

Misspecified Recovery*

Jaroslav Borovička

Lars Peter Hansen

New York University

University of Chicago and NBER

jaroslav.borovicka@nyu.edu

lhansen@uchicago.edu

José A. Scheinkman

Columbia University, Princeton University and NBER

joses@princeton.edu

July 2, 2015

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. To understand this identification challenge, we extract a positive martingale component from the stochastic discount factor process using Perron–Frobenius Theory. When this martingale is degenerate, probabilities that govern investor beliefs are recovered from the prices of Arrow securities. When the martingale component is not trivial, using this same approach recovers a probability measure, but not the one that is used by investors. We refer to this outcome as “misspecified recovery.” We show that the resulting misspecified probability measure absorbs long-term risk adjustments. Many structural models of asset prices have stochastic discount factors with martingale components. Also empirical evidence on asset prices suggests that the recovered measure differs from the actual probability distribution. While this probability measure is of substantive interest, interpreting it as the true probability distribution may bias our inference about investor aversion to risk as reflected in risk-return tradeoffs.

*We thank Fernando Alvarez, David Backus, Ravi Bansal, Anmol Bhandari, Peter Carr, Xiaohong Chen, Ing-Haw Cheng, Mikhail Chernov, Kyle Jurado, Stavros Panageas, Karthik Sastry, Kenneth Singleton, Johan Walden, Wei Xiong and the anonymous referees for useful comments.

1 Introduction

Asset prices are forward looking and encode information about investors' beliefs. This leads researchers and policy makers to look at financial market data to gauge the views of the private sector about the future of the macroeconomy. It has been known, at least since the path-breaking work of Arrow, that asset prices reflect a combination of investors' risk aversion and the probability distributions used to assess risk. In dynamic models, investors' risk aversion is expressed by stochastic discount factors that encode compensations for risk exposures. In this paper, we ask what can be learned from the Arrow prices about investors' beliefs. Data on asset prices alone are not sufficient to identify both the stochastic discount factor and transition probabilities without imposing additional restrictions. This additional information could be time series evidence on the evolution of the Markov state, or it could be information on the market-determined stochastic discount factors.

In a Markovian environment, Perron–Frobenius Theory selects a single transition probability compatible with asset prices. This Perron–Frobenius apparatus has been used in previous research in at least two manners. First, [Hansen and Scheinkman \(2009\)](#) showed that this tool identifies a probability measure that reflects the long-term implications for risk pricing under rational expectations. We refer to this probability as the long-term risk neutral probability since the use of this measure renders the long-term risk-return tradeoffs degenerate. [Hansen and Scheinkman \(2009\)](#) purposefully distinguish this constructed transition probability from the underlying time series evolution. The ratio of the recovered to the true probability measure is manifested as a non-trivial martingale component in the stochastic discount factor process. Second, [Ross \(2015\)](#) applied Perron–Frobenius Theory to identify or to “recover” investors' beliefs. Interestingly, this recovery does not impose rational expectations, thus the resulting Markov evolution could reflect subjective beliefs of investors and not necessarily the actual time series evolution.

In this paper we delineate the connection between these seemingly disparate results. We make clear the special assumptions that are needed to guarantee that the transition probabilities recovered using Perron–Frobenius Theory are equal to the subjective transition probabilities of investors or to the actual probabilities under an assumption of rational expectations. An extension of the Perron–Frobenius approach to include seemingly unnecessary state variables avoids the degeneracy in the implied long-term risk-return tradeoff. Unfortunately, full identification of the subjective beliefs is no longer feasible.

Section 2 illustrates the challenge of identifying the correct probability measure from asset market data in a finite-state space environment. While the finite-state Markov environment is too constraining for many applications, the discussion in this section provides an overview of some of the main results in this paper. In particular, we show that:

- the Perron–Frobenius approach recovers a probability measure that reflects long-term pricing;
- the relative density of the Perron–Frobenius probability to the physical probability gives rise to a martingale component to the stochastic discount factor process;

- the stochastic discount factor process used by Ross (2015) implies that this martingale component is a constant.

To place these results in a substantive context, we provide prototypical examples of asset pricing models that show that a nontrivial martingale component arises from (i) permanent shocks to the consumption process or (ii) continuation value adjustment that appear when agents have recursive utilities.

In subsequent sections we establish these insights in greater generality, a generality rich enough to include many existing structural Markovian models of asset pricing. The framework for this analysis, which allows for continuous state spaces and a richer information structure, is introduced in Section 3. In Section 4, we extend the Perron–Frobenius approach to this more general setting. Provided we impose an additional ergodicity condition, this approach identifies a unique probability measure captured by a martingale component to the stochastic discount factor process. In Section 5, we show the consequences of using the probability measure recovered by the use of the Perron–Frobenius Theory when making inferences on the risk–return tradeoff. We show how the recovered probability measure absorbs long-term risk adjustments and implies that assets are priced *as if* long-run risk premia were trivial. We are thus lead to call this probability specification the long-term risk neutral measure. Section 6 illustrates the impact of a martingale component to the stochastic discount factor in a workhorse asset pricing model that features long-run risk.

Starting in Section 7, we characterize the challenges in identifying subjective beliefs from asset prices. Initially we pose the *fundamental identification problem*: data on asset prices can only identify the stochastic discount factor up to an arbitrary strictly positive martingale, and thus the probability measure associated with a stochastic discount factor remains unidentified without imposing additional restrictions or using additional data. We also extend the analysis of Ross (2015) to this more general setting. By connecting to the results in Section 4, we demonstrate in Section 7 that the martingale component to the stochastic discount factor process must be identically equal to one for Ross (2015)’s procedure to reveal the subjective beliefs of investors. Under these beliefs, the long-term risk–return tradeoff is degenerate. In Section 8 we extend the Perron–Frobenius problem to encompass a larger state vector. This extension avoids the degeneracy in the long-term risk–return tradeoff, but it also eliminates the ability to identify subjective beliefs from asset prices without imposing rational expectations. In Section 9, we provide a unifying discussion of the empirical approaches that quantify the impact of the martingale components to stochastic discount factors when an econometrician does not use the full array of Arrow prices. We also suggest other approaches that connect subjective beliefs to the actual time series evolution of the Markov states. Section 10 concludes.

2 Illustrating the identification challenge

There are multiple approaches for extracting probabilities from asset prices. For instance, risk neutral probabilities (e.g., see Ross (1978) and Harrison and Kreps (1979)) and closely related

forward measures are frequently used in financial engineering. More recently, Perron–Frobenius Theory has been applied by [Backus et al. \(1989\)](#), [Hansen and Scheinkman \(2009\)](#) and [Ross \(2015\)](#) to study asset pricing — the last two references featuring the construction of an associated probability measure. [Hansen and Scheinkman \(2009\)](#) and [Ross \(2015\)](#) have rather different interpretations of this measure, however. [Ross \(2015\)](#) identifies this measure with investors’ beliefs while [Hansen and Scheinkman \(2009\)](#) use it to characterize long-term contributions to risk pricing. Under this second interpretation, Perron–Frobenius Theory features an eigenvalue that *dominates* valuation over long investment horizons, and the resulting probability measure targets the valuation of assets that pay off in the far future as a point of reference. Following [Hansen and Scheinkman \(2014\)](#), in this section we illustrate the construction of the alternative probability measure using matrices associated with finite-state Markov chains and we explore some simple example economies to understand better the construction of a probability measure based on Perron–Frobenius Theory.

