Misspecified Recovery[*]

Jaroslav Borovička

New York University

jaroslav.borovicka@nyu.edu

Lars Peter Hansen

University of Chicago and NBER

lhansen@uchicago.edu

José A. Scheinkman

Columbia University, Princeton University and NBER

joses@princeton.edu

May 22, 2014

Abstract

Asset prices contain information about the probability distribution of future states and the stochastic discounting of these states. Without additional assumptions, probabilities and stochastic discounting cannot be separately identified. Ross (2013) introduced a set of assumptions that restrict the dynamics of the stochastic discount factor in a way that allows for the recovery of the underlying probabilities. We use decomposition results for stochastic discount factors from Hansen and Scheinkman (2009) to explain when this procedure leads to misspecified recovery. We also argue that the empirical evidence on asset prices indicates that the recovered measure would differ substantially from the actual probability distribution and that interpreting this measure as the true probability distribution may severely bias our inference about risk premia, investors' aversion to risk, and the welfare cost of economic fluctuations.

1 Introduction

It has been known, at least since the path-breaking work of Arrow, that asset prices reflect a combination of stochastic discounting and probability distributions. In a Markovian environment, prices of assets that pay a single dividend next period reflect both how consumers discount payoffs in the different states and the transition probabilities from the current state. Data on asset prices alone are not sufficient to identify non-parametrically the stochastic discount factor and transition probabilities without imposing additional restrictions. For instance, one may use additional time series evidence on the evolution of the Markov state which gives a separate way to identify the transition probabilities. The state-specific discount factors can then be inferred from asset prices.

Hansen and Scheinkman (2009) showed that if the transition probabilities satisfy an ergodicity property, one may use Perron–Frobenius theory to identify a distorted probability measure that reflects the long-term implications for risk pricing. They point out that this measure is typically distinct from the *physical* probability measure. In contrast, Ross (2013) uses Perron–Frobenius theory to claim a recovery result — a full identification of the transition probabilities from asset prices.

In this paper we connect the two results to make clear the special assumptions that are needed to guarantee that the distorted transition probabilities recovered using Perron–Frobenius theory equal the actual transition probabilities. In the general case, the ratio of the distorted to the true probability measure will be manifested as a non-trivial martingale component in the stochastic discount factor. Several of the structural models of asset pricing used in macroeconomics imply a non-trivial martingale component and existing empirical evidence suggests that this martingale component is quantitatively important.

We start in Section 2 by illustrating the problem of identifying the correct probability measure in a discrete-state space environment. In continuous-state spaces, additional technical challenges arise that may lead to multiple solutions of the Perron–Frobenius problem. The potential for multiplicity leads us in Section 3 to introduce a stochastic stability condition. As we show in Section 4, stochastic stability uniquely picks a particular candidate solution. We demonstrate in Section 5 that a Perron–Frobenius approach leads naturally to the construction of a martingale component to the stochastic discount factor process, a component that must be identically equal to one for Ross (2013)'s analysis to apply. We also show that this martingale component is not degenerate in several well known models of asset pricing models. In Section 6, we unify and extend results from a literature that bounds the magnitude of the martingale component that can be recovered from the asset pricing data. Section 7 concludes.

2 Illustrating the identification challenge

Hansen and Scheinkman (2014) gave a counting argument that illustrates the identification challenge that is present even in the simple setting of a finite state Markov chain. We review and expand on that discussion in what follows.

Let X be a discrete time, n-state Markov chain with transition matrix $\mathbf{P} \doteq [p_{ij}]$. We may identify state i with a coordinate vector u^i with a single one in the i-th entry. Suppose the analyst can infer the prices of one-period Arrow claims. We represent this input as a matrix $\mathbf{Q} = [q_{ij}]$ that maps a unit payoff tomorrow specified as a function of tomorrow's state into a price today. Since there are only a finite number of states, the payoff and price can both be represented as vectors. In particular, q_{ij} is the price in state $X_t = u^i$ of a security that pays one unit of consumption in state $X_{t+1} = u^j$.

Notice that $\mathbf{Q}$ has $n \times n$ entries. $\mathbf{P}$ has $n \times (n-1)$ entries because row sums have to add up to one. The realized one-period stochastic discount factors are entries of a matrix $\mathbf{S} \doteq [s_{ij}]$ where s_{ij} discounts one unit of consumption in state u^j tomorrow given that the current state is u^i. The discounting is state-dependent to adjust valuation to uncertainty in the next-period payout. Risk adjustments are encoded by this state dependence because each future state may be discounted differently. In general the stochastic discount factor introduces $n \times n$ free parameters s_{ij}, $i, j = 1, \ldots, n$. The Arrow prices are the products:[1]

$$q_{ij} = s_{ij}p_{ij}. \tag{1}$$

Additional data is needed to non-parametrically identify $\mathbf{P} = [p_{ij}]$ and $\mathbf{S} = [s_{ij}]$ from asset prices $\mathbf{Q} = [q_{ij}]$ without imposing additional restrictions. For instance, one may use time series evidence on the evolution of the Markov state to infer the transition probabilities. This gives a separate way to identify $\mathbf{P}$, and then the state-specific discount factors $\mathbf{S}$ can be inferred from the Arrow prices.

What follows are two ways to construct probabilities from the Arrow prices without appealing to time series data to identify the probabilities.

---

[1]The simple counting requires some qualification when $\mathbf{Q}$ has zeros. For instance, when $q_{ij} = 0$, then $p_{ij} = 0$ in order to prevent arbitrage opportunities. In this case the numerical value assigned to s_{ij} is inconsequential.

2.1 Risk-neutral probabilities

The financial engineering literature extensively uses the concept of risk-neutral probabilities, which are constructed as:

$$p_{ij}^* = \frac{q_{ij}}{\bar{q}_i}$$

where $\bar{q}_i = \sum_j^n q_{ij}$ is the price of a one-period discount bond in state i. In order to satisfy equation (1), we may define

$$s_{ij}^* = \bar{q}_i.$$

The risk-neutral probabilities p_{ij}^* can always be constructed and used in conjunction with discount factors $[s_{ij}^*]$. The discount factors are independent of state j, reflecting the absence of risk adjustments conditioned on the current state. In contrast, one-period discount bond prices can still be state-dependent. While $[p_{ij}^*]$ is a probability matrix, there is no claim that it represents actual transition probabilities, or that $[s_{ij}^*]$ represents actual valuation of consumption in states tomorrow by an investor.

2.2 Long-term pricing

We follow Backus et al. (1989) in studying long-term pricing of cash flows associated with fixed income securities using Perron–Frobenius theory. When there exists a $\lambda > 0$ such that the matrix $\sum_{t=0}^{\infty} \lambda^t \mathbf{Q}^t$ has all entries that are strictly positive, the largest eigenvalue $\exp(\eta)$ of $\mathbf{Q}$ is unique and positive and the associated eigenvector e has strictly positive entries. We denote the i^{th} entry of e as e_i. Typically, $\eta < 0$ to reflect time discounting of future payoffs.

Recall that we may evaluate t-period claims by applying the matrix $\mathbf{Q}$ for t times in succession. From the Perron–Frobenius theory for positive matrices:

$$\lim_{t \to \infty} \exp(-\eta t) \mathbf{Q}^t f = (f \cdot e^*) e$$

where e^* is the corresponding positive row eigenvector of $\mathbf{Q}$ and $e^* \cdot e$ is normalized to be one. The positive eigenvector e thus provides state-dependent valuations for payoffs maturing in the distant future. The price in current state u^i of an arbitrary payoff f maturing in $t \to \infty$ periods is, up to a proportionality constant, equal to e_i, and the average rate of discount of this long-horizon cash flow is $-\eta$.

