal significance level $1 - P(t^\circ | A)$. The specification A is rejected if $1 - P(t^\circ | A)$ is too small. More realistically t° or $1 - P(t^\circ | A)$ may simply be presented as a useful summary of evidence against model specification A .

Underlying this familiar use of a test statistic $t^\circ = t(\mathbf{y}^\circ)$ is the belief, perhaps unarticulated, that there is some plausible alternative specification $p(\mathbf{y} | \theta_B, B)$, $\theta_B \in B$, and that for at least some $\theta_B \in B$, $P(t^\circ | \theta_B, B) < P(t^\circ | A)$. It is this belief that makes it reasonable to reject A in a test of, say, size 5% if $1 - P(t^\circ | A) < 0.05$, as opposed to rejecting A if (say) $0.376 < 1 - P(t^\circ | A) < 0.426$. In formal terms, the power as well as the size of a test must be considered. But the concept of power requires an alternative to the model under consideration. Model evaluation cannot be divorced from model comparison: underlying any specification test is a hypothesis test against a class of alternative models.

The choice of the test statistic $t = t(\mathbf{y})$ typically involves both science and art. Finding a pivotal test statistic is possible only in special simple circumstances: Student's t for a single linear restriction in a linear model is the leading example. In most cases the best a classical econometrician can do is to establish that $t = t(\mathbf{y})$ is pivotal asymptotically. This technical requirement substantially limits the range of test statistics that can be considered. Determination of $P(t | A)$ in finite sample for pivotal t or in large sample for asymptotically pivotal t is then straightforward, for simulation can be used to create the necessary tables if an analytic derivation is not at hand. Especially in complex models, showing that $t = t(\mathbf{y})$ is asymptotically pivotal is the dominant technical concern. The choice of t within this class is more art: whether there are other tests more powerful than a given $t = t(\mathbf{y})$ is generally unknown.

3.2 Bayesian econometrics

The formal Bayesian approach to model comparison is grounded in complete specification of the joint distribution of models, parameters, and observables. If two models, A and B , are being considered, this complete specification can be achieved by stipulating a density $p(\theta_A | A)$ with support Θ_A , a density $p(\theta_B | B)$ with support Θ_B , and model probabilities $p(A)$ and $p(B)$ with $p(A) + p(B) = 1$. Then

$$p(A, B, \theta_A, \theta_B, \mathbf{y}) = p(A)p(\theta_A | A)p(\mathbf{y} | \theta_A, A) + p(B)p(\theta_B | B)p(\mathbf{y} | \theta_B, B). \quad (9)$$

Given (9), the evidence in the data about models and parameters follows from the

laws of conditional probability. For example, the marginal likelihood of model A is

$$p(\mathbf{y}^o | A) = \int_{\Theta_A} p(\boldsymbol{\theta}_A | A) p(\mathbf{y}^o | \boldsymbol{\theta}_A, A) d\boldsymbol{\theta}_A,$$

the posterior distribution of $\boldsymbol{\theta}_A$ is

$$p(\boldsymbol{\theta}_A | A) p(\mathbf{y}^o | \boldsymbol{\theta}_A, A) / p(\mathbf{y}^o | A),$$

and the posterior odds ratio in favor of model A is

$$\frac{p(A | \mathbf{y}^o)}{p(B | \mathbf{y}^o)} = \frac{p(A) p(\mathbf{y}^o | A)}{p(B) p(\mathbf{y}^o | B)}. \quad (10)$$

This approach generalizes immediately to any finite number of models.

The posterior odds ratio directly addresses the question so frequently asked in applied economics, “having looked at the data, what do we now think about alternative theories A and B ?” The clarity of (10) in answering this question reflect the demands of Bayesian econometrics for an explicit representation of beliefs about models $p(A)$, $p(B)$, $p(\boldsymbol{\theta}_A | A)$ and $p(\boldsymbol{\theta}_B | B)$. Technical work in this approach focuses on two tasks. One is the tractable representation of prior beliefs in densities like $p(\boldsymbol{\theta}_A | A)$, and the related problem of communicating information about posterior odds ratios to audiences in which prior beliefs may be heterogeneous. The second task is the computation of (10) and the evaluation of posterior moments of generic form

$$E[g(\mathbf{y}^o, \boldsymbol{\theta}_A, A) | \mathbf{y}^o, A] = \int_{\Theta_A} g(\mathbf{y}^o, \boldsymbol{\theta}_A) p(\boldsymbol{\theta}_A | \mathbf{y}^o, A) d\boldsymbol{\theta}_A,$$

which generally entails constructing a posterior simulator – an algorithm that produces pseudo-random vectors $\boldsymbol{\theta}_A^{(m)}$ with common density $p(\boldsymbol{\theta}_A | \mathbf{y}^o, A)$.

3.3 Failure of the DSGE equity premium models

In the Mehra-Prescott and Rietz variants of the DSGE equity premium models, the number of states is finite. Given n states, there can be at most n^2 different observable combinations of consumption growth and asset returns. Obviously this property does not characterize the data in any literal way. A trivial classical specification test would reject the model, and the marginal likelihood in any Bayesian formulation would be zero. Less obviously, and more substantively, the observed combinations of the risk free rate and equity premium have a very dispersed support, as illustrated in Figure 1 (a).

Similar problems with respect to the support of the distribution of observables arise in the Labadie and Tsionas variants of the model. Corresponding to any growth

rate there is exactly one riskless return (7), and to any two successive growth rates, one risky return based on (8). Thus, for example, these models imply

$$\min_{(b_0, b_1, b_2)} \sum_{t=1}^T (r_t - b_0 - b_1 x_t^{b_2}) = 0.$$

This condition is violated in the data. In any classical specification test the model is rejected in a test of any positive size, and in a Bayesian model comparison the marginal likelihood of the model is zero. The observed combinations of growth rates and risk free rates, displayed in Figure 1(b), do not even suggest such a relationship.

The restriction of observables to a degenerate space of lower dimension is a well documented failure of most DSGE models. Watson (1993), for example, has illustrated that the reduction in dimension is not even approximately true as a characterization of the data in the one-sector neoclassical model of King, Plosser, and Rebelo (1988). The problem derives from the small number of shocks—often just one, as is the case here—and the larger number of observables. Smith (1993) presents a simple real business cycle model with two shocks and two observables, and employs a formal, likelihood-based approach to make inferences about parameter values. Since observables are not restricted to a space of lower dimension, his model is not trivially rejected under the strong econometric interpretation. The difficulty lies not in the economics of dynamic general equilibrium, but in the fact that the technology of building this kind of model is not generally developed to the point of accommodating a sufficiently large number of shocks in a credible way. A strong econometric interpretation of DSGE models requires an explicit accounting for the dimensions of variation observed in the data, that are not accounted for in the model.

4 Weak econometric interpretation

Most macroeconomists who work with DSGE models eschew the strong econometric interpretation. For example, Mehra and Prescott (1985) in constructing their model of consumption growth, the riskless return, and the equity premium, plainly state that the model is intended to explain the first moments in returns, but not the second moments. That is, the model purports to account for sample average values of the riskless return and the equity premium, but not for the volatility in returns (Mehra and Prescott, 1985, p 146). Kydland and Prescott (1996, p 69) also emphasize that the model economy is intended to “mimic the world along a carefully specified set of dimensions.”

To begin the process of formalizing this interpretation of DSGE models, let $\mathbf{z} = f(\mathbf{y})$ denote the dimensions of the model that are intended to mimic the real world. In the DSGE literature such dimensions are typically sample moments – means, variances, autocorrelations, and the like. The weak econometric interpretation of a DSGE model is that the model provides a predictive distribution for the

functions $\mathbf{z} = f(\mathbf{y})$ of the observable, finite sequences of quantities and/or prices \mathbf{y} . This section argues that this is the interpretation most frequently given to DSGE models by macroeconomists, including Kydland and Prescott (1996), and that it is in fact a special case of prior predictive analysis, articulated at least as early as Box (1980). While many applications of DSGE models resort to *ad hoc* comparison of predictive distributions with observed behavior, careful investigators, including Kydland and Prescott (1996), use the weak econometric interpretation of DSGE models presented here. This section illustrates this interpretation in the context of the equity premium model introduced in Section 2. Finally, this section shows that this implementation of the weak econometric interpretation in fact makes the same assumptions as the strong econometric interpretation. A corollary is that to deny the strong econometric interpretation of a DSGE model, while examining the implications of the weak econometric interpretation in the manner of most of the applied DSGE literature, is internally inconsistent.