Let X be a discrete time, n -state Markov chain with transition matrix $\mathbf{P} \doteq [p_{ij}]$. We suppose that these are the actual transition probabilities that govern the evolution of the Markov process. We identify state i with a coordinate vector u^i with a single one in the i -th entry. The analyst infers the prices of one-period Arrow claims from data. We represent this input as a matrix $\mathbf{Q} = [q_{ij}]$ that maps a payoff tomorrow specified as a function of tomorrow’s state j into a price in state i today. Since there are only a finite number of states, the payoff and price can both be represented as vectors. In particular, the vector of Arrow prices given the current state x is $x'\mathbf{Q}$. The entries of this vector give the prices of claims payable in each of the possible states tomorrow. Any state that cannot be realized tomorrow given the current state x is assigned a price of zero today.

Asset pricing implications are represented conveniently using stochastic discount factors. Stochastic discount factor encode adjustments for uncertainty by discounting the next-period state differentially. Risk premia are larger for states that are more heavily discounted. In this finite-state Markov environment, we compute

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

provided that $p_{ij} > 0$. The definition of s_{ij} is inconsequential if $p_{ij} = 0$. Given a matrix $\mathbf{S} = [s_{ij}]$, the stochastic discount factor process has the increment:

$$\frac{S_{t+1}}{S_t} = (X_t)'\mathbf{S}X_{t+1}. \tag{2}$$

The stochastic discount factor process $S = \{S_t : t = 0, 1, 2, \dots\}$ is initialized at $S_0 = 1$ and accumulates the increments given by (2):

$$S_t = \prod_{\tau=1}^t (X_{\tau-1})'\mathbf{S}X_{\tau}.$$

Observe that S_t depends on the history of the state from 0 to t . With this notation, we have two

ways to write the period-zero price of a claim to a vector of state-dependent payoffs $f \cdot X_t$ at time t .

$$\mathbf{Q}^t f \cdot x = E[S_t(f \cdot X_t) \mid X_0 = x].$$

Given the matrix \mathbf{P} , possibly determined under rational expectations by historical data and stochastic discount factors implied by an economic model, Arrow prices are given by inverting equation (1):

$$q_{ij} = s_{ij}p_{ij}. \tag{3}$$

A question that we explore is what we can learn about beliefs from Arrow prices. Market sentiments or beliefs are part of the discourse for both public and private sectors. We study this question by replacing the assumption of rational expectations with an assumption of subjective beliefs.

Unfortunately, there is considerable flexibility in constructing probabilities from the Arrow prices alone. Notice that \mathbf{Q} has $n \times n$ entries. \mathbf{P} has $n \times (n - 1)$ free entries because row sums have to add up to one. In general the stochastic discount factor introduces $n \times n$ free parameters s_{ij} , $i, j = 1, \dots, n$. Since the Arrow prices are the products given in formula (3), there are multiple solutions for probabilities and stochastic discount factors that are consistent with Arrow prices.¹

To depict this flexibility, represent alternative transition probabilities by

$$\tilde{p}_{ij} = h_{ij}p_{ij}$$

where $h_{ij} > 0$ and $\sum_{j=1}^n h_{ij}p_{ij} = 1$ for $i = 1, 2, \dots, n$. Form a matrix $\mathbf{H} = [h_{ij}]$ and a positive process $\{H_t : t = 0, 1, \dots\}$ with increments

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

The restrictions on the entries of \mathbf{H} restrict the increments to satisfy

$$E \left[\frac{H_{t+1}}{H_t} \mid X_t = x \right] = 1.$$

Accumulating the increments:

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

The initial distribution of X_0 together with the transition matrix \mathbf{P} define a probability P over realizations of the process X . 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 .

¹The 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.

Using the positive martingale H to induce a change of measure, we obtain the probability \tilde{P} :

$$\tilde{P}(X^t = x^t) = P(X^t = x^t)H_0 \prod_{\tau=1}^t (x_{\tau-1})' \mathbf{H} x_{\tau}.$$

for alternative possible realizations $x^t = (x_0, x_1, \dots, x_t)$ of $X^t = (X_0, X_1, \dots, X_t)$. In this formula, we presume that $EH_0 = 1$, and we use H_0 in order for \tilde{P} to include a change in the initial distribution of X_0 . Thus the random variable H_0 modifies the distribution of X_0 under \tilde{P} *vis-à-vis* P , and $\tilde{\mathbf{P}}$ specifies the altered transition probabilities. Most of our analysis conditions on X_0 , in which case the choice of H_0 is inconsequential and H_0 can be set to one for simplicity.

For each choice of the restricted matrix \mathbf{H} , we may form the corresponding state-dependent discount factors $\tilde{s}_{ij} = s_{ij}/h_{ij}$ by applying formula (1). By construction

$$q_{ij} = s_{ij}p_{ij} = \tilde{s}_{ij}\tilde{p}_{ij}. \quad (4)$$

Given flexibility in constructing $\mathbf{H} = [h_{ij}]$, we have multiple ways to recover probabilities from Arrow prices.

We may confront this multiplicity by imposing restrictions on the stochastic discount factors. As we shall argue, the resulting constructions provide valuable tools for asset pricing even when these probabilities are not necessarily the same as those used by investors. In what follows we consider two alternative restrictions:

i) *Risk-neutral pricing:*

$$\bar{s}_{i,j} = \bar{q}_i$$

for positive numbers $\bar{q}_i, i = 1, 2, \dots, n$. This restriction exploits the pricing of one-period discount bonds.

ii) *Long-term risk pricing:*

$$\hat{s}_{ij} = \exp(\eta) \frac{m_j}{m_i} \quad (5)$$

for positive numbers $m_i, i = 1, 2, \dots, n$ and a real number η that is typically negative. The m_i 's need only be specified up to a scale factor and the resulting vector can be normalized conveniently. As we will show later, this restriction helps us characterize long-term pricing implications.

In both cases we reduce the number of free parameters in the matrix \mathbf{S} from n^2 to n , making identification of the probabilities possible. As we will see, each approach has an explicit economic interpretation but the matrices of transition probabilities that are recovered do not necessarily coincide with those used by investors or with the actual Markov state dynamics. In the first case the difference between the inferred and true probabilities reflects a martingale that determines the one-period risk adjustments in financial returns. As we show below, in the second case the

difference between the inferred and true probabilities reflects a martingale that determines long-term risk adjustments in pricing stochastically growing cash flows.

2.1 Risk-neutral probabilities

Risk-neutral probabilities are used extensively in the financial engineering literature. These probabilities are a theoretical construct used to absorb the local or one-period risk adjustments and are determined by positing a fictitious “risk-neutral” investor. The resulting stochastic discount factor of this investor is therefore given by

$$\bar{s}_{i,j} = \bar{q}_i,$$

reflecting the fact that all states j tomorrow are discounted equally. In order to satisfy pricing restrictions (3), the risk-neutral transition probabilities must be given by

$$\bar{p}_{ij} = \frac{q_{ij}}{\bar{q}_i}.$$

Since rows of a probability matrix have to sum up to one, it necessarily follows that $\bar{q}_i = \sum_j^n q_{ij}$, which is the price of a one-period discount bond in state i .