The eigenvector e and the associated eigenvalue also provide a way to construct a probability transition matrix given $\mathbf{Q}$. Form

$$\widetilde{p}_{ij} = \exp(-\eta) q_{ij} \frac{e_j}{e_i} \tag{2}$$

Notice that

$$\sum_{j=1}^{n} \widetilde{p}_{ij} = \exp(-\eta)\frac{1}{e_i}\sum_{j=1}^{n} q_{ij}e_j = 1.$$

Thus $\widetilde{\mathbf{P}} = [\widetilde{p}_{ij}]$ is a transition matrix. Moreover,

$$q_{ij} = \exp(\eta)\frac{e_i}{e_j}\widetilde{p}_{ij}.$$

Hansen and Scheinkman (2009) and Ross (2013) both use this approach to construct a probability distribution, but they interpret it differently. Hansen and Scheinkman (2009) study multi-period pricing by compounding stochastic discount factors. They use the probability ratios for $\widetilde{p}_{ij}$ given by (2) and consider the following decomposition:

$$q_{ij} = \left[\exp(\eta)\frac{e_i}{e_j}\frac{\widetilde{p}_{ij}}{p_{ij}}\right]p_{ij} = \exp(\eta)\left(\frac{e_i}{e_j}\right)h_{ij}p_{ij}.$$

Hence,

$$s_{ij} = \exp(\eta)\left(\frac{e_i}{e_j}\right)h_{ij} \tag{3}$$

where $\mathbf{H} = [h_{ij}]$ and

$$h_{ij} = \widetilde{p}_{ij}/p_{ij}$$

provided that $p_{ij} > 0$. When $p_{ij} = 0$ the construction of h_{ij} is inconsequential.

Hansen and Scheinkman (2009) interpret the three components of the one-period stochastic discount factor displayed on the right-hand side of (3) and show how this representation can be used to study long-term valuation. The third term, which is a ratio of probabilities, is used as a change of probability measure in their analysis. In contrast, Ross (2013) uses (2) to identify the transition probability used by investors by presuming that $h_{ij} = 1$ for all (i, j), which implies that the transition probabilities $\mathbf{P}$ and $\widetilde{\mathbf{P}}$ coincide. In the next sections we address these issues under much more generality by allowing for continuous-state Markov processes. As we will see, some additional complications emerge.

To anticipate the counterpart to (3) for a more general probability model, consider the stochastic discount factor process $S = \{S_t : t = 0, 1, 2, ...\}$ that is obtained by compounding the one-period stochastic discount factor. In general, S_t depends on the history of the state from 0 to t since the increment between t and $t + 1$ is given by:

$$\frac{S_{t+1}}{S_t} = X_t'\mathbf{S}X_{t+1}.$$

Similarly, we define

$$\frac{H_{t+1}}{H_t} = X_t' \mathbf{H} X_{t+1}.$$

Then

$$\frac{S_{t+1}}{S_t} = \exp(\eta) \left(\frac{H_{t+1}}{H_t} \right) \left(\frac{e \cdot X_t}{e \cdot X_{t+1}} \right)$$

Compounding the three components in decomposition (3), we obtain:

$$S_t = \exp(\eta t) \left(\frac{H_t}{H_0} \right) \left(\frac{e \cdot X_0}{e \cdot X_t} \right) \tag{4}$$

where we initialize $S_0 = 1$ and

$$\frac{H_t}{H_0} = \prod_{\tau=1}^{t} X_{\tau-1}' \mathbf{H} X_\tau.$$

Because h_{ij} is obtained as a ratio of probabilities, H is a positive martingale under P for any positive specification of H_0 as a function of X_0. Thus from (4), the eigenvalue η contributes an exponential function of t and the eigenvector contributes a function of the Markov state to the stochastic discount factor process. In addition there is martingale component, whose logarithm has stationary increments. The assumption of Ross (2013) thus amounts to requiring that $H_t/H_0 \equiv 1$ or equivalently that S has no martingale component.

3 Setup

We start with a set of indices T (either the non-negative integers or the non-negative reals) and $X = \{X_t\}_{t \in T}$ a stationary Markov process on a probability space $\{\Omega, \mathcal{F}, P\}$. Write $\mathfrak{F} \doteq \{\mathcal{F}_t\}_{t \in T}$ for the (completed) filtration generated by the histories X_u, $u \leq t$.

3.1 Constructing other processes

We use X to build other processes. Consider a process M specified in discrete time by

$$\log M_{t+1} - \log M_t = g(X_{t+1}, X_t) \tag{5}$$

where $\log M_0 = 0$. When X is a continuous time diffusion in a space with a Brownian motion W, suppose that

$$d \log M_t = \beta(X_t)dt + \alpha(X_t) \cdot dW_t, \tag{6}$$

again with $\log M_0 = 0$. In particular, the $\log M$ process has stationary increments in both cases.

These two specifications are sufficiently general for us to discuss identification and recovery. The resulting constructions of M give rise to special cases of what are called *positive multiplicative functionals*.[2] In this paper a positive process is one that is strictly greater than zero almost surely for all t in T.

Definition 3.1. *A* **multiplicative functional** *is a process M that is adapted to $\mathfrak{F}$, $M_0 = 1$, and*

$$M_t(X) = M_\tau(X)M_{t-\tau}(\theta_\tau(X)), \tag{7}$$

where θ_τ is the shift operator that moves the time subscript of X by τ, that is,

$$(\theta_\tau(X))_s = X_{\tau+s}.$$

We will analyze alternative measures equivalent to P that we build using strictly positive martingales. Given an $\mathfrak{F}$-martingale H that is strictly positive and has unit expectation for all t, we may define a probability P^H such that if $A \in \mathcal{F}_\tau$ for some $\tau \geq 0$,

$$P^H(A) = E(1_A H_\tau). \tag{8}$$

The Law of Iterated Expectations guarantees that these definitions are consistent, that is, if $A \in \mathcal{F}_\tau$ and $t > \tau$ then

$$P^H(A) = E(1_A H_t) = E(1_A H_\tau).$$

3.2 Stochastic stability

For general state spaces, our analysis requires that stability restrictions are imposed when we change probability measures. We first specify our most general condition, and then we construct two more easily interpretable sufficient conditions.

Condition 3.2. *The process X is* **stochastically stable in averages** *under a probability distribution P^H if for any bounded (Borel measurable) function f,*

$$\lim_{N \to \infty} E^H \left[\frac{1}{N} \sum_{t=1}^{N} f\left(X_t\right) | X_0 = x \right] = E^H f(X_0).$$

---

[2]There are two extensions of this setup that are also of substantive interest. In Hansen and Scheinkman (2014), we have extended this analysis to allow for a richer specification of uncertainty by introducing explicit additional shocks that do not directly influence the state dynamics, but at the same time have an impact on valuation. Also, while the Brownian information specification (6) abstracts from jumps, these can be included without changing the implications of the analysis, see Hansen and Scheinkman (2009).

when T is the set of nonnegative integers (discrete time), or, when T is the set of nonnegative real numbers (continuous time),

$$\lim_{N \to \infty} E^H \left[\frac{1}{N} \int_0^N f(X_t) \, dt \, | X_0 = x \right] = E^H f(X_0)$$

pointwise in x for almost all x.

Condition 3.2 allows for periodic components to the Markov process. When such components are absent, a simpler but stronger version of Condition 3.2 is given by:

Condition 3.3. *The process X is **stochastically stable** under a probability distribution P^H if for any bounded (Borel measurable) function f,*

$$\lim_{t \to \infty} E^H \left[f(X_t) \, | X_0 = x \right] = E^H f(X_0).$$

pointwise in x for almost all x.

This condition was imposed by Hansen and Scheinkman (2009) and Hansen (2012) to select among Perron–Frobenius eigenfunctions.

Another sufficient condition uses the Law of Large Numbers (LLN) for Markov processes conditioned on an initial state $X_0 = x$. This version also allows for periodic components.