4.1 Formalizing the weak econometric interpretation

Given a complete, probabilistic specification of the model along the lines outlined in Section 3.2, $p(\mathbf{z} | \boldsymbol{\theta}_A, A)$ is implied by $p(\mathbf{y} | \boldsymbol{\theta}_A, A)$ and $\mathbf{z} = f(\mathbf{y})$. Hence there is a predictive density

$$p(\mathbf{z} | A) = \int_{\Theta_A} p(\boldsymbol{\theta}_A | A) p(\mathbf{z} | \boldsymbol{\theta}_A, A) d\boldsymbol{\theta}_A$$

for \mathbf{z} .

In the DSGE literature, this predictive density is typically investigated by means of simulation, or computational experiments. Often, $\boldsymbol{\theta}_A$ is fixed, or a few different values of $\boldsymbol{\theta}_A$ are considered to allow for uncertainty about $\boldsymbol{\theta}_A$. These are simply particular forms of the prior density $p(\boldsymbol{\theta}_A | A)$. Formally, a computational experiment is $\tilde{\boldsymbol{\theta}}_A^{(m)} \sim p(\boldsymbol{\theta}_A | A)$, $\tilde{\mathbf{y}}^{(m)} \sim p(\mathbf{y} | \boldsymbol{\theta}_A^{(m)}, A)$, $\tilde{\mathbf{z}}^{(m)} = f(\tilde{\mathbf{y}}^{(m)})$ ($m = 1, \dots, M$). The pseudo-random vectors $\tilde{\mathbf{z}}^{(m)}$ characterize the predictive distribution of the model, and can be compared with the observed value, \mathbf{z}^o .

Kydland and Prescott (1996, p 70) are quite clear about this process:

If the model has aggregate uncertainty. . . then the model will imply a process governing the random evolution of the economy. In the case of uncertainty, the computer can generate any number of independent realizations of the equilibrium stochastic process, and these relations, along with statistical estimation theory, are then used to measure the sampling distribution of any desired set of statistics of the model economy.

And, again (Kydland and Prescott, 1996, pp 75-76):

If the model economy has aggregate uncertainty, first a set of statistics that summarize relevant aspects of the behavior of the actual economy is selected. Then the computational experiment is used to generate many independent realizations of the equilibrium process for the model economy. In this way, the sampling distribution of this set of statistics can be determined to any degree of accuracy for the model economy and compared with the values of the set of statistics for the actual economy. In comparing the sampling distribution of a statistic for the model economy to the value of that statistic for the actual data, it is crucial that the same statistic be computed for the model and the real world. If, for example, the statistic for the real world is for a 50-year period, then the statistic for the model economy must also be for a 50-year period.

What aspects of the predictive density $p(\mathbf{z} | A)$, as represented by the $\tilde{\mathbf{z}}^{(m)}$, should be compared with \mathbf{z}^o ? In a classical approach, if \mathbf{z} were pivotal then $p(\mathbf{z} | A) = p(\mathbf{z} | \boldsymbol{\theta}_A^*)$ for any $\boldsymbol{\theta}_A^* \in \Theta_A$; the prior distribution for $\boldsymbol{\theta}_A$ would be irrelevant, and \mathbf{z}^o would be compared with the relevant tail of $p(\mathbf{z} | \boldsymbol{\theta}_A^*)$. In fact, however, the dimensions of the real world, \mathbf{z} , that the model is intended to describe are typically anything but pivotal. For example, in the equity premium model set forth in Section 2, the sample means of the riskless rate and equity premium that the model is intended to describe are sensitive to the risk aversion parameter α and discount parameter δ .

A formal Bayesian approach conditions on \mathbf{z}^o , the observed dimensions of the real world the model is intended to address. Given two competing models, A and B, the posterior odds ratio is then

$$\frac{p(A | \mathbf{z}^o)}{p(B | \mathbf{z}^o)} = \frac{p(A) p(\mathbf{z}^o | A)}{p(B) p(\mathbf{z}^o | B)}.$$

For purposes of model evaluation and comparison, therefore, it is the predictive density of \mathbf{z} at the observed value \mathbf{z}^o that matters. This has long been recognized in the Bayesian model evaluation literature. Box (1980, pp 385-386) describes exactly this approach, referring to f as a “relevant model checking function,” and traces its use in this fashion to Good (1956) and Roberts (1965) among others. Lancaster (2004, pp 79-81) and Geweke (2005, pp 262-267) describe computational methods for the systematic application of this approach, which they term prior predictive analysis, to model evaluation and comparison. The former text explicitly notes (p 81) that calibration as conventionally practiced is a special case of prior predictive analysis.

The key technical problem is the approximation of $p(\mathbf{z}^o | A)$ on the basis of the computational experiments that generated $\tilde{\mathbf{z}}^{(m)}$ ($m = 1, \dots, M$). Predictive means, $E(\mathbf{z} | A) \cong M^{-1} \sum_{m=1}^M \tilde{\mathbf{z}}^{(m)}$, which are so often presented and compared \mathbf{z}^o with in the DSGE macroeconomics literature, are irrelevant, under either a classical or a Bayesian approach.

If the order of the vector \mathbf{z} is small—say, three or less—then numerical approximation of $p(\mathbf{z}^o | A)$ is straightforward, and much simpler than the numerical ap-

proximation of the full marginal likelihoods in (10) under the strong econometric interpretation of the DSGE model. The latter requires backward simulation, for example by means of a Markov chain Monte Carlo algorithm. The former only requires the $\tilde{\mathbf{z}}^{(m)}$ produced through the forward simulation in the familiar computational experiments of the DSGE literature. Conventional smoothing procedures like kernel density methods will provide a numerical approximation to $p(\mathbf{z}^o | A)$:

$$p(\mathbf{z}^o | A) \cong M^{-1} \sum_{m=1}^M K(\tilde{\mathbf{z}}^{(m)}; \mathbf{z}^o),$$

where the kernel smoother $K(\mathbf{z}; \mathbf{z}^o)$ is a nonnegative function of \mathbf{z} that is concentrated near \mathbf{z}^o and integrates to one. This computational procedure also provides the numerical approximation

$$E[g(\boldsymbol{\theta}_A, \mathbf{y}^o) | \mathbf{z}^o, A] \cong \sum_{m=1}^M K(\tilde{\mathbf{z}}^{(m)}; \mathbf{z}^o) g(\boldsymbol{\theta}_A^{(m)}, \mathbf{y}^o) / \sum_{m=1}^M K(\tilde{\mathbf{z}}^{(m)}; \mathbf{z}^o)$$

to any posterior moment. This is also much less demanding than the posterior simulation required to approximate (3.2).

As a byproduct, the forward simulation exercise produces the full predictive distribution of \mathbf{z} , including points far from the observed value. Careful examination of these points can lead to further insights into the model.

4.2 Illustration in the equity premium model

The Mehra-Prescott and Labadie models completely specify the distribution of growth. The Rietz and mixture models each specify the distribution of growth up to a single, unknown parameter: λ_3 , growth in the event of a crash, and $\omega_{(2)}$, the high variance, respectively. For the Rietz model, adopt the prior distribution

$$\log[\lambda_3 / (1 - \lambda_3)] \sim N(.036, 1.185^2)$$

and for the mixture model use

$$1.33/\omega_{(2)} \sim \chi^2(1.66)$$

Deciles for both distributions are given in Table 1. The distribution of λ_3 is centered at $\lambda_3 = .509$, halving of expected output, which is the intermediate of the three examples taken up in Rietz (1988). The distribution of $\omega_{(2)}$ centers the standard deviation in the high variance state about 1.26, implying that in this state output is about as likely to be between one-third and triple its normal value as it is to be outside this range.