The risk-neutral probabilities $[\bar{p}_{ij}]$ can always be constructed and used in conjunction with discount factors $[\bar{s}_{ij}]$. By design 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 and this dependence is absorbed into the subjective discount factor of the fictitious-risk neutral investor. While it is widely recognized that the risk-neutral distribution is distinct from the actual probability distribution, some have argued that the risk-neutral dynamics remain interesting for macroeconomic forecasting precisely because they do embed risk adjustments.²

When short-term interest rates are state-dependent, forward measures are sometimes used in valuation. Prices of t -period Arrow securities, $q_{ij}^{[t]}$, are the entries of the t -th power of the matrix \mathbf{Q} . The t -period forward probability measure given the current state i is

$$\bar{\mathbf{P}}_t = \left[\frac{q_{ij}^{[t]}}{\sum_{j=1}^n q_{ij}^{[t]}} \right].$$

The denominator used for scaling the Arrow prices is now the price of a t -period discount bond. While the forward measure is of direct interest,

$$\bar{\mathbf{P}}_t \neq (\bar{\mathbf{P}})^t, \tag{6}$$

²Narayana Kocherlakota, President of the Federal Reserve Bank of Minneapolis, during a speech to the Mitsui Financial Symposium in 2012 asks and answers: “How can policymakers formulate the needed outlook for marginal net benefits? ... I argue that policymakers can achieve better outcomes by basing their outlooks on *risk-neutral probabilities* derived from the prices of financial derivatives.” See [Hilscher et al. \(2014\)](#) for a study of public debt using risk-neutral probabilities.

when one-period bond prices are state dependent. Variation in one-period interest rates contributes to risk adjustment over longer investment horizons, and as a consequence the construction of risk-neutral probabilities is horizon-dependent.

2.2 Long-term pricing

We study 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 (in absolute value) eigenvalue of \mathbf{Q} is unique and positive and thus can be written as $\exp(\hat{\eta})$, and has a unique associated right eigenvector \hat{e} , which has strictly positive entries. Every non-negative eigenvector of \mathbf{Q} is a scalar multiple of \hat{e} . We denote the i^{th} entry of \hat{e} as \hat{e}_i . Typically, $\hat{\eta} < 0$ to reflect time discounting of future payoffs over long investment horizons.

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

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

where \hat{e}^* is the corresponding positive left eigenvector of \mathbf{Q} . Applying this formula, the large t approximation to the rate of discount on an arbitrary security with positive payoff $f \cdot X_t$ in t periods is $-\hat{\eta}$. Similarly, the one-period holding-period return on this limiting security is:

$$\lim_{t \rightarrow \infty} \frac{\mathbf{Q}^{t-1} f \cdot X_1}{\mathbf{Q}^t f \cdot X_0} = \exp(-\hat{\eta}) \frac{\hat{e} \cdot X_1}{\hat{e} \cdot X_0}$$

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

$$\hat{p}_{ij} = \exp(-\hat{\eta}) q_{ij} \frac{\hat{e}_j}{\hat{e}_i}. \tag{7}$$

Notice that since $\mathbf{Q}\hat{e} = \exp(\hat{\eta})\hat{e}$,

$$\sum_{j=1}^n \hat{p}_{ij} = \exp(-\hat{\eta}) \frac{1}{\hat{e}_i} \sum_{j=1}^n q_{ij} \hat{e}_j = 1.$$

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

$$q_{ij} = \exp(\hat{\eta}) \frac{\hat{e}_i}{\hat{e}_j} \hat{p}_{ij} = \hat{s}_{ij} \hat{p}_{ij}$$

Thus we have used the eigenvector \hat{e} and the eigenvalue $\hat{\eta}$ to construct a stochastic discount factor that satisfies (5) together with a probability measure that satisfies (3). The probability measure constructed in this fashion absorbs the compensations for exposure to long-term components of risk. Conversely, if one starts with an $\hat{\mathbf{S}}$ and $\hat{\mathbf{P}}$ that satisfy (3) and (5) then it is straightforward

to show that the vector with entries $\tilde{e}_i = 1/m_i$ is an eigenvector of \mathbf{Q} .³

If we start with the t -period Arrow prices in the matrix \mathbf{Q}^t instead of the one-period Arrow prices in the matrix \mathbf{Q} , then

$$\mathbf{Q}^t \hat{e} = \exp(\hat{\eta}t) \hat{e}$$

for the same vector \hat{e} and $\hat{\eta}$. The implied matrix $\hat{\mathbf{P}}_t$ constructed from \mathbf{Q}^t satisfies:

$$\hat{\mathbf{P}}_t = \left(\hat{\mathbf{P}}\right)^t.$$

In contrast to the corresponding result (6) for risk-neutral probabilities, Perron–Frobenius Theory recovers the same t -period transition probability regardless whether we use one-period or t -period Arrow claims.

Hansen and Scheinkman (2009) and Ross (2015) 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 \hat{p}_{ij} given by (7) and consider the following decomposition:

$$q_{ij} = \left[\exp(\hat{\eta}) \frac{\hat{e}_i \hat{p}_{ij}}{\hat{e}_j p_{ij}} \right] p_{ij} = \exp(\hat{\eta}) \left(\frac{\hat{e}_i}{\hat{e}_j} \right) \hat{h}_{ij} p_{ij}.$$

Hence,

$$s_{ij} = \exp(\hat{\eta}) \left(\frac{\hat{e}_i}{\hat{e}_j} \right) \hat{h}_{ij} \tag{8}$$

where

$$\hat{h}_{ij} = \hat{p}_{ij}/p_{ij}$$

provided that $p_{ij} > 0$. When $p_{ij} = 0$ the construction of \hat{h}_{ij} is inconsequential.

Hansen and Scheinkman (2009) and Hansen (2012) discuss how the decomposition of the one-period stochastic discount factor displayed on the right-hand side of (8) 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. We call this the long-term risk neutral probabilities. Alternatively, we could follow Ross (2015) and use $\hat{\mathbf{S}} = [\hat{s}_{ij}]$ where

$$\hat{s}_{ij} = \exp(\hat{\eta}) \left(\frac{\hat{e}_i}{\hat{e}_j} \right)$$

to construct the stochastic discount factor process and to let $\hat{\mathbf{P}} = [\hat{p}_{ij}]$ denote the subjective beliefs of the investors for the Markov transition. It is easy to show that \hat{h}_{ij} cannot be written as $\hat{h}_{ij} = \exp(\tilde{\eta}) \tilde{e}_i / \tilde{e}_j$ for some number $\tilde{\eta}$ and a vector \tilde{e} with positive entries, and thus the decomposition in

³To see this, notice that $\tilde{s}_{ij}(1/m_j) = \exp(\eta)(1/m_i)$. The implied probabilities are given by $\hat{p}_{ij} = q_{ij}/\tilde{s}_{ij}$. Pre-multiplying by the probabilities \hat{p}_{ij} , summing over j , and stacking into the vector form, we obtain $\mathbf{Q}\tilde{e} = \exp(\tilde{\eta})\tilde{e}$ for a vector \tilde{e} with entries $\tilde{e}_i = 1/m_i$.