Condition 3.4. *The process X obeys the **LLN** under a probability distribution P^H if for any bounded (Borel measurable) function f,*

$$\lim_{N \to \infty} E^H \left[\left| \frac{1}{N} \sum_{t=1}^N f(X_t) - E^H f(X_0) \right| \, | X_0 = x \right] = 0$$

when T is the set of nonnegative integers (discrete time), or, when T is the set of nonnegative real numbers (continuous time),

$$\lim_{N \to \infty} E^H \left[\left| \frac{1}{N} \int_0^N f(X_t) \, dt - E^H \left[f(X_0) \right] \right| \, | X_0 = x \right] = 0$$

pointwise in x for almost all x.

To see that Condition 3.2 is implied by Condition 3.4, notice that

$$E^H \left[\left| \frac{1}{N} \sum_{t=1}^N f(X_t) - E^H f(X_0) \right| \, | X_0 = x \right] \geq \left| E^H \left[\frac{1}{N} \sum_{t=1}^N f(X_t) \, | X_0 = x \right] - E^H f(X_0) \right|$$

If the left-hand side converges to zero, then so does the right-hand side. An entirely similar argument applies for the continuous-time case.

Since we condition on X_0 in our statement of Conditions 3.2 and 3.3, the conditional change of measure is represented by H/H_0:

$$E^H \left[\frac{1}{N} \sum_{t=1}^{N} f(X_t) \, | X_0 = x \right] = E \left[\left(\frac{H_N}{H_0} \right) \frac{1}{N} \sum_{t=1}^{N} f(X_t) \, | X_0 = x \right].$$

The random variable H_0 that initializes H defines the limit points for the LLN and makes X stationary under the probability P^H implied by H. The random variable $H_0 = h(X_0)$ must satisfy the equation:

$$E \left[f(X_t) \left(\frac{H_t}{H_0} \right) h(X_0) \right] = E \left[f(X_0) h(X_0) \right].$$

for any bounded (Borel measurable) f.

We have purposely stated these conditions for a probability measure P^H associated with H. We suppose that the stability Condition 3.2 is satisfied for the probability P. As we change probability measures, stochastic stability will not necessarily be satisfied, but checking for this stability under the probability P^H will be used in our analysis.

4   Stochastic discount factor, probability measure and recovery

A *stochastic discount factor process* S is a positive multiplicative functional with finite first moments (conditioned on $X_0 = x$) such that the date zero price of any claim $\psi(X_t)$ payable at t is

$$[\mathbb{Q}_t \psi](x) \doteq E[S_t \psi(X_t) | X_0 = x]. \tag{9}$$

The multiplicative property of S allows us to price consistently at intermediate dates. In discrete time we may build the t-period operator $\mathbb{Q}_t$ by applying the one-period operator $\mathbb{Q}_1$ t times in succession. In continuous time, $\{\mathbb{Q}_t : t \in T\}$ forms what is called a semigroup of operators. The counterpart to a one-period operator is a *generator* of this semigroup that governs instantaneous valuation and which acts as a time derivative of $\mathbb{Q}_t$ at $t = 0$. Thus in discrete time it suffices to study the one-period operator and in continuous time the generator of the family of operators $\{\mathbb{Q}_t : t \in T\}$.

We are typically interested in pricing a richer collection of cash flows including ones that

display stochastic growth and ones that are history-dependent. Given a probability measure P, the pricing operator in (9) applied to all (bounded, measurable) functions $\psi(x)$ for every $t \geq 0$ is sufficient to determine the stochastic discount factor process S, which can be used to assign prices to these other cash flows.

Notice that equation (9) involves simultaneously the probability P on $\mathfrak{F}$-measurable sets and the process S. In other words, a stochastic discount factor process is only well-defined for a given probability measure.

Definition 4.1. *The pair (S, P) explains asset prices if equation (9) gives the date zero price of any claim $\psi(X_t)$ payable at any time $t \in T$.*

Consider now a strictly positive martingale H and the associated probability measure P^H defined through (8). We define:

$$S^H = S\frac{H_0}{H}.$$

The following proposition is immediate:

Proposition 4.2. *Suppose that $E(H_0) = 1$, H is a positive martingale, and $\frac{H}{H_0}$ is a multiplicative functional. If the pair (S, P) explains asset prices then the pair (S^H, P^H) also explains asset prices.*

This proposition captures the notion that stochastic discount factors are only well-defined for a given probability distribution. When we change the probability distribution, we must also change the stochastic discount factor in order to represent the same asset prices. An analogous observation carries over to so-called risk prices that are assigned by the stochastic discount factor to alternative risk exposures. In other words we have multiple ways to represent $\mathbb{Q}$:

$$[\mathbb{Q}_t \psi](x) = E[S_t \psi(X_t)|X_0 = x] = E^H\left[S_t^H \psi(X_t)|X_0 = x\right] \tag{10}$$

where E^H is the expectation operator associated with P^H.

So far we have imposed one probability measure P and shown what happens when we consider other probability measures. Suppose, however, that this probability distribution is not known to an external analyst. Given that H can be any positive multiplicative martingale, we are left with a *fundamental identification problem*. From the Arrow prices alone we cannot distinguish (S, P) from (S^H, P^H). In particular, we cannot recover P from the Arrow prices alone. To achieve identification, either we have to restrict the stochastic discount factor process S or we have to restrict the probability distribution used to represent the valuation operators $\mathbb{Q}$.

There are multiple ways we might address this lack of identification. First, we might impose rational expectations, observe time series data, and let the Law of Large Numbers for stationary distributions determine the probabilities. Then observations for a complete set of Arrow securities allows us to identify S. See Hansen and Richard (1987) for an initial discussion of the stochastic discount factors and the Law of Large Numbers, and see Hansen and Singleton (1982) for an econometric approach that imposes a parametric structure on the stochastic discount factor and avoids assuming that the analyst has access to data on the complete set of Arrow securities.

Second, as in Ross (2013), we may impose a special structure on S by assuming:

Assumption 4.3. *Let*

$$S_t = \exp(-\delta t)\frac{m(X_t)}{m(X_0)}$$

for some positive function m and some real number δ.

The pair (S^H, P^H) that explains asset prices and satisfies Assumption 4.3 is unique provided that stochastic stability is preserved under the measure P^H. Formally,

Proposition 4.4. *Suppose (S, P) explain asset prices and S satisfies Assumption 4.3. Let $\frac{H}{H_0}$ be a positive multiplicative martingale such that (S^H, P^H) also explain asset prices and X is stochastically stable in averages (satisfies Condition 3.2) under P^H. If S^H also satisfies Assumption 4.3, then $H \equiv 1$.*

The proof of this theorem is similar to the proof of a related uniqueness result in Hansen and Scheinkman (2009) and is detailed in Appendix A.

Remarkably, Proposition 4.4 gives an identification result, which is a counterpart of the *recovery result* in Ross (2013). Under the conditions of this proposition, we may infer both the stochastic discount factor and a probability distribution associated with that stochastic discount factor from the pricing operator $\mathbb{Q}_t$. In the next section we relax Assumption 4.3, and we ask "what does this approach actually recover?"

To illustrate the role the stability in selecting a single function that satisfies Assumption 4.3, and thus a single probability distribution, consider the following example.

Example 4.5. *Suppose that X is a Feller square root process:*

$$dX_t = -\kappa(X_t - \bar{\mu})dt + \bar{\sigma}\sqrt{X_t}dW_t$$

where $\kappa > 0$, $\bar{\mu} > 0$ and $\kappa\bar{\mu} \geq \frac{1}{2}(\bar{\sigma})^2$. With these restrictions, the process X is stochastically stable and strictly positive.