All other parameters in the models pertain to the consumption growth process. In the exercises reported here these parameters were held fixed at their calibrated

values, which are chosen to reproduce the mean, standard deviation, and first order autocorrelation of the consumption growth rate. Modifying the analysis by introducing prior distributions for these parameters increases the technical complexity of the exercise, because the dimension of the predictive distribution is increased from two to five, but should have little effect on the final results.

None of the models fix the relative risk aversion parameter α or the subjective rate of discount δ . This analysis employs priors that should provide substantial probabilities to the ranges most economists would regard as plausible. For α , take

$$\log \alpha \sim N (.4055, 1.3077^2),$$

and for δ ,

$$\log [\delta / (1 - \delta)] \sim N (3.476, 1.418^2).$$

Deciles for these prior distributions are also shown in Table 1. The prior distribution for α is centered at $\alpha = 1.5$, and a centered 80% prior credible interval for α is (0.281, 8.0). The prior distribution for δ is centered at 0.97, and a centered 80% prior credible interval is (0.84, 0.995).

These prior distributions, together with the data densities described in Section 2, provide predictive densities for all four models: Mehra-Prescott, Rietz, Labadie, and mixtures. For each model, draws from predictive densities for output growth rate and asset returns can be made by (1) drawing from the prior distributions of the unknown parameters; (2) conditional on the drawn parameters, generating a sample of 90 successive years of growth rates from the probability density for $\{x_i\}$; (3) solving for the riskless and risky returns in each year as indicated in Section 2. Draws from the predictive density for any function of output growth rate and asset returns, are then just the corresponding functions of this generated, synthetic sample. For example, to draw from the predictive density for 90-year means of the risk free rate and the equity premium, following step (3), just construct these functions and record them. Notice that the predictive density for the mean risk free rate and mean equity premium accounts for both uncertainty about parameter values (by means of the draws from the prior) and sampling variation due to 90-year averaging (by means of the 90-year simulation).

Figures 2-3 show the predictive distributions for 90-year averages of the risk free return and equity premium, as represented by 1,000 points $\tilde{\mathbf{z}}^{(m)}$ drawn from $p(\mathbf{z} | A)$ for each model. In each panel, the vertical line indicates the observed value of .008 for the risk free rate and the horizontal line indicates the observed value of .0618 for the equity premium. The supports of the Rietz and mixture model predictive densities include the observed values, but those of the Mehra-Prescott and Labadie model predictive densities do not. This is qualitative corroboration of the failure of the latter two models to explain the equity premium puzzle in the weak econometric DSGE literature, and the ability of the Rietz model to account for the observed means. It extends this corroboration to the mixture model.

Figures 2-3 can be used to explore the workings of all four models in some detail, by examining the parameter values or the history of consumption growth underlying the points plotted. This can be done by color coding the points for values of one of the parameters, like the relative risk aversion parameter α , or by software that provides a full display of parameter values and the simulated time series after clicking on a point. For example, high values of the risk free rate correspond to low values of δ . Negative values of the risk free rate and high equity premia in the Rietz and mixture models typically reflect high risk aversion in conjunction with a low probability of very negative growth rates. The values of the risk free rate and equity premium in the Rietz and mixture models close to the historical averages typically correspond to situations in which very negative growth rates were possible but did not occur during the simulated 90-year history.

Table 2 provides approximations of the log marginal likelihood, $\log [p(\mathbf{z}^o | A)]$, of each of the four models under the weak econometric interpretation. A independent symmetric bivariate Gaussian density kernel was centered at the observed sample mean for the riskless return and equity premium, and points were drawn from the predictive density $p(\mathbf{z} | A)$. The density kernel was symmetric. Various standard deviations were used as indicated in the left column of Table 2. As one moves down the rows, approximations show greater bias (because they include values of $\tilde{\mathbf{z}}^{(m)}$ farther from the data point) but less variance (because more points are given weight). Asymptotic standard errors are indicated parenthetically. Table 2 shows that the marginal likelihoods of the Mehra-Prescott and Labadie models are zero. The mixture model is favored over the Rietz model, the Bayes factor being about 3:1.

4.3 Logical problems with the weak econometric interpretation

As the DSGE literature emphasizes, all models are approximations of reality, and it is important to clarify which aspects of reality a model is intended to mimic. In the strong econometric interpretation of a model, this limited scope is recognized in the choice of the random vector \mathbf{y} . If, subsequently, attention is shifted to only a subset of the original variables, there are no conceptual difficulties: one simply works with the marginal distributions of the included variables.

The dimensions of reality addressed by DSGE models entail a reduction of a different kind. For example, the equity premium models are intended to explain sample means of the riskless return and equity premium, but no other aspects of these returns. This is not possible. If the model accounts for $(T + 1)$ -year averages as well as T -year averages, then the model also has implications for the year-to-year returns. One cannot choose to believe the former but not the latter.

More significantly, the DSGE calibration literature takes the short-run dynamics of these models literally, in establishing the sampling distribution of the set of statistics \mathbf{z} that summarize the relevant aspects of the behavior of the actual economy.

This fact is emphasized in Kydland and Prescott (1996). It is made quite clear in careful calibration studies, for example Gregory and Smith (1991, p 298) and Christiano and Eichenbaum (1992, pp 436 and 439). Bayesian prior predictive analyses of complete models also incorporate all of the dynamics of the model, and the use of the sampling distribution of \mathbf{z} in the DSGE calibration literature is equivalent to such an analysis.

There is an important distinction in context between the DSGE calibration literature and Bayesian prior predictive analyses, however. The latter treats the model as descriptive of all dimensions of reality, including the distribution as well as the first moment of \mathbf{z} . The former regards the model as predictive for the first moment but no other aspects of \mathbf{z} . The sampling distribution of the statistic \mathbf{z} is often a function of profoundly unrealistic aspects of DSGE models, aspects that lie outside the dimensions of reality the models were intended to mimic. For example, in the equity premium models the sampling distribution of average asset returns over the 90-year period are closely related to the variances of these returns, through the usual arithmetic for the standard deviation of a sample mean. In establishing the sampling distribution of these means through repeated simulation of the model, one is taking literally the second moments of returns inherent in the model. These are precisely the dimensions the original model was not intended to capture (Mehra and Prescott, 1985, p 146), and the models are unrealistic in these dimensions. For example, the sample standard deviation of the equity premium is .164 over the period 1889-1979, whereas at prior median values the standard deviation is .055 in the Mehra-Prescott model and .258 in the Rietz model. The weak econometric interpretation of DSGE models leads to formal methods for model comparison that are easy to implement and have an unambiguous interpretation. As a byproduct, there are some interesting and useful visual displays. But the assumptions that underlie the weak econometric interpretation are in fact the same as those made in the strong econometric interpretation: the model is assumed to account for all aspects of the observed sequence of quantities and/or prices.

5 Minimal econometric interpretation

The logical problems encountered in the claim that DSGE models account for only a few sample moments of observed sequences of quantities and prices prevents the development of this notion into coherent methods of inference about these models. To broaden the claim to assert that DSGE models in fact provide likelihood functions leads to outright dismissal of many of these models (Section 3). This section considers a more modest claim for DSGE models, also studied by DeJong, Ingram and Whiteman (1996): only that they account for population moments of specified, observable functions of sequences of prices and/or quantities. A DSGE model, A , with a given parameter vector $\boldsymbol{\theta}_A$, implies population moments $\mathbf{m} = E(\mathbf{z} \mid \boldsymbol{\theta}_A, A)$, where $\mathbf{z} = f(\mathbf{y})$ is the same vector of sample moments considered under the weak

econometric interpretation. If A is endowed with a prior distribution $p(\boldsymbol{\theta}_A | A)$, then A provides a distribution for \mathbf{m} as well.