(8) is unique.⁴ In particular we have that:

$$s_{ij} = \exp(\hat{\eta}) \left(\frac{\hat{e}_i}{\hat{e}_j} \right) \text{ for some vector } \hat{e} \iff \hat{h}_{ij} \equiv 1.$$

While the asset price data in \mathbf{Q} uniquely determine $(\hat{\eta}, \hat{e})$ and thus the transition matrix $\hat{\mathbf{P}}$, they contain no information about $[\hat{h}_{ij}]$ and therefore about \mathbf{P} . This highlights the crucial role of Restriction (5). Additional information or assumptions are needed to separate the right-hand side terms in $\hat{p}_{ij} = \hat{h}_{ij} p_{ij}$, and imposing $\hat{h}_{ij} \equiv 1$ provides such an assumption. Throughout the paper, we study the role of $[\hat{h}_{ij}]$ in structural models of asset pricing and ways of identifying it in empirical data.

2.3 Compounding one-period stochastic discounting

An equivalent statement of equation (8) is

$$\frac{S_{t+1}}{S_t} = \exp(\hat{\eta}) \left(\frac{\hat{e} \cdot X_t}{\hat{e} \cdot X_{t+1}} \right) \left(\frac{\hat{H}_{t+1}}{\hat{H}_t} \right).$$

Compounding over time and initializing $S_0 = 1$, we obtain

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

Thus the eigenvalue $\hat{\eta}$ 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 a martingale component \hat{H} , whose logarithm has stationary increments. Imposing restriction (5) on the stochastic discount factor used by investors with subjective beliefs implies that the martingale component under rational expectations is absorbed into the probabilities used by investors. If investors have rational expectations and (5) is not imposed, the martingale implies a change of measure that absorbs the long-term compensations for exposure to growth rate uncertainty.

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.

2.4 Examples

The behavior of underlying shocks is of considerable interest when constructing stochastic equilibrium models. There is substantial time series literature on the role of permanent shocks in multivariate analysis and there is a related macroeconomic literature on models with balanced

⁴For if $\hat{h}_{ij} = \exp(\tilde{\eta}) \tilde{e}_i / \tilde{e}_j$ for some number $\tilde{\eta}$ and a vector \tilde{e} with positive entries, there would exist another Perron–Frobenius eigenvector for \mathbf{Q} with entries given by $\hat{e}_i \tilde{e}_i$ and an eigenvalue $\exp(\hat{\eta} + \tilde{\eta})$. The Perron–Frobenius Theorem guarantees that there is only one eigenvector with strictly positive entries (up to scale) implying that \hat{e} must be a vector of constants and $\hat{\eta} = 0$.

growth behavior, allowing for stochastic growth. The martingale components in stochastic discount factors characterize durable components to risk adjustments in valuation over alternative investment horizons. As we will see, one source for these durable components are permanent shocks to the macroeconomic environment. But valuation models have other sources for this durability including investor preferences. The following examples illustrate that even in this n -state Markov-chain context it is possible to obtain a non-trivial martingale component for the stochastic discount factor.

Example 2.1. *Consumption-based asset pricing models assume that the stochastic discount factor process is a representation of investors' preferences over consumption. Suppose that the growth rate of equilibrium consumption is stationary and that investor preferences can be depicted using a power utility function. For the time being, suppose we impose rational expectations. Thus the marginal rate of substitution is*

$$\exp(-\delta) \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} = \phi(X_{t+1}, X_t).$$

With this formulation, we may write

$$s_{ij} = \phi(X_{t+1} = u_j, X_t = u_i).$$

Stochastic growth in consumption as reflected in the entries s_{ij} will induce a martingale component to the stochastic discount factor. An exception occurs when $C_t = \exp(g_c t)(c \cdot X_t)$ for some vector c with strictly positive entries and a known constant g_c and hence

$$\frac{C_{t+1}}{C_t} = \exp(g_c) \left(\frac{c \cdot X_{t+1}}{c \cdot X_t} \right).$$

Here g_c governs the deterministic growth in consumption and is presumably revealed from time-series data. In this case,

$$s_{ij} = \exp(-\delta - \gamma g_c) \frac{(c_j)^{-\gamma}}{(c_i)^{-\gamma}}$$

implying that $\tilde{\eta} = -(\delta + \gamma g_c)$ and $\tilde{e}_j = (c_j)^\gamma$.

Under subjective beliefs and a stochastic discount factor of the form:

$$\hat{s}_{ij} = \exp(-\delta - \gamma g_c) \frac{(c_j)^{-\gamma}}{(c_i)^{-\gamma}},$$

we may recover subjective probabilities using formula (7). This special case is featured in [Ross \(2015\)](#), but here, except for a deterministic trend, consumption is stationary. Once the consumption process is exposed to permanent shocks, the stochastic discount factor inherits a martingale component that reflects this stochastic contribution under the subjective Markov evolution.

Example 2.2. *Again let $C_t = \exp(g_c t)(c \cdot X_t)$ be a trend-stationary consumption process where c is*

an $n \times 1$ vector that represents consumption in individual states of the world. The (representative) investor is now endowed with recursive preferences of [Kreps and Porteus \(1978\)](#) and [Epstein and Zin \(1989\)](#). We consider a special case of unitary elasticity of substitution and initially impose rational expectations. The continuation value for these preferences satisfies the recursion

$$V_t = [1 - \exp(-\delta)] \log C_t + \frac{\exp(-\delta)}{1 - \gamma} \log E_t[\exp((1 - \gamma)V_{t+1})], \quad (9)$$

where γ is a risk aversion coefficient and δ is a subjective rate of discount. For this example, $V_t = V(t, X_t = u^i) = v_i + g_c t$ where v_i is the continuation value for state $X_t = u^i$ net of a time trend. Let v be the vector with entry i given by v_i and $\exp[(1 - \gamma)v]$ be the vector with entry i given by $\exp[(1 - \gamma)v_i]$. The (translated) continuation values satisfy the fixed-point equation:

$$v_i = [1 - \exp(-\delta)] \log c_i + \frac{\exp(-\delta)}{1 - \gamma} \log [\mathbf{P}_i \exp[(1 - \gamma)v] + \exp(-\delta)g_c] \quad (10)$$

where \mathbf{P}_i is the i -th row of the transition matrix \mathbf{P} . This equation gives the current-period continuation values as a function of the current-period consumption and the discounted risk-adjusted future continuation values. We are led to a fixed-point equation because of our interest in an infinite-horizon solution. Given the solution v of this equation, denote $v^* = \exp[(1 - \gamma)v]$.

The implied stochastic discounting is captured by the following equivalent depictions:

$$s_{ij} = \exp[-(\delta + g_c)] \left(\frac{c_i}{c_j} \right) \left(\frac{v_j^*}{\mathbf{P}_i v^*} \right),$$

or, compounding over time,

$$S_t = \exp[-(\delta + g_c)t] \left(\frac{c \cdot X_0}{c \cdot X_t} \right) \left(\frac{H_t^*}{H_0^*} \right) \quad (11)$$

where

$$\frac{H_{t+1}^*}{H_t^*} = \frac{X_{t+1} \cdot v^*}{X_t \cdot (\mathbf{P}v^*)}.$$

The process H^* is a martingale. Perron–Frobenius Theory applied to \mathbf{P} implies that $\mathbf{P}v^* = v^*$ if, and only if, v^* has constant entries. As long as the solution v of equation (10) satisfies $v_i \neq v_j$ for some pair (i, j) , we conclude that $\mathbf{P}v^* \neq v^*$.