Let $m(x) = \exp(\zeta x)$ so that the stochastic discount factor satisfies

$$S_t = \exp\left[-\delta t + \zeta(X_t - X_0)\right], \tag{11}$$

or, in differential form,

$$d \log S_t = \left[-\delta - \zeta \kappa (X_t - \bar{\mu})\right] dt + \zeta \bar{\sigma} \sqrt{X_t} dW_t.$$

We choose a multiplicative martingale of the form:

$$\frac{H_t}{H_0} = \exp\left[-\int_0^t \xi(X_s) dW_s - \frac{1}{2}\int_0^t \xi(X_s)^2 ds\right].$$

and thus

$$S_t^H = S_t \frac{H_0}{H_t}$$

must satisfy:

$$d \log S_t^H = \left(-\delta - \zeta \kappa (X_t - \bar{\mu}) + \frac{1}{2}[\xi(X_t)]^2\right) dt + \left[\zeta \bar{\sigma}\sqrt{X_t} + \xi(X_t)\right] dW_t. \tag{12}$$

We will show that there is a choice of ξ such that

$$S_t \neq S_t^H = \exp\left(-\widetilde{\delta} t\right) \frac{\exp\left(\widetilde{\zeta} X_t\right)}{\exp\left(\widetilde{\zeta} X_0\right)},$$

or

$$d \log S_t^H = \left[-\widetilde{\delta} - \widetilde{\zeta} \kappa (X_t - \bar{\mu})\right] dt + \widetilde{\zeta} \bar{\sigma} \sqrt{X_t} dW_t.$$

Equating coefficients:

$$\xi(x) = \left(\widetilde{\zeta} - \zeta\right) \bar{\sigma} \sqrt{x},$$

and

$$-\zeta \kappa x + \frac{1}{2}\left(\widetilde{\zeta} - \zeta\right)^2 \bar{\sigma}^2 x = -\widetilde{\zeta} \kappa x.$$

Consequently, either $\widetilde{\zeta} = \zeta$ or,

$$\widetilde{\zeta} = \zeta - \frac{2\kappa}{\bar{\sigma}^2}.$$

In this case,

$$\xi(x) = -\frac{2\kappa}{\bar{\sigma}} \sqrt{x}.$$

For this second solution, under the probability P^H the Brownian motion W has a local drift

$-\xi(X_t)$ *and hence the local mean for* X *is*

$$-\kappa \left(x - \bar{\mu} \right) + 2\kappa x = \kappa x + \kappa \bar{\mu}.$$

Thus the dynamics for X *are*

$$dX_t = \kappa(X_t + \bar{\mu})dt + \bar{\sigma}\sqrt{X_t}d\widetilde{W}_t$$

where $\widetilde{W}_t = W_t + \int_0^t \xi(X_s)ds$ *is a Brownian motion under* P^H. *Since* $\kappa > 0$, X *is not stochastically stable when* $\widetilde{\zeta} = \zeta - \frac{2\kappa}{\bar{\sigma}^2}$.

5   What is recovered?

This section reviews results in Hansen and Scheinkman (2009), Hansen (2012) and Hansen and Scheinkman (2014). These references and Ross (2013) use Perron–Frobenius theory in their analyses. At this juncture, we adopt a more general starting point and do not impose Assumption 4.3.

Suppose there exists a function $e(x)$ that solves the following *Perron–Frobenius* problem:

Problem 5.1 (Perron–Frobenius)**.** *Find a scalar* η *and a function* $e > 0$ *such that for every* $t \in T$,

$$[\mathbb{Q}_t e](x) = \exp(\eta t)e(x).$$

When the state space is finite as in Section 2, functions of x can be identified with vectors in $\mathbb{R}^n$, and the operator $\mathbb{Q}_1$ can be identified with a matrix $\mathbf{Q}$. If the matrix $\sum_{t=0}^{\infty} \lambda^t \mathbf{Q}_t$ is strictly positive for some positive λ, then the Perron–Frobenius Theorem states that there exists a unique (up to scale) solution to Problem 5.1.

Existence and uniqueness are more complicated in the case of general state spaces. Hansen and Scheinkman (2009) present sufficient conditions for the existence of a solution, but even in examples commonly used in applied work, multiple (scaled) positive solutions are a possibility. See Hansen and Scheinkman (2009), Hansen (2012) and our subsequent discussion for such examples. When we have a solution of the Perron–Frobenius problem, we follow Hansen and Scheinkman (2009), and define a process $\widetilde{H}$ that satisfies:

$$\frac{\widetilde{H}_t}{\widetilde{H}_0} = \exp(-\eta t)S_t \frac{e(X_t)}{e(X_0)}. \tag{13}$$

The process $\widetilde{H}$ is an $\mathfrak{F}$-martingale under the probability measure P. In fact,

$$E\left[\frac{\widetilde{H}_t}{\widetilde{H}_0} \mid \mathcal{F}_\tau\right] = \frac{\exp(-\eta t)}{e(X_0)} S_\tau e(X_\tau) E\left[\frac{S_t e(X_t)}{S_\tau e(X_\tau)} \mid \mathcal{F}_\tau\right] = \frac{\exp(-\eta \tau)}{e(X_0)} S_\tau e(X_\tau) = \frac{\widetilde{H}_\tau}{\widetilde{H}_0},$$

where in the second equality we used the property from equation (7) of the definition of a multiplicative process.

Note that $\widetilde{H}$ inherits much of the mathematical structure of the original process S. For instance, if S has the form given by equation (5) above, then in discrete time:

$$\log \widetilde{H}_{t+1} - \log \widetilde{H}_t = g(X_{t+1}, X_t) + \log e(X_{t+1}) - \log e(X_t) - \eta$$
$$= \widetilde{g}(X_{t+1}, X_t)$$

In logarithms, this constructed process has stationary increments.

As we noted, such a construction may not be unique. (For instance, see Example 4.5.) If, however, we restrict $\widetilde{H}$ so that X under the implied probability distribution is stochastically stable in averages (Condition 3.2), then there is at most one such $\widetilde{H}$. The proof of this uniqueness result is essentially the same as that of Proposition 4.4, and is a result reported in Hansen and Scheinkman (2009). While this recovers a single probability, there is no claim that this constructed probability is the one that generates the data.

Note that if $\widetilde{H} \equiv 1$ then S necessarily satisfies Assumption 4.3. On the other hand, if S satisfies Assumption 4.3, then $e(x) = \frac{1}{m(x)}$ solves the Perron–Frobenius Problem 5.1 and the associated $\widetilde{H} \equiv 1$. Thus we obtain:

Proposition 5.2. *S satisfies Assumption 4.3 if and only if there exists a solution $e > 0$ to the Perron–Frobenius problem and the associated η such that:*

$$\frac{\widetilde{H}_t}{\widetilde{H}_0} = \exp(-\eta t) S_t \frac{e(X_t)}{e(X_0)} \equiv 1.$$

This proposition reveals the recovery result as a special case of the decomposition (13). Recall that the primitive of our analysis is the pair (S, P) where the probability measure P satisfies the stochastic stability condition. If the Perron–Frobenius Problem 5.1 yields a solution $\widetilde{H} \equiv 1$, the uniqueness result implies that P is the single stochastically stable measure that will be recovered. In the general case, the stochastic discount factor S contains a martingale component $\widetilde{H} \neq 1$ for which the associated measure $\widetilde{P}$ satisfies the stability condition. In this case the recovery procedure will *not* identify the underlying probability

measure. Instead, it will recover the measure $\widetilde{P}$ implied by $\widetilde{H}$.[3]

5.1   Long-term pricing

The alternative probability associated with $\widetilde{H}$ turns out to be useful in representing long-term values. Consider

$$[\mathbb{Q}_t\psi](x) = \exp(\eta t)e(x)\widetilde{E}\left[\frac{\psi(X_t)}{e(X_t)} \mid X_0 = x\right]$$

for some positive ψ. We use the notation $\widetilde{E}$ to denote expectations computed with the probability $\widetilde{P}$ implied by $\widetilde{H}$.[4] For pricing pure discount bonds we can set $\psi(x) \equiv 1$. Taking the negative logarithms and dividing by t gives the yield on the payoff $\psi(X_t)$:

$$-\frac{1}{t}\log[\mathbb{Q}_t\psi](x) = -\eta - \frac{1}{t}\log e(x) - \frac{1}{t}\log\widetilde{E}\left[\frac{\psi(X_t)}{e(X_t)} \mid X_0 = x\right].$$

Taking the limit $t \to \infty$ and using the stochastic stability Condition 3.3 under the probability $\widetilde{P}$ shows that $-\eta$ is the long term yield on an arbitrary cash flow $\psi(x)$ maturing in the distant future, provided

$$\widetilde{E}\left[\frac{\psi(X_t)}{e(X_t)}\right] < \infty.$$

Under the same assumption, the long maturity limit of a holding period return $R_{t,t+1}^\tau$ between period t and period $t+1$ on bond with maturity τ is

$$R_{t,t+1}^\infty = \lim_{\tau\to\infty} R_{t,t+1}^\tau = \lim_{\tau\to\infty} \frac{[\mathbb{Q}_{\tau-1}\psi](X_{t+1})}{[\mathbb{Q}_\tau\psi](X_t)} = \exp(-\eta)\frac{e(X_{t+1})}{e(X_t)}.$$

Moreover, the "Euler-equation error" associated with the limit return is

$$\left(\frac{S_{t+1}}{S_t}\right)R_{t,t+1}^\infty = \frac{\widetilde{H}_{t+1}}{\widetilde{H}_t}, \tag{14}$$

which reveals the martingale increment in the stochastic discount factor. We will use this formula in Section 6 when we discuss empirical methods and evidence.