By not claiming to predict sample moments, the minimal econometric interpretation avoids the logical pitfall that predicting sample moments and carrying out inference based on the model's predictive distribution leads inevitably back to a conventional likelihood function. The cost entailed in this retreat is that the DSGE model, by itself, now has no implications for anything that might be observed. To endow such a model with empirical content it is necessary to posit, separately, a link between the population moments \mathbf{m} and the observable sequence of prices and/or quantities \mathbf{y} . DeJong, Ingram and Whiteman (1996) also noted the need for such a link. This section shows how to do this formally, and provides some examples of the procedure. The result is an integration of atheoretical, data-based macroeconomic models with DSGE models.

5.1 Formal development

Let A and B denote two alternative DSGE models, each describing the same vector of population moments \mathbf{m} by means of the respective densities $p(\mathbf{m} | A)$ and $p(\mathbf{m} | B)$. The densities could be degenerate at a point but in general are not because of subjective uncertainty about parameter values in both models.

Introduce a third, econometric model E , that specifies a conditional distribution of observables $p(\mathbf{y} | \boldsymbol{\theta}_E, \mathbf{m}, E)$ together with a proper conditional prior distribution $p(\boldsymbol{\theta}_E | \mathbf{m}, E)$. Model E is incomplete because it provides no prior distribution $p(\mathbf{m} | E)$: in this sense it may be said to be atheoretical. The prior distribution for \mathbf{m} is provided either by model A through $p(\mathbf{m} | A)$ or by model B through $p(\mathbf{m} | B)$. Thus

Condition 1 *Conditional on the DSGE model A and the econometric model E ,*

$$p(\mathbf{m}, \boldsymbol{\theta}_E, \mathbf{y} | A, E) = p(\mathbf{m} | A) \cdot p(\boldsymbol{\theta}_E | \mathbf{m}, E) \cdot p(\mathbf{y} | \boldsymbol{\theta}_E, \mathbf{m}, E).$$

Conditional on the DSGE model B and the econometric model E ,

$$p(\mathbf{m}, \boldsymbol{\theta}_E, \mathbf{y} | B, E) = p(\mathbf{m} | B) \cdot p(\boldsymbol{\theta}_E | \mathbf{m}, E) \cdot p(\mathbf{y} | \boldsymbol{\theta}_E, \mathbf{m}, E).$$

The sole function of the DSGE is to provide a prior distribution for \mathbf{m} . We have immediately

Proposition 2 *Given Condition 1,*

$$\begin{aligned} p(\mathbf{y} | \mathbf{m}, A, E) &= \frac{\int p(\mathbf{m}, \boldsymbol{\theta}_E, \mathbf{y} | A, E) d\boldsymbol{\theta}_E}{p(\mathbf{m} | A, E)} \\ &= \frac{\int p(\mathbf{m} | A) p(\boldsymbol{\theta}_E | \mathbf{m}, E) p(\mathbf{y} | \boldsymbol{\theta}_E, \mathbf{m}, E) d\boldsymbol{\theta}_E}{p(\mathbf{m} | A)} \\ &= \int p(\boldsymbol{\theta}_E | \mathbf{m}, E) p(\mathbf{y} | \boldsymbol{\theta}_E, \mathbf{m}, E) d\boldsymbol{\theta}_E = p(\mathbf{y} | \mathbf{m}, E) \end{aligned}$$

and $p(\mathbf{y} \mid \mathbf{m}, B, E) = p(\mathbf{y} \mid \mathbf{m}, E)$, as well.

From Condition 1 and Proposition 2 we have

Proposition 3 *Given Condition 1,*

$$\frac{p(A \mid \mathbf{y}^o, E)}{p(B \mid \mathbf{y}^o, E)} = \frac{p(A \mid E) p(\mathbf{y}^o \mid A, E)}{p(B \mid E) p(\mathbf{y}^o \mid B, E)} = \frac{p(A \mid E) \int p(\mathbf{m} \mid A) p(\mathbf{y}^o \mid \mathbf{m}, E) d\mathbf{m}}{p(B \mid E) \int p(\mathbf{m} \mid B) p(\mathbf{y}^o \mid \mathbf{m}, E) d\mathbf{m}}. \quad (11)$$

The evidence about models A and B , in the context of the econometric model E , is in the convolutions

$$\int p(\mathbf{m} \mid A) p(\mathbf{y}^o \mid \mathbf{m}, E) d\mathbf{m} \text{ and } \int p(\mathbf{m} \mid B) p(\mathbf{y}^o \mid \mathbf{m}, E) d\mathbf{m}. \quad (12)$$

whose ratio in (11) is the Bayes factor in favor of model A . Loosely speaking, if $p(\mathbf{y}^o \mid \mathbf{m}, E)$ overlaps more with $p(\mathbf{m} \mid A)$ than with $p(\mathbf{m} \mid B)$, the Bayes factor favors model A . This looser interpretation underlies the confidence interval criterion proposed in DeJong, Ingram and Whiteman (1996), for univariate m . The odds ratio (11) provides an exact interpretation, that also extends to multivariate \mathbf{m} .

The expression $\int p(\mathbf{m} \mid A) p(\mathbf{y}^o \mid \mathbf{m}, E) d\mathbf{m}$ in (12) is the convolution of two densities, the first of which is the density for \mathbf{m} implied by the DSGE model, A . It can be accessed by means of the conventional simulations used in the calibration literature: if $\boldsymbol{\theta}_A^{(r)}$ is a random sample from $p(\boldsymbol{\theta}_A \mid A)$, then $\mathbf{m}_A^{(r)} = \mathbf{m}(\boldsymbol{\theta}_A^{(r)})$ is a random sample from $p(\mathbf{m} \mid A)$; similarly for $p(\mathbf{m} \mid B)$.

To access $p(\mathbf{y}^o \mid \mathbf{m}, E)$, define the auxilliary model E^* with the improper prior distribution $p(\mathbf{m} \mid E^*) \propto \text{constant}$, $p(\boldsymbol{\theta}_E \mid \mathbf{m}, E^*) = p(\boldsymbol{\theta}_E \mid \mathbf{m}, E)$ and $p(\mathbf{y} \mid \boldsymbol{\theta}_E, \mathbf{m}, E) = p(\mathbf{y} \mid \boldsymbol{\theta}_E, \mathbf{m}, E^*)$. Then

$$\begin{aligned} p(\mathbf{m} \mid \mathbf{y}^o, E^*) &\propto \int p(\mathbf{m} \mid E^*) p(\boldsymbol{\theta}_E \mid \mathbf{m}, E^*) p(\mathbf{y}^o \mid \boldsymbol{\theta}_E, \mathbf{m}, E^*) d\boldsymbol{\theta}_E \\ &\propto \int p(\boldsymbol{\theta}_E \mid \mathbf{m}, E) p(\mathbf{y}^o \mid \boldsymbol{\theta}_E, \mathbf{m}, E) d\boldsymbol{\theta}_E = p(\mathbf{y}^o \mid \mathbf{m}, E). \end{aligned} \quad (13)$$

A posterior simulator for \mathbf{y}^o and the auxilliary model E^* provides a sample $\mathbf{m}_{E^*}^{(s)}$ whose density kernel is $p(\mathbf{y}^o \mid \mathbf{m}, E)$.