For this example

$$q_{ij} = p_{ij} \exp[-(\delta + g_c)] \left(\frac{c_i}{c_j} \right) \left(\frac{v_j^*}{\mathbf{P}_i v^*} \right).$$

Solving

$$\mathbf{Q}\hat{e} = \exp(\hat{\eta})\hat{e}$$

for a vector e with positive entries yields

$$\begin{aligned}\widehat{e}_j &= c_j, \quad j = 1, \dots, n \\ \widehat{\eta} &= -(\delta + g_c).\end{aligned}$$

This Perron–Frobenius solution (7) recovers the transition matrix $\widehat{\mathbf{P}}$ given by

$$\widehat{p}_{ij} = p_{ij} \left(\frac{v_j^*}{\mathbf{P}_i v^*} \right).$$

The recovered transition matrix $\widehat{\mathbf{P}}$ absorbs the risk adjustment that arises from fluctuations in the continuation value v . In particular, when $\gamma > 1$, transition probabilities \widehat{p}_{ij} are overweighted for low continuation value states v_j next period. When $\gamma = 1$, the two transition matrices coincide because v^* is necessarily constant across states.⁵

Consider now the case of subjective beliefs. Suppose that an analyst mistakenly assumes $\gamma = 1$ even though it is not. Then the martingale component in the actual stochastic discount factor is absorbed into the probability distribution the analysts attributes to the subjective beliefs. Alternatively, suppose that the analyst correctly recognizes that beliefs are subjective. Then recursion (10) holds with the subjective transition matrix $\widehat{\mathbf{P}}$ replacing \mathbf{P} . Any attempt to recover probabilities would have to take account of the impact of subjective beliefs on the value function and hence the stochastic discount factor construction. This impact is in addition to equation (4) that links Arrow prices to probabilities and to state-dependent discounting.

3 General framework

We now introduce a framework which encompasses a large class of relevant asset pricing models. Consistent intertemporal pricing together with a Markovian property lead us to use a class of stochastic processes called multiplicative functionals. These processes are built from the underlying Markov process in a special way and will be used to model stochastic discount factors. Alternative structural economic models will imply further restrictions on stochastic discount factors.

We start with a probability space $\{\Omega, \mathcal{F}, P\}$ and a set of indices T (either the non-negative integers or the non-negative reals). On this probability space, there is an n -dimensional, stationary Markov process $X = \{X_t : t \in T\}$ and a k -dimensional process W with increments that are jointly stationary with X and initialized at $W_0 = 0$. Although we are interested as before in the transition probabilities of the process X we use the process W to model the dynamics of X and provide a source for aggregate risks. The increments to W represent shocks to the economic dynamics and could be independently distributed over time with mean zero. We start with the discrete-time case,

⁵In the limiting case as $\delta \rightarrow 0$, the continuation value v_j converges to a constant independent of the state j due to the trend stationarity specification of the consumption process, and the risk adjustment embedded in the potential fluctuations of the continuation values becomes immaterial.

postponing the continuous-time case until Section 3.6. We suppose that

$$X_{t+1} = \phi_x(X_t, \Delta W_{t+1})$$

for a known function ϕ_x where $\Delta W_{t+1} \doteq W_{t+1} - W_t$. Furthermore we assume:

Assumption 3.1. *The process X is ergodic under P and the distribution of ΔW_{t+1} conditioned on X_t is time-invariant and independent of the past realizations of ΔW_s , $s \leq t$ conditioned on X_t .*

Write $\mathfrak{F} = \{\mathfrak{F}_t : t \in T\}$ for the filtration generated by histories of W and the initial condition X_0 .

3.1 Information

In what follows we assume that X_t is observable at date t but that the shock vector ΔW_{t+1} is not directly observable at date $t + 1$. Many of the results in this paper can be fully understood considering only the case where the shock vector ΔW_{t+1} can be inferred from (X_t, X_{t+1}) . In some examples that we provide later, however, the vector ΔW_{t+1} cannot be inferred from (X_t, X_{t+1}) . These are economic models for which there are more sources of uncertainty ΔW_{t+1} pertinent to investors than there are relevant state variables. In this case we consider a stationary increment process

$$Y_{t+1} - Y_t = \phi_y(X_t, \Delta W_{t+1})$$

that together with (X_{t+1}, X_t) reveals ΔW_{t+1} . Thus given X_t and knowledge of the functions ϕ_x and ϕ_y , the shock vector ΔW_{t+1} can be inferred from X_{t+1} and $Y_{t+1} - Y_t$. The observed histories of the joint process $Z \doteq (X, Y)$ thus generate the same filtration \mathfrak{F} as the histories of the shocks ΔW and the initial condition X_0 . Our construction implies that Z is a Markov process with a triangular structure because the distribution of $(X_{t+1}, Y_{t+1} - Y_t)$ conditioned on \mathcal{F}_t depends only on X_t .

3.2 Growth, discounting and martingales.

We introduce a valuable collection of scalar processes M that can be constructed from Z . The evolution of $\log M$ is restricted to have Markov increments of the form:

Restriction 3.2. *M satisfies*

$$\log M_{t+1} - \log M_t = \kappa(X_t, \Delta W_{t+1}).$$

Given our invertibility restriction, we may write:

$$\log M_{t+1} - \log M_t = \kappa^*(X_t, X_{t+1}, \Delta Y_{t+1})$$

Processes satisfying Restriction 3.2 are restricted versions of what we call *multiplicative functionals* of the process Z . (See Appendix A for the formal definition of a multiplicative functional.)

In what follows we will refer to the processes satisfying Restriction 3.2 simply as multiplicative functionals, because all of the results can be extended to this larger class of processes.

The process M is strictly positive. For two such functionals M^1 and M^2 , the product $M^1 M^2$ and the reciprocal $1/M^1$ are also strictly positive multiplicative functionals. Examples of such functionals are the exponential of linear combinations of the components of Y .

In light of Assumption 3.1, the logarithm of a multiplicative functional $\log M$ has stationary increments, and thus M itself can display geometric growth or decay along stochastic trajectories. The process M also could be a martingale whose expectation is invariant across alternative forecasting horizons, and in this sense does not grow or decay over time. We use multiplicative functionals to construct stochastic discount factors, stochastic growth factors and positive martingales that represent alternative probability measures.

3.3 An example

In the following example, we show how multiplicative functionals relate to the Markov chain framework analyzed in Section 2. This example consists of a Markov switching model that has state dependence in the conditional mean and in the exposure to normally distributed shocks. To include this richer collection of models, we allow our multiplicative functional to depend on a shock vector ΔW that is the normally distributed shock vector and hence not fully revealed by the evolution of the state X . Nevertheless, X_t still serves as the relevant state vector at date t .

Example 3.3. *Let X be a discrete-time, n -state Markov chain, and*

$$\Delta W_{t+1} = \begin{bmatrix} X_{t+1} - E(X_{t+1} | X_t) \\ \Delta \widehat{W}_{t+1} \end{bmatrix}$$

where $\Delta \widehat{W}_{t+1}$ is a k -dimensional standard normally distributed random vector that is independent of \mathfrak{F}_t and X_{t+1} .