---

[3]Example 3.1 in Hansen and Scheinkman (2014) points out this connection in a finite-state Markov chain setting.

[4]Note that we use $\widetilde{P}$ and $\widetilde{E}$, instead of the more cumbersome $P^{\widetilde{H}}$ and $E^{\widetilde{H}}$.

5.2 Examples

In the previous discussion, we described two issues arising in the recovery procedure. First, the positive candidate solution for $e(x)$ may not be unique. Our stability criterion allows us to pick the single solution that preserves stability. Second, and more importantly, even this unique choice may not uncover the true probability distribution if there is a martingale component in the stochastic discount factor. We illustrate these challenges using stylized versions of widely used asset pricing models. The next example shows that in a simplified version of a stochastic volatility model one always recovers an incorrect probability distribution that, for certain parameter values, actually coincides with the risk-neutral probability measure.

Example 5.3. *Consider a stochastic discount factor model with state-dependent risk prices. Suppose that*

$$d \log S_t = \bar{\beta}dt - \frac{1}{2}X_t (\bar{\alpha})^2 dt + \sqrt{X_t}\bar{\alpha}dW_t$$

where X has the square root dynamics given in Example 4.5 and $\bar{\beta} < 0$. Guess a solution:

$$e(x) = \exp(\upsilon x)$$

Since $\{\exp(-\eta t)S_t e(X_t) : t \geq 0\}$ is a martingale, its local mean should be zero:

$$\bar{\beta} - \frac{1}{2}(\bar{\alpha})^2 - \upsilon\kappa x + \upsilon\kappa\bar{\mu} + \frac{1}{2}x (\upsilon\bar{\sigma} + \bar{\alpha})^2 = 0.$$

In particular, the coefficient on x should satisfy

$$-\upsilon\left[-\kappa + \frac{1}{2}\upsilon (\bar{\sigma})^2 + \bar{\sigma}\bar{\alpha}\right] = 0.$$

There are two solutions: $\upsilon = 0$ and

$$\upsilon = \frac{2\kappa - 2\bar{\alpha}\bar{\sigma}}{(\bar{\sigma})^2} \tag{15}$$

For this example, the risk neutral dynamics for X corresponds to the solution $\upsilon = 0$ and the instantaneous risk-free rate is constant and equal to $-\bar{\beta}$. The resulting X process remains a square root process, but with κ replaced by

$$\kappa_n = \kappa - \bar{\alpha}\bar{\sigma}.$$

Although κ is positive, κ_n could be positive or negative. If $\kappa_n > 0$, then the Perron–Frobenius

problem that we feature extracts the risk-neutral dynamics, but this is distinct form the actual probability evolution for X. Suppose instead that $\kappa_n < 0$. This occurs when $\kappa < \bar{\sigma}\bar{\alpha}$. In this case we choose υ according to (15), implying that κ is replaced by

$$\kappa_{pf} = -\kappa + \bar{\sigma}\bar{\alpha} = -\kappa_n > 0.$$

The resulting dynamics are distinct from both the risk-neutral dynamics and the original dynamics for the process X.

This example was designed to keep the algebra simple, but there are straightforward extensions that are described in Hansen (2012). In this example there is a single shock that shifts both the stochastic discount factor and the state variable X that governs volatility. The so-called "local risk price" for this shock is given by $-\sqrt{X_t}\sigma_s$. There is a straightforward extension that extends the dimension of the Brownian increment and reproduces analogous findings. In addition, a predictable component can be included in the local mean for the $\log S$. Hansen (2012) includes these generalizations using a model with affine dynamics of the type featured by Duffie and Kan (1994).

The continuous-time model of Campbell and Cochrane (1999) has some implications that parallel those captured by Example 5.3.

Example 5.4. *Suppose that*

$$d\log S_t = \bar{\beta}dt + \bar{\alpha}dW_t + d\log Z_t$$

where $Z_t = k(X_t) > 0$ and X is a stationary Markov process that captures the contribution of external or social habits to investor preferences. For the continuous-time Campbell and Cochrane (1999) model, there are special functional forms for k and the evolution of X. By design, their model has a constant instantaneous interest rate ρ. A candidate for a Perron–Frobenius eigenvalue is $\eta = -\rho$ and an associated positive eigenfunction is constant. This eigenfunction induces unstable state dynamics in their specification. There is a second eigenfunction $\log e(x) = -\log k(x)$, and this one induces stable dynamics, but these are different from the actual evolution of X unless $\bar{\alpha} = 0$. See Borovička et al. (2011) for more discussion of these issues.

Our next example illustrates that even when the macroeconomy is modeled as stationary, the stochastic discount factor can still inherit a martingale component.

Example 5.5. *Here we illustrate a point made by Alvarez and Jermann (2005) concerning recursive utility of the type featured by Kreps and Porteus (1978) and Epstein and Zin (1989).*

For an infinite-horizon discrete-time economy, The recursive utility stochastic discount factor in discrete time obeys:

$$\frac{S_{t+1}}{S_t} = \exp(-\delta) \left(\frac{C_{t+1}}{C_t}\right)^{-\rho} \left[\frac{V_{t+1}}{\mathcal{R}_t\left(V_{t+1}\right)}\right]^{\rho-\gamma}$$

where $\rho > 0$ and $\frac{1}{\rho}$ is the elasticity of intertemporal substitution, $\delta > 0$ is a subjective rate of discount and γ adjusts the for the riskiness in the continuation value V_{t+1} in the formula:

$$\mathcal{R}_t\left(V_{t+1}\right) = \left(E\left[\left.(V_{t+1})^{1-\gamma}\right|\mathcal{F}_t\right]\right)^{\frac{1}{1-\gamma}}.$$

Suppose that $C = k(X) > 0$ implying a stationary consumption process, that $\gamma > 1$ and for simplicity suppose that $\rho = 1$.[5] Then

$$\left[\frac{V_{t+1}}{\mathcal{R}_t\left(V_{t+1}\right)}\right]^{1-\gamma} = \frac{(V_{t+1})^{1-\gamma}}{E\left[(V_{t+1})^{1-\gamma}|\mathcal{F}_t\right]} \tag{16}$$

which has conditional expectation equal to unity and is the increment to a positive martingale. For this specification V_{t+1} can be expressed as a time invariant function of the Markov state X_{t+1} and the risk adjustment term $\mathcal{R}_t\left(V_{t+1}\right)$, a time invariant function of X_t. Thus the ratio on the right-hand side of (16) is the increment to a positive multiplicative martingale. Even though C is stationary, the stochastic discount factor process has a positive martingale component. While we have not shown that this martingale necessarily induces stable dynamics, this may be checked in actual applications.

Often, $\log C$ is modeled as a process with stationary increments. In this case the ratio $\frac{V_t}{C_t}$ is computed as a time-invariant function of the Markov state X_t. Both consumption growth and the ratio of the continuation value to its risk-adjusted counterpart contribute to the increment to a martingale component of S.