Thus from (11)

$$\frac{p(A \mid \mathbf{y}^o, E)}{p(B \mid \mathbf{y}^o, E)} = \frac{p(A \mid E) \int p(\mathbf{m} \mid A) p(\mathbf{m} \mid \mathbf{y}^o, E^*) d\mathbf{m}}{p(B \mid E) \int p(\mathbf{m} \mid B) p(\mathbf{m} \mid \mathbf{y}^o, E^*) d\mathbf{m}}.$$

Models A and B can be compared on the basis of three simulations of the moment vector \mathbf{m} : $\mathbf{m}_A^{(r)}$ ($r = 1, \dots, N_A$) drawn from $p(\mathbf{m} \mid A)$, $\mathbf{m}_B^{(r)}$ ($r = 1, \dots, N_B$) drawn

from $p(\mathbf{m} | B)$, and $\mathbf{m}_{E^*}^{(r)}$ ($r = 1, \dots, N_{E^*}$) drawn from $p(\mathbf{m} | \mathbf{y}^o, E^*)$. Informal comparison can be based on a visual inspection of the clouds of points from these three models. A more formal comparison can be made by means of the kernel density approximation,

$$p(A | \mathbf{y}^o, E) \propto p(A | E) (N_A N_{E^*})^{-1} \sum_{r=1}^{N_A} \sum_{s=1}^{N_{E^*}} K(\mathbf{m}_A^{(r)}, \mathbf{m}_{E^*}^{(s)}). \quad (14)$$

5.2 Illustration in the equity premium model

In the equity premium example the vector \mathbf{m} consists of the population means for the risk free rate and the equity premium. Here we use perhaps the simplest econometric model E with implications for \mathbf{m} : a first-order Gaussian bivariate autoregression for the risk free rate at equity premium, with stationarity imposed:

$$\mathbf{y}_t - \mathbf{m} = \mathbf{F}(\mathbf{y}_{t-1} - \mathbf{m}) + \boldsymbol{\varepsilon}_t \quad (t = 1879, \dots, 1978)$$

where the 2×1 vector \mathbf{y}_t consists of the observed risk free rate and equity premium in the indicated year, and $\boldsymbol{\varepsilon}_t \stackrel{iid}{\sim} N(\mathbf{0}, \boldsymbol{\Sigma})$.

Draws of \mathbf{m} from the posterior distribution were obtained using a Metropolis within Gibbs posterior simulation algorithm. An improper prior for $(\mathbf{m}, \mathbf{F}, \boldsymbol{\Sigma})$, flat subject to the stationarity condition on \mathbf{F} , was employed. This prior satisfies the conditions for the model E^* discussed above. Some posterior moments for the parameters are indicated in Table 3. There is modest autocorrelation in the riskless rate (about .4), less in the equity premium (about .2), and very little cross correlation between the two time series. The innovation variance in the equity premium exceeds that of the risk free rate by a factor of more than 10. The implied standard deviation for the equity premium is over .16, and that for the riskless rate is over .05.

The posterior distribution of \mathbf{m} is by the crosses in each panel of Figures 4- 5. The range of values well within the support of the posterior distribution extends far beyond the observed sample means, indicated by the horizontal and vertical lines in each panel. A centered 90% posterior credible interval for the mean of the risk free rate extends from -0.9% to 2.6%. For the equity premium the range is much larger: from 2.2% to 9.7%. Even with 90 years of data, there is great uncertainty about the population mean of the equity premium. This uncertainty is due to the great variance in the equity premium from year to year. It is not due to drift: there has been no tendency for the equity premium to rise or fall secularly (Mehra and Prescott, 1985, Table 1), and the autocorrelation in the simple model used here is only 0.2.

The dots in each panel of Figures 4- 5 represent $p(\mathbf{m} | A)$ for the indicated model A . For each model, parameters were drawn from the same prior distributions used in the weak econometric interpretation. (These priors are summarized in Table 1.) Then, the corresponding population moments were computed. For the Mehra-Prescott and Rietz models there are closed form expressions for these moments. For the Labadie

and mixture models, a simulation of 1,000 periods was made corresponding to each set of parameter values drawn from the prior. Then a second antithetic simulation (i.e., shocks with signs reversed) was made. Then the mean of the risk free rate and equity premium averaged over the two 1,000-period simulations was used in lieu of the population mean.

Comparisons of the corresponding panels of Figures 2-3 and Figures 4- 5 reveal similar patterns. But the distributions in Figures 2-3 are more diffuse, relative to those in the latter figures which are more neatly demarcated and somewhat more compact. The difference reflects the sampling variation in 90-year averages, which is present in Figures 2-3 but not Figures 4- 5.

In the minimal econometric interpretation, a model receives support to the extent that the posterior density $p(\mathbf{m} | \mathbf{y}^o, E^*)$, represented by the crosses in Figures 4- 5, overlaps with the prior density $p(\mathbf{m} | A)$, represented by the dots, in a manner that is made explicit in equation (13). Because the posterior density $p(\mathbf{m} | \mathbf{y}^o, E)$ is so diffuse, there is substantial overlap between $p(\mathbf{m} | \mathbf{y}^o, E^*)$ and $p(\mathbf{m} | A)$ for each of the four alternative models A . This is true even of the Mehra-Prescott and Labadie models, which received no support under the weak econometric interpretation. The weak econometric interpretation takes literally the unrealistically small variance in the riskless return and equity premium implied by these two models and concludes that the cannot account for the observed averages. The minimal econometric interpretation utilizes the much greater sampling variation implied by the bivariate autoregression, and interprets the historical evidence about population moments as being much weaker. This finding underscores the point made forcefully by Eichenbaum (1991, p 611) that assuming the population moment is equal to the sample moment can be treacherous. “Data” and “facts” are not the same.

Formal approximation using (14) underscores these informal findings: see Table 4, which uses some of the same Gaussian kernels employed in Table 2. The ordering by Bayes factors is mixture model over Rietz model over Labadie model over Mehra-Prescott model, with the ratios being roughly 5 : 2 : 1 : 0.55. These conclusions are robust over the bandwidths indicated in Table 4.

6 Summary and conclusion

This study has examined three ways in which computational experiments can be used to see how well dynamic, stochastic general equilibrium (DSGE) models explain observed behavior. Since these models imply distributions for the paths of prices and quantities, a straightforward, likelihood based approach—termed the strong econometric interpretation in this study—is perhaps the most obvious. Many DSGE models fail under this interpretation because they predict exact relations that are not found in the data.

A widespread interpretation of DSGE models in the macroeconomics literature is that they are intended only to mimic the world along a carefully specified set of

dimensions. This interpretation is sometimes reduced to a list of sample moments, on the one hand, and a list of corresponding moments of the model's predictive distribution, on the other. Careful investigators recognize that some basis for comparison of these two sets of moments is needed. Kydland and Prescott (1996) clearly indicate that what is at stake is whether the sample moments are consistent with the predictive distribution of the model for those moments. This inherently Bayesian approach—termed the weak econometric interpretation in this study, and equivalent to a Bayesian prior predictive analysis—takes the period-to-period dynamics of the models literally. While it confines itself to just a few dimensions of the data, in accounting for sampling variation it makes the same assumptions as does the strong econometric interpretation. It is therefore subject to the same criticism—that those assumptions are inconsistent with what is observed.

To isolate the idea that DSGE models explain only certain dimensions, in a way that does not run afoul of the literal incredibility of these models, this study examined the implications of the claim that DSGE models predict only certain specified population moments of observable data. Since population moments are never observed, an auxiliary econometric model must also be assumed if the DSGE model is to have any refutable implications. The study showed that atheoretical econometric models with a Bayesian interpretation can perform this function, with the DSGE providing the prior distribution for the dimensions it claims to explain. Under this set of assumptions—termed the minimal econometric interpretation in this study—formal Bayesian model comparison is possible, and is free of the logical problems associated with the weak econometric interpretation.

These ideas were illustrated using the “equity premium puzzle” models of Mehra and Prescott (1985), Rietz (1988), Labadie (1989), and Tsionas (1994). The weak econometric interpretation reaffirmed both the inability of the Mehra-Prescott and Labadie models to account for the sample average risk free rate and equity premium in the U.S., and the ability of the Rietz model and the mixture models developed by Tsionas to do so. This reflects the fact that it is the weak econometric interpretation that is dominant in the DSGE literature of macroeconomics. This application provided a rich graphical interpretation of these models as well as Bayes factors for the comparison of models.