The first block of the shock vector, $X_{t+1} - E(X_{t+1} | X_t)$, is by construction revealed by the observed realizations of the Markov chain X . In addition, we construct the vector process Y whose j -th coordinate evolves as

$$Y_{j,t+1} - Y_{j,t} = X_t \cdot \left[\bar{\mu}_j + \bar{\sigma}_j \left(\Delta \widehat{W}_{t+1} \right) \right]$$

where $\bar{\mu}_j$ is a vector of length n , and $\bar{\sigma}_j$ is an $n \times k$ matrix. The matrices $\bar{\sigma}_j$ are restricted to insure that $\Delta \widehat{W}_{t+1}$ can be computed from $Y_{t+1} - Y_t$ and X_t . Thus $Z = (X, Y)$ reveals W . We allow for date $t+1$ Arrow contracts to be written as functions of Y_{t+1} as well as X_{t+1} along with the relevant date t information. In this environment, we represent the evolution of a multiplicative functional M as:

$$\log M_{t+1} - \log M_t = X_t \cdot \left[\bar{\beta} + \bar{\alpha} (\Delta W_{t+1}) \right]$$

where $\bar{\beta}$ is a vector of length n and $\bar{\alpha}$ is an $n \times (n+k)$ matrix.

3.4 Stochastic discount factors

A *stochastic discount factor process* S is a positive multiplicative functional with $S_0 = 1$ and finite first moments (conditioned on X_0) such that the date τ price of any bounded \mathfrak{F}_t -measurable claim Φ_t for $t > \tau$ is:

$$\Pi_{\tau,t}(\Phi_t) \doteq E \left[\frac{S_t}{S_\tau} \Phi_t \mid \mathfrak{F}_\tau \right]. \quad (12)$$

As a consequence, for a bounded claim $f(X_t)$ that depends only on the current Markov state, the time-zero price is

$$[\mathbb{Q}_t f](x) \doteq E [S_t f(X_t) \mid X_0 = x]. \quad (13)$$

We view \mathbb{Q}_t as the pricing operator for payoff horizon t . By construction,

$$\Pi_{\tau,t}[f(X_t)] = [\mathbb{Q}_{t-\tau} f](X_t).$$

The operator \mathbb{Q}_t is well defined at least for bounded functions of the Markov state, but often for a larger class of functions, depending on the tail behavior of the stochastic discount factor S_t .

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 and thus it suffices to study the one-period operator.⁶

3.5 Multiplicative martingales and probability measures

Alternative probability measures equivalent to P are built using strictly positive martingales. Given an \mathfrak{F} -martingale H that is strictly positive with $E(H_0) = 1$, define a probability P^H such that if $A \in \mathcal{F}_\tau$ for some $\tau \geq 0$,

$$P^H(A) = E(1_A H_\tau). \quad (14)$$

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

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

Now suppose that H is a multiplicative martingale, a multiplicative functional that is also a martingale with respect to the filtration \mathfrak{F} , modeled as

$$\log H_{t+1} - \log H_t = h(X_t, \Delta W_{t+1}).$$

For the martingale restriction to be satisfied, impose

$$E(\exp[h(X_t, \Delta W_{t+1})] \mid X_t = x) = 1.$$

⁶There is a different stochastic discount factor process that we could use for much of our analysis. Let $\bar{\mathfrak{F}}$ denote the (closed) filtration generated by X . Compute $\bar{S}_t = E[S_t \mid \bar{\mathfrak{F}}_t]$. Then \bar{S}_t is a stochastic discount factor process pertinent for pricing claims that depend on the history of X . It is also a multiplicative functional constructed from X .

Under the implied change of measure, the probability distribution for $(X_{t+1}, \Delta W_{t+1})$ conditioned on \mathfrak{F}_t continues to depend only on X_t .

While we normalize $S_0 = 1$, we will not do the same for H_0 . For some of our subsequent discussion, we use H_0 to alter the initial distribution of X_0 in a convenient way. Thus we allow H_0 to depend on X_0 , but we restrict it to have expectation equal to unity.

An examination of (13) makes it evident that by using $S^H = S \frac{H}{H_0}$ as the stochastic discount factor and P^H as the corresponding probability measure, we will represent the same family of pricing operators $\{\mathbb{Q}_t : t \geq 0\}$ over bounded functions of x . This flexibility in how we represent pricing extends what we observed in (4) for the finite-state economies.

3.6 Continuous-time diffusions

We impose an analogous structure when X is a continuous-time diffusion. The process W is now an underlying n -dimensional Brownian motion and we suppose X_0 is independent of W and let \mathfrak{F} be the (completed) filtration associated with the Brownian motion augmented to include date-zero information revealed by $Z_0 = (X_0, Y_0)$. Then X , Y and $\log M$ processes evolve according to:⁷

$$\begin{aligned} dX_t &= \mu_x(X_t)dt + \sigma_x(X_t)dW_t \\ dY_t &= \mu_y(X_t)dt + \sigma_y(X_t)dW_t \\ d \log M_t &= \beta(X_t)dt + \alpha(X_t) \cdot dW_t. \end{aligned} \tag{15}$$

Notice that the conditional distribution of $(X_{t+\tau}, Y_{t+\tau} - Y_t)$ conditioned on \mathfrak{F}_t depends only on X_t analogous to the assumption that we imposed in the discrete-time specification. In addition we suppose that

$$\sigma = \begin{bmatrix} \sigma_x \\ \sigma_y \end{bmatrix}$$

is nonsingular, implying that the Brownian motion history is revealed by the $Z := (X, Y)$ history and \mathfrak{F} is also the filtration associated with the diffusion Z . For the continuous-time specification (15), the drift term of a multiplicative martingale H satisfies⁸

$$\beta(x) = -\frac{1}{2}\alpha(x) \cdot \alpha(x).$$

The definition of a stochastic discount factor and of the family of operators \mathbb{Q}_t when t is continuous is the direct counterpart to the constructs used in Section 3.4 for the discrete-time models. 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 $t = 0$.

⁷While this Brownian information specification abstracts from jumps, these can be included without changing the implications of the analysis, see Hansen and Scheinkman (2009).

⁸This restriction implies that H is a local martingale and additional restrictions may be required to ensure that H is a martingale.

Under the change of probability induced by H , W_t has a drift $\alpha(X_t)$ and is no longer a martingale. As in our discrete-time specification, under the change of measure, Z will remain a Markov process and the triangular nature of $Z = (X, Y)$ will be preserved. Furthermore, we can represent the same operator $\{\mathbb{Q}_t : t \geq 0\}$ over bounded functions of x using $S^H = S \frac{H}{H_0}$ as the stochastic discount factor and P^H as the corresponding probability measure.

4 What is recovered

We now review and extend previous results on long-term valuation. In so doing we will exploit the triangular nature of the Markov process and feature the state vector X_t . Later we will explore what happens when we extend the state vector to include Y_t in our analysis of long-term pricing.

4.1 Perron–Frobenius approach to valuation

Consider a solution to the the following *Perron–Frobenius* problem:

Problem 4.1 (Perron–Frobenius). *Find a scalar $\hat{\eta}$ and a function $\hat{e} > 0$ such that for every $t \in T$,*

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

A solution to this problem will necessarily satisfy the conditional moment restriction:

$$E[S_t \hat{e}(X_t) \mid \mathfrak{F}_\tau] = \exp[(t - \tau)\hat{\eta}] S_\tau \hat{e}(X_\tau) \tag{16}$$

for $t \geq \tau$. Since \hat{e} is an eigenfunction, it is only well-defined up to a positive scale factor. When we make reference to a unique solution to this problem, we mean that \hat{e} is unique up to scale.