5.3   Additional insights about the factorization

One of the aims of Hansen and Scheinkman (2009) was to provide a formal justification for the factorization of the stochastic discount factor process featured by Alvarez and Jermann (2005):

$$\frac{S_t}{S_0} = \exp(\eta t) \left(\frac{\widetilde{H}_t}{\widetilde{H}_0}\right) \left(\frac{e(X_0)}{e(X_t)}\right)$$

---

[5]See Proposition 9 of Alvarez and Jermann (2005) for a closely related example in which consumption is independent and identically distributed and $\rho \neq \gamma$.

Sometimes this factorization may be difficult to compute. Nevertheless, there are circumstances in which we can still draw conclusions without a full characterization.

5.3.1 Permanent shocks

Alvarez and Jermann (2005) use the martingale component as evidence for the impact of permanent shocks to the stochastic discount factor process. To understand this link more fully consider an alternative martingale extraction that is familiar from time series analysis. Recall that $\log M_t$ has stationary increments. Often it is most convenient to compute or model $\log M$ as a precursor to the study of M. The martingale component of $\log M$ can be computed from the decomposition

$$\log S_t = \rho t + \log \widehat{H}_t - g(X_t) + g(X_0)$$

and reveals a permanent shock $\log \widehat{H}_t$ up to a scale normalization. When $\log S - \log S_0$ has a non-degenerate martingale component with stationary increments, then S/S_0 has a multiplicative martingale component and *vice versa*. When M is globally log-normal, there are simple links between these martingales; but outside the log-normal environment this link may cease to be direct. In general, the difference $\eta - \rho$ between the extracted drift terms is linked to the average relative entropy of M, which reflects the magnitude of the martingale component:

$$
\begin{aligned}
\eta - \rho &\equiv \lim_{t \to \infty} \frac{1}{t} \left[\log E \left(\frac{S_t}{S_0} \mid X_0 = x \right) - E \left(\log \frac{S_t}{S_0} \mid X_0 = x \right) \right] \\
&= -E \left[\log \frac{\widetilde{H}_{t+1}}{\widetilde{H}_t} \right].
\end{aligned}
$$

Thus the difference between η and ρ from two different martingale constructions reveal the average of the one-period log-likelihood ratio between the original probability measure and the $\widetilde{}$ probability measure.[6]

5.3.2 A more general factorization

Bansal and Lehmann (1997) note that in many examples the stochastic discount factor can be expressed as:
$$\frac{S_t}{S_0} = \left(\frac{S_t^*}{S_0^*} \right) \left[\frac{m^*(X_t)}{m^*(X_0)} \right] \tag{17}$$

---

[6]See Hansen (2012) for a further elaboration of this relationship, and see Backus et al. (2014) for discussion of a term structure counterpart applicable to all investment horizons.

where S^*/S_0^* is a positive multiplicative process with a direct interpretation. Sometimes, it will be easy to characterize the martingale component of S^*. The stochastic discount factor process S will have the same martingale component as S^* with the same Perron–Frobenius eigenvalue but with a different Perron–Frobenius eigenfunction. This formula gives us a way to represent implications of a class of models. Suppose that two models share the same same S^* but differ in m^*. Both models will inherit the long-horizon pricing properties from S^* although risk-return tradeoffs over short investment horizons could be quite different. Consider two illustrations. The stochastic discount factor for the Campbell–Cochrane model in Example 5.4 can be written as:

$$\frac{S_t}{S_0} = \exp\left[\bar{\beta}t + \int_0^t \bar{\alpha}dW_u\right]\left[\frac{k(X_t)}{k(X_0)}\right],$$

and thus satisfies (17) with S^* being a geometric Brownian and $m^* = k$. The contribution of the geometric Brownian motion can be interpreted as a marginal utility process for a corresponding preference specification that abstracts from the contribution of external habits.

Hansen (2012) uses this approach to compare a model with power utility preferences with risk-aversion coefficient γ to a more general recursive utility model from Example 5.5 with $\rho \neq \gamma$.[7] Both preference specifications share the same martingale component holding fixed the consumption process provided that the subjective rate of discount is pushed to an appropriately defined limiting value.

5.3.3 Risk-return tradeoff and welfare consequences

Since the prices of all Arrow claims are taken as given, the recovery of an incorrect probability does not affect the pricing of any derivative security. The recovery of an incorrect probability measure *does* lead to incorrect inference about expected returns and as a consequence to a misspecification of the risk-return tradeoff faced by investors. Incorrect inference about expected cash flows affect the calculation of the welfare cost of economic fluctuations, since Hansen et al. (1999) and Alvarez and Jermann (2004) showed a link between local measures of welfare cost of uncertainty and prices and expectations of the appropriately constructed cash flows.

---

[7]See Theorem 9.1 of Hansen (2012).

6   Measuring the martingale component

We next consider methods for extracting evidence from asset market data that supports the existence of a non-trivial martingale component in stochastic discount factors. We build on the approach initiated by Hansen and Jagannathan (1991) aimed at non-parametric characterizations of stochastic discount factors without using a full set of Arrow prices. In such circumstances, full identification is not possible; but nevertheless the data from financial markets can be informative. For this discussion we use pedagogically useful characterization of Almeida and Garcia (2013) and Hansen (2014), but adapted to misspecified beliefs along the lines suggested in Gosh et al. (2012) and Hansen (2014). In so doing we build on a key insight of Alvarez and Jermann (2005).

Consider strictly convex functions ϕ_θ defined on the positive real numbers such as:

$$\phi_\theta(\mathsf{r}) = \frac{1}{\theta(1+\theta)} \left[(\mathsf{r})^{1+\theta} - 1 \right] \tag{18}$$

for alternative choices of the parameter θ. By design $\phi_\theta(1) = 0$ and $\phi_\theta''(1) = 1$. The function ϕ_θ remains well defined for $\theta = 0$ and $\theta = -1$ by taking pointwise limits in r as θ approaches these two values. Thus $\phi_0(\mathsf{r}) = \mathsf{r} \log \mathsf{r}$ and $\phi_{-1}(\mathsf{r}) = -\log \mathsf{r}$. The functions ϕ_θ are used to construct discrepancy measures between probability densities as in the work of Cressie and Read (1984). We are interested in such measures as a way to quantify the martingale component to stochastic discount factors. Recall that

$$E\left[\frac{\widetilde{H}_{t+1}}{\widetilde{H}_t} | X_t = x \right] = 1$$

and that $\widetilde{H}_{t+1}/\widetilde{H}_t$ defines a conditional density of the $\widetilde{P}$ distribution relative to the P distribution. This leads us to apply the discrepancy measures to $\widetilde{H}_{t+1}/\widetilde{H}_t$.

Since ϕ_θ is strictly convex and $\phi_\theta(1) = 0$, from Jensen's inequality:

$$E\left[\phi_\theta\left(\frac{\widetilde{H}_{t+1}}{\widetilde{H}_t} \right) | X_t = x \right] \geq 0,$$

with equality only when $\widetilde{H}_{t+1}/\widetilde{H}_t$ is identically one. There are three special cases that receive particular attention.

i) $\theta = 1$ in which case the implied measure of discrepancy is equal to one-half times the conditional variance of $\widetilde{H}_{t+1}/\widetilde{H}_t$;

ii) $\theta = 0$ in which case the implied measure of discrepancy is based on conditional relative

entropy:

$$E\left[\left(\widetilde{H}_{t+1}/\widetilde{H}_t\right)\left(\log \widetilde{H}_{t+1} - \log \widetilde{H}_t\right)|X_t = x\right]$$

which is the expected log-likelihood under the $\widetilde{P}$ probability measure.

iii) $\theta = -1$ in which case the discrepancy measure is:

$$-E\left[\log \widetilde{H}_{t+1} - \log \widetilde{H}_t | X_t = x\right]$$

which the negative of the expected log-likelihood under the original probability measure.