The minimal econometric interpretation of the same models greatly changed the nature of the findings, and underscores that point that the methodological issues raised in this study have substantive implications for macroeconomics. The most important finding was that we in fact have limited information about the population mean of the equity premium, because year-to-year fluctuations have been so great. The posterior distribution for the mean of the risk free rate and the equity premium supports values consistent with the original Mehra-Prescott model, the other models considered in this study, and quite likely with all other DSGE models designed to address this question. In short, there is no evidence of an equity premium puzzle.

These substantive findings illustrate the dangers of informal or undisciplined read-

ings of the evidence in macroeconomic time series. To push ahead with methods based on assumptions known to be incompatible with what is observed, or to ignore uncertainty that is manifest in the record, can lead large and well-intentioned groups of investigators astray. The benefits of an analytically rigorous economic theory will be realized only when harnessed to the same high standards for measurement.

References

- Box, G.E.P., 1980, "Sampling and Bayes Inference in Scientific Modeling and Robustness", *Journal of the Royal Statistical Society Series A* 143, 383-430.
- Christiano, L.J., and M. Eichenbaum, 1992, "Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations," *American Economic Review* 82, 430-450.
- DeJong, D.N., B.F. Ingram, and C.H. Whiteman, 1996, "A Bayesian Approach to Calibration," *Journal of Business and Economic Statistics* 14: 1-10.
- Eichenbaum, M., 1991, "Real Business Cycle Theory: Wisdom or Whimsy?" *Journal of Economic Dynamics and Control* 15: 607-626.
- Geweke, J., 2001, "A Note on Some Limitations of CRRA Utility," *Economics Letters* 71: 341-346.
- Geweke, J., 2005, *Contemporary Bayesian Econometrics and Statistics* (Wiley, Hoboken NJ).
- Good, I.J., 1956, "The Surprise Index for the Multivariate Normal Distribution", *Annals of Mathematical Statistics* 27: 1130-1135.
- Gregory, A.W. and G.W. Smith, 1991, "Calibration as Testing: Inference in Simulated Macroeconomic Models," *Journal of Business and Economic Statistics* 9: 297-304.
- King, R. G., C.I. Plosser, and S. T. Rebelo, 1988, "Production, Growth and Business Cycles: II. New Directions," *Journal of Monetary Economics* 21: 309-341.
- King, R. G., C.I. Plosser, and S. T. Rebelo, 1990, "Production, Growth and Business Cycles: Technical Appendix," University of Rochester mimeo.
- Kocherlakota, N., 1996, "The Equity Premium: It's Still a Puzzle," *Journal of Economic Literature* 34: 42-71.
- Kydland, F.E., and E.C. Prescott, 1996, "The Computational Experiment: An Econometric Tool," *Journal of Economic Perspectives* 10:1, 69-86.
- Labadie, P., 1989, "Stochastic Inflation and the Equity Premium," *Journal of Monetary Economics* 24: 277-298.
- Lancaster, T., 2004, *An Introduction to Modern Bayesian Econometrics* (Blackwell Publishing, Malden MA).
- Mehra, R., and E.C. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15: 145-161.
- Rietz, T.A., 1988, "The Equity Premium: A Solution," *Journal of Monetary Economics* 22: 17-33.
- Roberts, H.V., 1965, "Probabilistic Prediction," *Journal of the American Statistical Association* 60, 50-62.
- Smith, A.A., 1993, "Estimating Nonlinear Time-series Models using Simulated Vector Autoregressions," *Journal of Applied Econometrics* 8: S63-S84.
- Tsionas, E.G., 1994, "Asset Returns in General Equilibrium with Scale-Mixture-of-Normals Endowment Processes," University of Minnesota unpublished Ph.D. dissertation.

Watson, M.W., 1993, "Measures of Fit for Calibrated Models," *Journal of Political Economy* 101: 1-11-1041.

Table 1				
Deciles of prior distributions				
	λ_3	$\omega_{(2)}$	α	δ
0.1	.185	.332	.281	.840
0.2	.277	.491	.499	.907
0.3	.358	.677	.756	.939
0.4	.434	.917	1.077	.958
0.5	.509	1.26	1.500	.970
0.6	.583	1.77	2.089	.979
0.7	.659	2.68	2.978	.986
0.8	.738	4.61	4.509	.991
0.9	.826	11.12	8.016	.995

Table 2 ^a				
Weak Econometric Interpretation: Relative Log Marginal Likelihoods				
Gaussian smoothing kernel standard deviation				
Random draws	Mehra-Prescott	Rietz	Labadie	Mixture
	10^8	10^8	10^6	10^6
.0001	$-\infty$	1.512 (0.011)	$-\infty$	3.04 (0.76)
.0003	$-\infty$	1.887 (0.036)	$-\infty$	2.90 (0.21)
.0010	-114	1.902 (0.011)	-455	3.07 (0.06)
.0030	-19.03 (0.39)	1.9228 (0.0036)	-53.72 (0.68)	3.01 (0.02)
.0100	-3.0153 (.0006)	2.1558 (0.0009)	-2.894 (0.004)	3.086 (0.006)

^aModel predictive densities evaluated at sample means.

Table 3				
Bivariate first order autoregression				
Risk-free rate (r_t) and equity premium (e_t)				
$\begin{pmatrix} r_t - m_1 \\ e_t - m_2 \end{pmatrix}$	$=$	$\begin{vmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{vmatrix}$	$\begin{pmatrix} r_{t-1} - m_1 \\ e_{t-1} - m_2 \end{pmatrix}$	$+ \begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix};$
$\begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix}$	$\stackrel{iid}{\sim} N$	$\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{vmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{vmatrix} \right)$		

Posterior moments		
	Mean	Stan. dev.
m_1	.0088	.0106
m_2	.0591	.0227
f_{11}	.4362	.0907
f_{12}	-.0972	.0303
f_{21}	-.0065	.3143
f_{22}	.2003	.1077
σ_{11}	.0022	.0003
σ_{22}	.0268	.0041
σ_{12}	-.0009	.0008

Table 4 ^a				
Minimal Econometric Interpretation: Relative Log Marginal Likelihoods				
Gaussian smoothing kernel standard deviation	Mehra-Prescott	Rietz	Labadie	Mixture
.001	1.548 (0.007)	2.825 (0.016)	2.141 (0.006)	3.709 (0.008)
.003	1.699 (0.005)	2.837 (0.011)	2.053 (0.005)	3.696 (0.008)
.010	1.842 (0.005)	2.928 (0.008)	2.144 (0.005)	3.629 (0.007)

^aApproximations use (14) with $N_A = 10^6$, $N_E^* = 10^5$.