When the state space is finite as in Section 2, functions of x can be represented as vectors in \mathbb{R}^n , and the operator \mathbb{Q}_1 can be represented as a matrix \mathbf{Q} . In this case, the existence and uniqueness of a solution to Problem 4.1 is well understood. 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. If the Perron–Frobenius problem has a solution, we follow Hansen and Scheinkman (2009), and define a process \hat{H} that satisfies:

$$\frac{\hat{H}_t}{\hat{H}_0} = \exp(-\hat{\eta}t) S_t \frac{\hat{e}(X_t)}{\hat{e}(X_0)}. \tag{17}$$

The process \hat{H} is a positive \mathfrak{F} -martingale under the probability measure P , since, for $t \geq \tau$,

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

where in the second equality we used equation (16).

The process \widehat{H} inherits much of the mathematical structure of the original stochastic discount factor process S and is itself a multiplicative martingale. For instance, if S has the form given by Restriction 3.2, then:

$$\begin{aligned}\log \widehat{H}_{t+1} - \log \widehat{H}_t &= \kappa(X_t, \Delta W_{t+1}) + \log \widehat{e}(X_{t+1}) - \log \widehat{e}(X_t) - \widehat{\eta} \\ &\doteq \widehat{h}(X_t, \Delta W_{t+1})\end{aligned}$$

where we have used the fact that $X_{t+1} = \phi_x(X_t, \Delta W_{t+1})$.

When we change measures using the martingale \widehat{H} , to be consistent with the family of pricing operators $\{\mathbb{Q}_t : t \geq 0\}$ the associated stochastic discount factor must be:

$$\widehat{S}_t = S_t \frac{\widehat{H}_0}{\widehat{H}_t} = \exp(\widehat{\eta}) \frac{\widehat{e}(X_0)}{\widehat{e}(X_t)}.$$

Under this change of measure, the discounting of a payoff at time t to time 0 is independent of the path of the state between 0 and t .

As we change probability measures, stationarity and ergodicity of X will not necessarily continue to hold. But checking for this stability under the probability P^H induced by the martingale H will be featured in our analysis. Thus in establishing a uniqueness result, we impose the following restriction on the stochastic evolution of X under the probability distribution P^H .

Restriction 4.2. *The process X is stationary and ergodic under P^H .*

Stationarity and ergodicity requires the choice of an appropriate $H_0 = h(X_0)$ that induces a stationary distribution under P^H for the Markov process X .⁹

If X satisfies Restriction 4.2 then it satisfies a Strong Law of Large Numbers.¹⁰ In the discrete-time case, if a function ψ has finite expected value, then

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N \psi(X_t) = E^H \psi(X_0)$$

almost surely. The process X also obeys another version of Law of Large Numbers that considers convergence in means:¹¹

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

⁹Given H_t/H_0 , the random variable $H_0 = h(X_0)$ must satisfy the equation:

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

for any bounded (Borel measurable) ψ and any $t \in T$.

¹⁰E.g., Breiman (1982), Corollary 6.23 on page 115.

¹¹Ibid., Corollary 6.25 on page 117.

As a consequence of both versions of the Law of Large Numbers, times-series averages of conditional expectations also converge:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N E^H [\psi(X_t) | X_0] = E^H \psi(X_0)$$

almost surely. Corresponding results hold in continuous time.

We now show that the Perron–Frobenius Problem 4.1 has a unique solution under which X is stationary and ergodic under the probability measure implied by \hat{H} .

Proposition 4.3. *There is at most one solution $(\hat{e}, \hat{\eta})$ to Problem 4.1 such that X is stationary and ergodic under the probability measure $P^{\hat{H}}$ induced by the multiplicative martingale \hat{H} given by (17).*

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 B.¹² In what follows we will use \hat{P} and \hat{E} instead of the more cumbersome $P^{\hat{H}}$ and $E^{\hat{H}}$.

4.2 An illustration of what is recovered

In the previous discussion, we described two issues arising in the recovery procedure. First, the positive candidate solution for $\hat{e}(x)$ may not be unique. Our Restriction 4.2 allows us to pick the single solution that preserves stationarity and ergodicity. Second, even this unique choice may not uncover the true probability distribution if there is a martingale component in the stochastic discount factor. The following example shows that in a simplified version of a stochastic volatility model one always recovers an incorrect probability distribution.

Example 4.4. *Consider a stochastic discount factor model with state-dependent risk prices.*

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

where $\bar{\beta} < 0$ and X has the square root dynamics

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

Guess a solution for a positive eigenfunction:

$$\hat{e}(x) = \exp(vx).$$

Since $\{\exp(-\hat{\eta}t) S_t \hat{e}(X_t) : t \geq 0\}$ is a martingale:

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

¹²Hansen and Scheinkman (2009) and Hansen (2012) used an implication of the SLLN in their analysis.

In particular, the coefficient on x should satisfy

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

There are two solutions: $v = 0$ and

$$v = \frac{2\kappa - 2\bar{\alpha}\bar{\sigma}}{(\bar{\sigma})^2} \quad (18)$$

In this example, the risk neutral dynamics for X corresponds to the solution $v = 0$ 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 Restriction 4.2 picks the risk neutral dynamics, which is distinct from original dynamics for X . Suppose instead that $\kappa_n < 0$, which occurs when $\kappa < \bar{\sigma}\bar{\alpha}$. In this case Restriction 4.2 selects v given by (18), 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).

Multiplicity of solutions to the Perron–Frobenius problem is prevalent in models with continuous states. In confronting this multiplicity, Hansen and Scheinkman (2009) show that the eigenvalue $\hat{\eta}$ that leads to a stochastically stable probability measure \hat{P} gives a lower bound to the set of eigenvalues associated with strictly positive eigenfunctions. For a univariate continuous-time Brownian motion setup, Walden (2014) and Park (2014) construct positive solutions e for every candidate eigenvalue $\eta > \hat{\eta}$. However, none of these solution pairs (e, η) leads to a probability measure that satisfies Restriction 4.2.

5 Long-term pricing

Risk neutral probabilities absorb short-term risk adjustments. In contrast, we now show that the probability measure identified by the application of Perron–Frobenius theory absorbs risk adjustments over long horizons. We call this latter probability measure the long-term risk neutral measure.

Perron–Frobenius Theory features an eigenvalue $\hat{\eta}$ and an associated eigenfunction \hat{e} which determine the limiting behavior of securities with payoffs far in the future. We exploit this domination to study long-term risk-return tradeoffs building on the work of Hansen et al. (2008) and Hansen and Scheinkman (2009) and long-term holding period returns building on the work of Alvarez and

Jermann (2005). We will show that under the \widehat{P} probability measure, risk-premia on long term cash flows that grow stochastically are zero; but not under the P measure. We also show that the holding period return on long-term bonds is the increment in the stochastic discount factor under the \widehat{P} measure.

For some of the results in this section, we impose the following refinement of ergodicity.