We describe how to compute lower bounds for these discrepancy measures. We are led to the study lower bounds because we prefer not to compel an econometrician to use a full array of Arrow prices. Let Y_{t+1} be a vector of asset payoffs and Q_t the corresponding vector of prices. Recall the formula

$$R^{\infty}_{t,t+1} = \exp(-\eta)\frac{e(X_{t+1})}{e(X_t)}.$$

and thus

$$\frac{S_{t+1}}{S_t} = \left(\frac{\widetilde{H}_{t+1}}{\widetilde{H}_t}\right)\left(\frac{1}{R^{\infty}_{t,t+1}}\right). \tag{19}$$

As in Alvarez and Jermann (2005), suppose that the limiting holding-period return $R^{\infty}_{t,t+1}$ can be well approximated. In this case, one could *test* directly for the absence of the martingale component by assessing whether

$$E\left[\left(\frac{1}{R^{\infty}_{t,t+1}}\right)(Y_{t+1})|X_t = x\right] = (Q_t)'.$$

since in this case:

$$\frac{S_{t+1}}{S_t} = \left(\frac{1}{R^{\infty}_{t,t+1}}\right)$$

prices assets correctly.

More generally, we express the pricing restrictions as

$$E\left[\left(\frac{\widetilde{H}_{t+1}}{\widetilde{H}_t}\right)\left(\frac{1}{R^{\infty}_{t,t+1}}\right)(Y_{t+1})'|X_t = x\right] = (Q_t)',$$

where $\widetilde{H}$ is now treated as unobservable to an econometrician. To bound a discrepancy

measure, let a random variable J_{t+1} be a potential specification for the martingale increment:

$$J_{t+1} = \frac{\widetilde{H}_{t+1}}{\widetilde{H}_t}.$$

Solve

$$\lambda_\theta(x) = \inf_{J_{t+1}>0} E\left[\phi_\theta\left(J_{t+1}\right)|X_t = x\right]$$

subject to the linear constraints:

$$E\left[J_{t+1}|X_t = x\right] - 1 = 0$$

$$E\left[J_{t+1}\left(\frac{1}{R_{t,t+1}^\infty}\right)(Y_{t+1})'|X_t = x\right] - (Q_t)' = 0$$

A strictly positive $\lambda_\theta(x)$ implies a nontrivial martingale component to the stochastic discount factor.

For the limiting continuous-time diffusion case, the choice of θ is irrelevant. Suppose that

$$d\log\widetilde{H}_t = -\frac{1}{2}|\tilde{\alpha}(X_t)|^2\,dt + \tilde{\alpha}(X_t)\cdot dW_t.$$

As a consequence,

$$\frac{d\widetilde{H}_t}{\widetilde{H}_t} = \tilde{\alpha}(X_t)\cdot dW_t$$

Thus the local mean of $\left(\widetilde{H}\right)^{\theta+1}$ is

$$\frac{\theta(\theta+1)}{2}\left(\widetilde{H}_t\right)^{\theta+1}|\tilde{\alpha}(X_t)|^2$$

and the discrepancy measure is $\frac{1}{2}|\tilde{\alpha}(X_t)|^2$ independent of θ. The discrepancies for all values of θ are equal to one-half times the local variance of $\log\widetilde{H}$. This equivalence of the discrepancy measures is special to the continuous-time diffusion model, however.

To compute λ_θ in practice requires that we estimate conditional distributions. There is an unconditional counterpart to these calculations obtained by solving:

$$\bar{\lambda}_\theta = \inf_{J_{t+1}>0} E\phi_\theta(J_{t+1}) \tag{20}$$

subject to:

$$E\left[J_{t+1}\right] - 1 = 0$$

$$E\left[J_{t+1}\left(\frac{1}{R_{t,t+1}^{\infty}}\right)(Y_{t+1})' - (Q_t)'\right] = 0 \qquad (21)$$

This bound, while more tractable, is weaker in the sense that $\bar{\lambda}_{\theta} \leq E\lambda_{\theta}(X_t)$. To guarantee a solution to optimization problem (20) it is sometimes convenient to include random variables J_{t+1} that are zero with positive probability. Since the aim is to produce bounds, this augmentation can be justified for mathematical and computational convenience. Although this problem optimizes over an infinite-dimensional family of random variables J_{t+1}, the dual problem that optimizes over the Lagrange multipliers associated with the pricing constraint (21) is often quite tractable. See Hansen et al. (1995) and the literature on implementing generalized empirical likelihood methods for further discussion.

For the case in which $\theta = 1$, Hansen and Jagannathan (1991) study a mathematically equivalent problem by constructing volatility bounds for stochastic discount factors and deduce quasi-analytical formulas for the solution obtained when ignoring the restriction that stochastic discount factors should be nonnegative. Bakshi and Chabi-Yo (2012) apply the latter methods to obtain $\theta = 1$ bounds (volatility bounds) for the martingale component to the stochastic discount factor process. Similarly, Bansal and Lehmann (1997) study bounds on the stochastic discount factor process for the case in which $\theta = -1$ and show the connection with a maximum growth rate portfolio. Alvarez and Jermann (2005) apply these methods to produce the corresponding bounds for the martingale component to the stochastic discount factor process. Both Alvarez and Jermann (2005) and Bakshi and Chabi-Yo (2012) exploit equation (19) and approximate the return $R_{t,t+1}^{\infty}$ in order to target their analysis to the martingale component. These papers provide empirical evidence in support of a substantial martingale component to the stochastic discount factor process. Bakshi and Chabi-Yo (2012) summarize results from both papers in their Table 1 and contrast differences in the $\theta = 1$ and $\theta = -1$ discrepancy measures. To our knowledge, the $\theta = 0$ discrepancy measure has not been used to quantify the magnitude of the martingale component to the stochastic discount factor processes.

7   Conclusion

Perron–Frobenius theory applies to an operator used to represent Arrow prices. Recent research uses this insight in varied ways. As we have argued, when the state-space is contin-

uous, the positive eigenfunction of the pricing operator may not be unique (up to scale), and further assumptions are needed to determine a unique decomposition. In our previous work, Hansen and Scheinkman (2009), we were interested in long-run risk and for this reason we emphasized a stochastic stability condition. In this paper we gave a generalization of that stability condition to establish uniqueness by replacing the convergence of a sequence of conditional expectations by the convergence of partial averages of that sequence. We show that this more general condition is implied by an ergodicity assumption commonly used in the study of rational expectations models of financial markets.

Other researchers have taken different approaches to the uniqueness. In a continuous-time Brownian information setup, alternative conditions on the boundary behavior of the underlying Markov process also uniquely identify a probability measure. These conditions often utilize linkages of the Perron–Frobenius theorem to the Sturm–Liouville problem in the theory of second-order differential equations. Carr and Yu (2012) and Dubynskiy and Goldstein (2013) impose conditions on reflecting boundaries, while Walden (2013) analyzes natural boundaries. A recent working paper by Qin and Linetsky (2014) provides a more comprehensive treatment of the continuous state space environment, connections to our own previous results in Hansen and Scheinkman (2009) and Hansen and Scheinkman (2014), and an extensive range of examples utilizing models in mathematical finance.

While these technical conditions deliver uniqueness, we argue that there is a more fundamental identification problem. If the stochastic discount factor includes a martingale component, then use of the Perron–Frobenius eigenvalue and function recovers a distorted probability measure that provides insights into the pricing of long-term bonds and other cash flows that do not grow stochastically over time. Many structural models of asset pricing that are motivated by empirical evidence have martingale components in stochastic discount factor processes. These martingales characterize what probability is actually recovered by application of Perron-Frobenius theory.

While fixed income security valuation is of interest in its own right, Hansen and Scheinkman (2009) and Hansen (2012) use Perron–Frobenius methods to study more general valuation problems in which cash flows can grow stochastically over time. These extensions are critical for many macro-finance applications because, empirically, many macro time series display stochastic growth. Long-term valuation is only a component to a more systematic study of pricing implications over alternative investment horizons. Recent work by Borovička et al. (2011) and Borovička et al. (2014) deduces methods that extend impulse response functions to characterize the pricing of exposures to shocks to stochastically growing cash flows over alternative investment horizons.