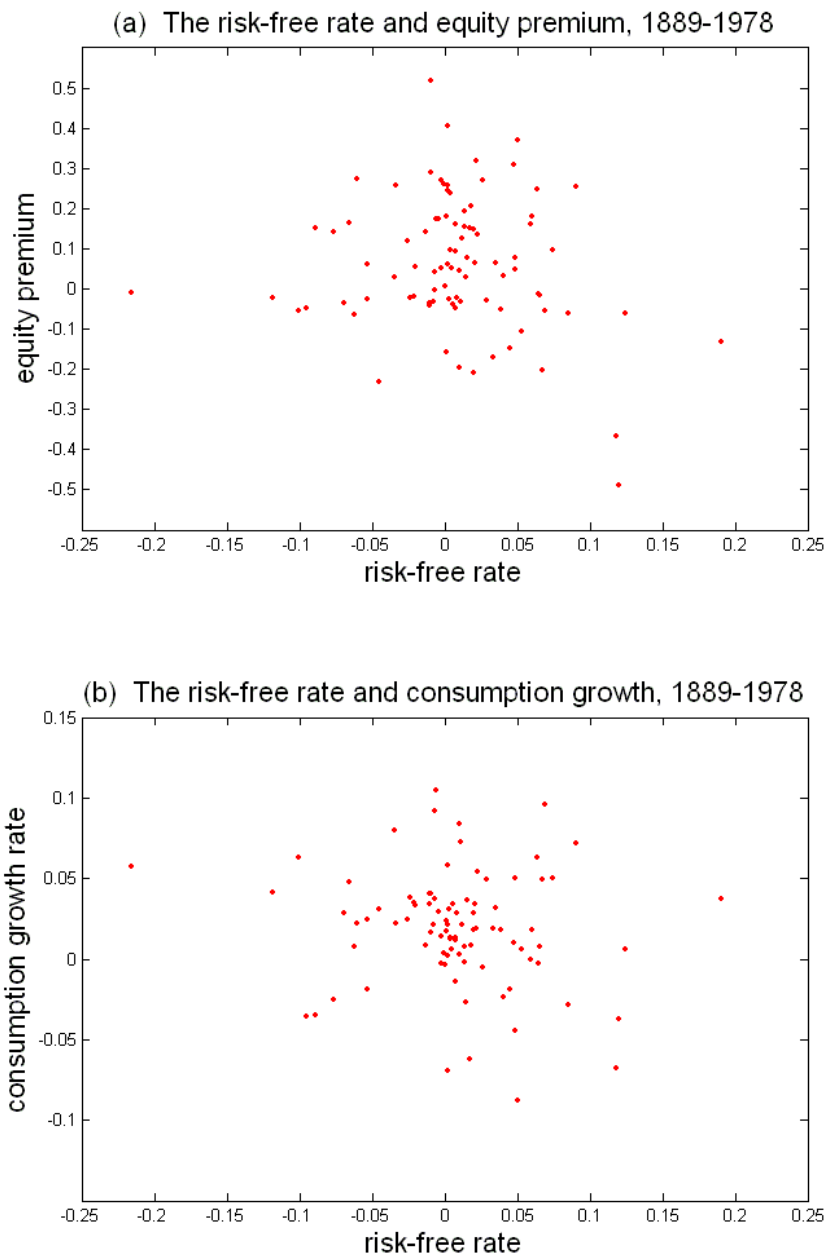


Figure 1: Annual combinations of the risk free rate and the equity premium and consumption growth rate, 1889-1978

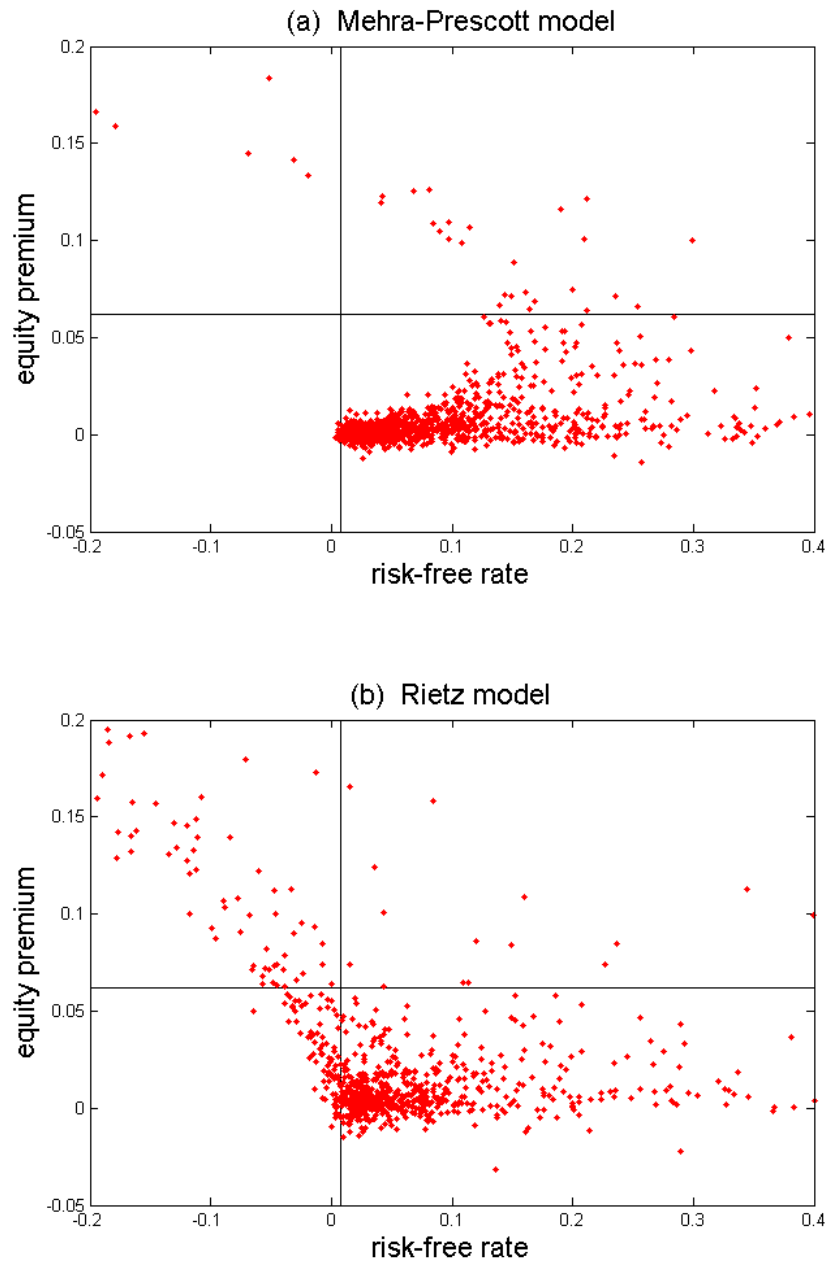


Figure 2: Predictive distributions for the risk-free rate and the equity premium under the weak econometric interpretation

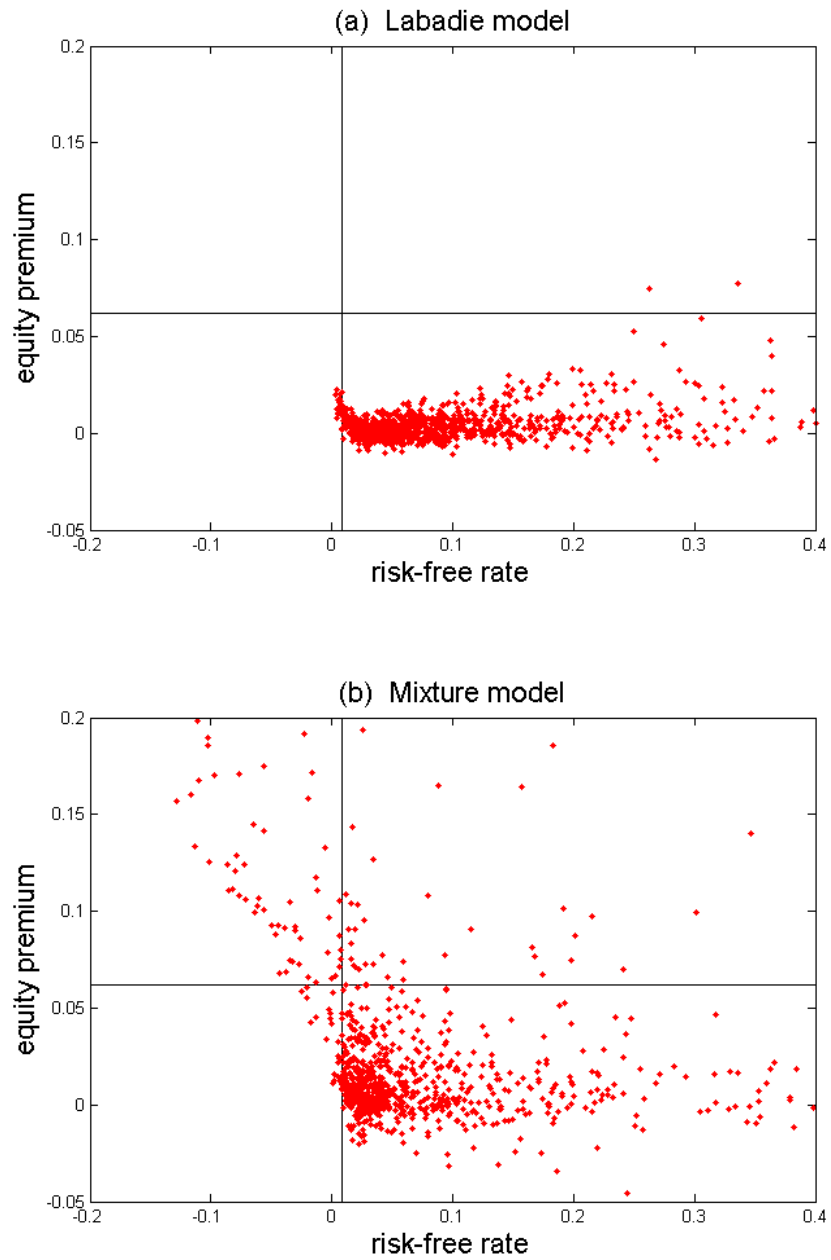


Figure 3: Predictive distributions for the risk-free rate and the equity premium under the weak econometric interpretation

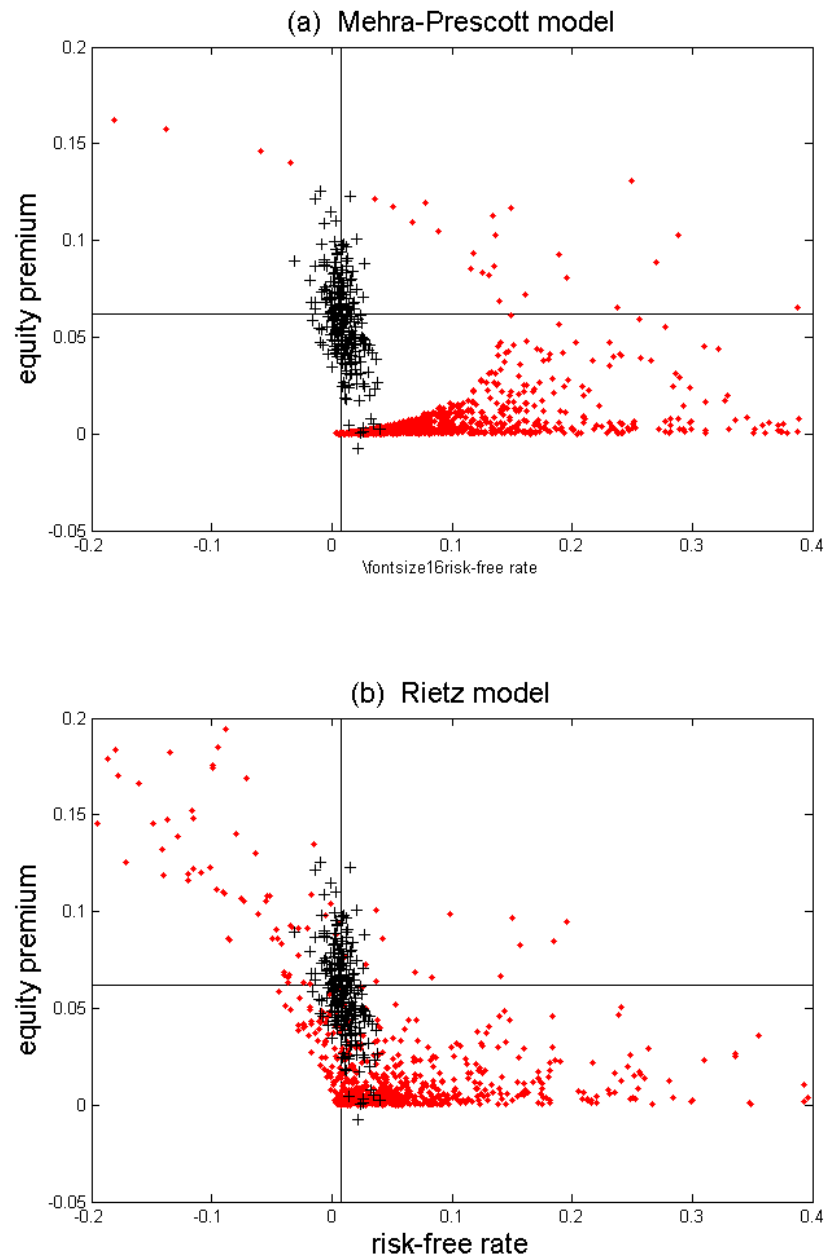


Figure 4: Distributions under the minimal econometric interpretation. Pluses represent the posterior distribution in the econometric model E^* .

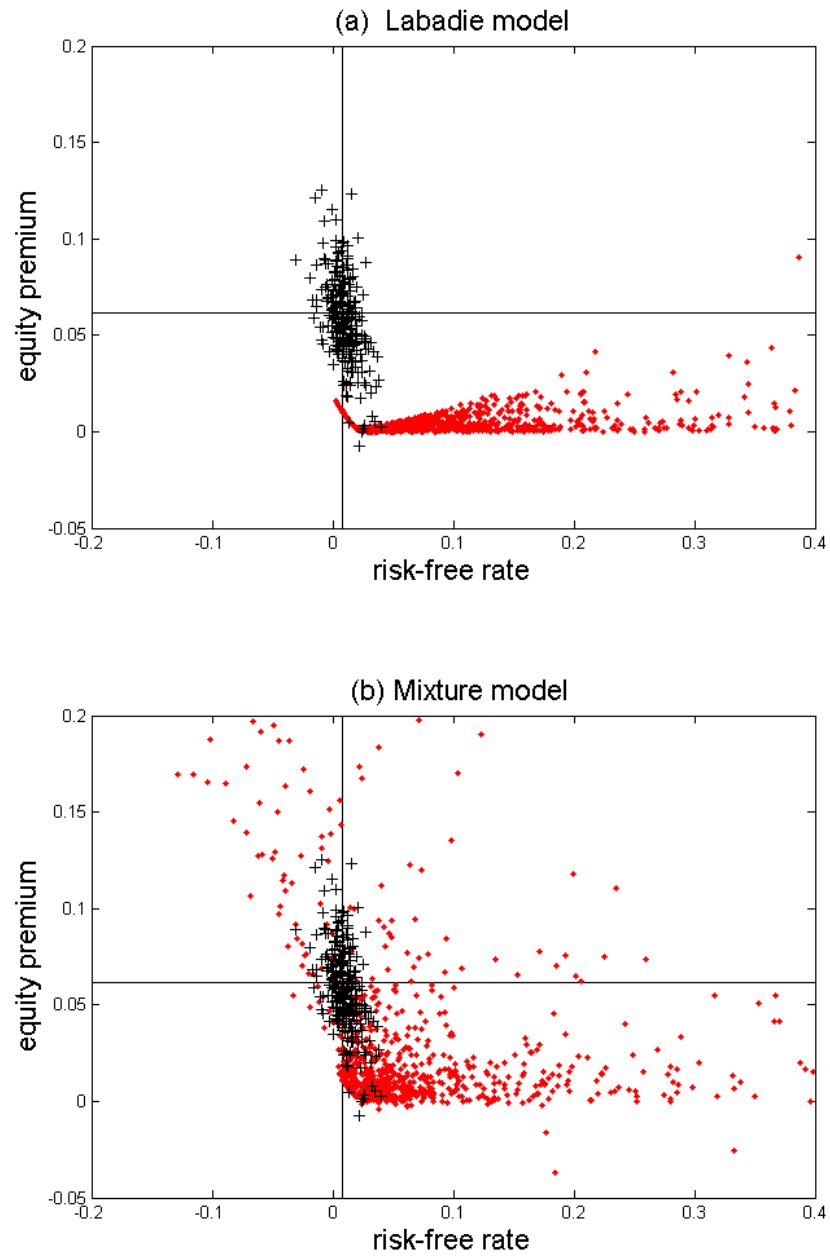


Figure 5: Distributions under the minimal econometric interpretation. Pluses represent the posterior distribution in the econometric model E^* .