Appendix

A   Perron–Frobenius theory

Proof of Proposition 4.4. Write

$$S^H = \exp\left(-\widehat{\delta}t\right) \frac{\widehat{m}(X_t)}{\widehat{m}(X_0)},$$

for a positive function $\widehat{m}$. Thus

$$\frac{H_t}{H_0} = \exp\left[-\left(\delta - \widehat{\delta}\right)t\right] \frac{\frac{m(X_t)}{\widehat{m}(X_t)}}{\frac{m(X_0)}{\widehat{m}(X_0)}} := \exp\left(\eta t\right) \frac{k(X_t)}{k(X_0)},$$

where $\eta = \delta - \widehat{\delta}$ and $k = \frac{m}{\widehat{m}}$. In what follows we consider the discrete time case. The continuous time case uses an identical approach, with the obvious changes.

Consider three cases. First suppose that $\eta > 0$. Since $\frac{H_t}{H_0}$ is a martingale,

$$E\left[k(X_t)|X_0 = x\right] = \exp(-\eta t)k(x) \tag{22}$$

Form

$$\hat{k}(x) = \min\{1, k(x)\} > 0, \text{ for all } x.$$

Since $\eta > 0$, the right-hand side of (22) converges to zero for each x as $t \to \infty$. Thus,

$$0 = \lim_{N \to \infty} \frac{1}{N} \sum_{t=1}^{N} E\left[k(X_t)|X_0 = x\right] \geq \lim_{N \to \infty} \frac{1}{N} \sum_{t=1}^{N} E\left[\hat{k}(X_t)|X_0 = x\right] = E\left[\hat{k}(X_0)\right] > 0.$$

Thus we have established a contradiction.

Next suppose that $\eta < 0$. Note that

$$E^H\left[\frac{1}{k(X_t)}|X_0 = x\right] = \exp(\eta t)\frac{1}{k(x)}. \tag{23}$$

Form

$$\widetilde{k}(x) = \min\left\{1, \frac{1}{k(x)}\right\} > 0, \text{ for all } x.$$

Since $\eta < 0$, the right-hand side of (23) converges to zero for each x as $t \to \infty$. Thus,

$$0 = \lim_{N \to \infty} \frac{1}{N} \sum_{t=1}^{N} E^H\left[\frac{1}{k(X_t)}|X_0 = x\right] \geq \lim_{N \to \infty} \frac{1}{N} \sum_{t=1}^{N} E^H\left[\widetilde{k}(X_t)|X_0 = x\right] = E^H\left[\widetilde{k}(X_0)\right] > 0.$$

We have again established a contradiction.

Finally, suppose $\eta = 0$. Then

$$E^H \left[\frac{1}{k(X_t)} | X_0 = x \right] = \frac{1}{k(x)}$$

for all x. Since X is stochastically stable in averages:

$$\lim_{N \to \infty} \frac{1}{N} \sum_{t=1}^{N} E^H [k^n(X_t) | X_0 = x] = E^H k^n(X_0).$$

for $k^n = \min\{1/k, n\}$. The same equality applies to the limit as $n \to \infty$ whereby k^n is replaced by $1/k$. Consequently, $E^H \left[\frac{1}{k(X_0)} \right] = \frac{1}{k(x)}$ for almost all x. Thus $k(x)$ is a constant, and $H_t/H_0 \equiv 1$ with probability one. $\square$

References

Almeida, Caio and René Garcia. 2013. Robust Economic Implications of Nonlinear Pricing Kernels.

Alvarez, Fernando and Urban J. Jermann. 2004. Using Asset Prices to Measure the Cost of Business Cycles. *Journal of Political Economy* 112 (6):1223–1256.

———. 2005. Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth. *Econometrica* 73 (6):1977–2016.

Backus, David K., Allan W. Gregory, and Stanley E. Zin. 1989. Risk Premiums in the Term Structure : Evidence from Artificial Economies. *Journal of Monetary Economics* 24 (3):371–399.

Backus, David K., Mikhail Chernov, and Stanley E. Zin. 2014. Sources of Entropy in Representative Agent Models. *The Journal of Finance* 59 (1):51–99.

Bakshi, Gurdip and Fousseni Chabi-Yo. 2012. Variance Bounds on the Permanent and Transitory Components of Stochastic Discount Factors. *Journal of Financial Economics* 105 (1):191–208.

Bansal, Ravi and Bruce N. Lehmann. 1997. Growth-optimal Portfolio Restrictions on Asset Pricing Models. *Macroeconomic Dynamics* 1 (2):333–354.

Borovička, Jaroslav, Lars Peter Hansen, Mark Hendricks, and José A. Scheinkman. 2011. Risk Price Dynamics. *Journal of Financial Econometrics* 9:3–65.

Borovička, Jaroslav, Lars Peter Hansen, and José A. Scheinkman. 2014. Shock Elasticities and Impulse Responses. Mimeo.

Campbell, John Y. and John H. Cochrane. 1999. By Force of Habit. *Journal of Political Economy* 107 (2):205–251.

Carr, Peter and Jiming Yu. 2012. Risk, Return and Ross Recovery. *The Journal of Derivatives* 20 (1):38–59.

Cressie, Noel and Timothy R. C. Read. 1984. Multinomial Goodness-of-Fit Tests. *Journal of the Royal Statistical Society. Series B (Methodological)* 46 (3):440–464.

Dubynskiy, Sergey and Robert S. Goldstein. 2013. Recovering Drifts and Preference Parameters from Financial Derivatives.

Duffie, Darrell and Rui Kan. 1994. Multi-Factor Term Structure Models. *Philosophical Transactions: Physical Sciences and Engineering* 347 (1684):577–586.

Epstein, Larry G. and Stanley E. Zin. 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica* 57 (4):937–969.

Gosh, Anisha, Christian Julliard, and Alex Taylor. 2012. What is the Consumption-CAPM Missing? An Information-Theoretic Framework for the Analysis of Asset Pricing Models. Working paper, Tepper School of Business, Carnegie Mellon University.

Hansen, Lars Peter. 2012. Dynamic Valuation Decomposition within Stochastic Economies. *Econometrica* 80 (3):911–967. Fisher–Schultz Lecture at the European Meetings of the Econometric Society.

———. 2014. Uncertainty Outside and Inside Economic Models. Nobel Prize paper.

Hansen, Lars Peter and Ravi Jagannathan. 1991. Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy* 99 (2):225–262.

Hansen, Lars Peter and Scott F. Richard. 1987. The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models. *Econometrica* 55 (3):587–613.

Hansen, Lars Peter and José A. Scheinkman. 2009. Long Term Risk: An Operator Approach. *Econometrica* 77 (1):177–234.

———. 2014. Stochastic Compounding and Uncertain Valuation.

Hansen, Lars Peter and Kenneth J. Singleton. 1982. Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models. *Econometrica* 50 (5):1269–1286.

Hansen, Lars Peter, John Heaton, and Erzo G. J. Luttmer. 1995. Econometric Evaluation of Asset Pricing Models. *The Review of Financial Studies* 8 (2):237–274.

Hansen, Lars Peter, Thomas J. Sargent, and Jr. Tallarini, Thomas D. 1999. Robust Permanent Income and Pricing. *The Review of Economic Studies* 66 (4):873–907.

Kreps, David M. and Evan L. Porteus. 1978. Temporal Resolution of Uncertainty and Dynamic Choice Theory. *Econometrica* 46 (1):185–200.

Qin, Likuan and Vadim Linetsky. 2014. Ross Recovery in Continuous Time. Mimeo, Northwestern University.

Ross, Stephen A. 2013. The Recovery Theorem. Forthcoming in the Journal of Finance.

Walden, Johan. 2013. Recovery with Unbounded Diffusion Processes.