Restriction 5.1. *The Markov process X is aperiodic, irreducible and positive recurrent under the measure \widehat{P} .*

We refer to this restriction as *stochastic stability*, and it implies that

$$\lim_{t \rightarrow \infty} \widehat{E}[f(X_t) | X_0 = x] = \widehat{E}[f(X_0)]$$

almost surely provided that $\widehat{E}[f(X_0)] < \infty$.¹³ In this formula, we use the notation \widehat{E} to denote expectations computed with the probability \widehat{P} implied by \widehat{H} .¹⁴

5.1 Long-term yields

We first show that the characterization of the eigenvalue $\widehat{\eta}$ in Section 2.2 extends to this more general framework. Consider,

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

for some positive payoff $\psi(X_t)$ expressed as a function of the Markov state. For instance, to price pure discount bonds we should set $\psi(x) \equiv 1$.

Consider the implied yield to this investment under the measure \widehat{P} :

$$\widehat{y}_t[\psi(X)](x) \doteq \frac{1}{t} \log \widehat{E}[\psi(X_t) | X_0 = x] - \frac{1}{t} \log [\mathbb{Q}_t \psi](x)$$

Taking the limit as $t \rightarrow \infty$:

$$\lim_{t \rightarrow \infty} \widehat{y}_t[\psi(X)](x) = -\widehat{\eta} + \lim_{t \rightarrow \infty} \frac{1}{t} \log \widehat{E}[\psi(X_t) | X_0 = x] - \lim_{t \rightarrow \infty} \frac{1}{t} \log \widehat{E} \left[\frac{\psi(X_t)}{\widehat{e}(X_t)} | X_0 = x \right]$$

This limit shows that $-\widehat{\eta}$ is the long term yield maturing in the distant future, provided that the last two terms vanish. These last two terms vanish under the stochastic stability Restriction 5.1 provided that¹⁵

$$\widehat{E}[\psi(X_0)] < \infty, \quad \widehat{E} \left[\frac{\psi(X_0)}{\widehat{e}(X_0)} \right] < \infty.$$

¹³ For discrete-time models, see Meyn and Tweedie (2009) Theorem 14.0.1 on page 334 for an even stronger conclusion. We prove the results in this section and appendices only for the discrete-time case. Analogous results for the continuous-time case would use propositions in Meyn and Tweedie (1993).

¹⁴Note that we use \widehat{P} and \widehat{E} instead of the more cumbersome $P^{\widehat{H}}$ and $E^{\widehat{H}}$.

¹⁵Since the logarithms of the conditional expectations are divided by t , $-\widehat{\eta}$ is the long-term yield under more general circumstances. See Appendix C for details.

When we use the original measure P instead of \widehat{P} , the finite horizon yields will differ. The limiting yield would still equal $-\widehat{\eta}$ under the original probability measure, provided that

$$E[\psi(X_t)] < \infty.$$

In summary, stationary cash flow risk does not alter the long-term yield since $-\widehat{\eta}$ is also the yield on a long-term discount bond. The limiting risk premium is zero under both probability measures because of the transient nature of the cash-flow risk. Next we introduce payoffs for which stochastic growth implies nonvanishing limiting risk premia.

5.2 Long-term risk-return tradeoff

Consider a positive cash flow process G that grows stochastically over time. We model such a cash flow as a multiplicative functional satisfying Restriction 3.2 with payoff G_t at time t . By design, the growth rate in logarithms fluctuates randomly depending on the Markov state and the shocks modeled as martingale increments. Given the multiplicative nature of G , the impact of growth compounds over time. While we expect the long-term growth rate of G to be positive, the overall exposure to shocks will increase with the payoff date t . Cash-flow risk will no longer be transient as in Section 5.1. For convenience we initialize $G_0 = 1$.

The yield on the cash flow G under the long-term risk neutral probability model $(\widehat{S}, \widehat{P})$ is

$$\widehat{y}_t[G](x) = \frac{1}{t} \log \frac{\widehat{E}[G_t | X_0 = x]}{\widehat{E}[\widehat{S}_t G_t | X_0 = x]} = \frac{1}{t} \log \frac{E\left[\frac{\widehat{H}_t}{\widehat{H}_0} G_t | X_0 = x\right]}{E[S_t G_t | X_0 = x]} \quad (19)$$

where

$$\frac{\widehat{H}_t}{\widehat{H}_0} = \exp(-\eta t) S_t \frac{\widehat{e}(X_t)}{\widehat{e}(X_0)},$$

and $\widehat{S}_t = S_t \frac{\widehat{H}_0}{\widehat{H}_t}$. Using this formula for \widehat{H} in equation (19), we obtain:

$$\begin{aligned} \widehat{y}_t[G](x) &= \frac{1}{t} \log E\left[\frac{\widehat{H}_t}{\widehat{H}_0} G_t | X_0 = x\right] - \frac{1}{t} \log E[S_t G_t | X_0 = x] \\ &= -\eta + \frac{1}{t} \log E\left[S_t G_t \frac{\widehat{e}(X_t)}{\widehat{e}(X_0)} | X_0 = x\right] - \frac{1}{t} \log E[S_t G_t | X_0 = x] \\ &= -\eta, \end{aligned}$$

provided that additional moment restrictions are imposed. See Appendix C for details. The limiting yield computed under \widehat{P} remains the same even after we have introduced stochastic growth in the payoff. In particular, the long-term risk premia on cash flows are zero under \widehat{P} even when the cash flows display stochastic growth.

This conclusion, however, is altered when we compute the yield under the original probability

measure. Now the horizon t yield is:

$$y_t[G](x) = \frac{1}{t} \log E[G_t | X_0 = x] - \frac{1}{t} \log E[S_t G_t | X_0 = x]$$

If the martingale components of S and G are correlated over long horizons, the limiting yield of the payoff G under P differs from $-\eta$.

When S and G have non-trivial martingale components, the expected rate of growth of the multiplicative functional SG does *not* typically equal the sum of the expected rate of growth of S plus the expected rate of growth of G . Hence the limiting yield of the payoff G under P differs from $-\eta$. While the long-term risk premia on stochastically growing cash flows are zero under \widehat{P} , these same long-term risk premia under P are often not degenerate. By construction, the probability measure associated with Perron–Frobenius Theory makes the long-term risk-return tradeoff vanish.

5.3 Forward measures and holding-period returns to long-term bonds

Holding-period returns on long-term bonds inform us about the solution to the Perron–Frobenius Problem (Problem 4.1) for the stochastic discount factor process. The long maturity limit of a holding period return $R_{t,t+1}^\tau$ from period t to period $t+1$ on a bond with maturity τ is (almost surely)

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

provided that the stochastic stability Restriction 5.1 is imposed and

$$\widehat{E} \left[\frac{1}{\widehat{e}(X_0)} \right] < \infty.$$

Since prices of discount bonds at alternative investment horizons are used to construct forward measures, this same computation allows us to characterize the limiting forward measure. As Hansen and Scheinkman (2014) argue, the limit of the forward probability measures defined in Section 2.1 above coincides with the measure recovered using Perron–Frobenius Theory. To see why, consider the forward measure at date t for a maturity τ . We represent this measure using the positive random variable

$$F_{t,t+\tau} = \frac{S_{t+\tau}/S_t}{E[S_{t+\tau}/S_t | X_t]}.$$

with conditional expectation one given the date- t Markov state X_t . The associated conditional expectations computed using this forward measure are formed by first multiplying by $F_{t,t+\tau}$ prior to computing the conditional expectations using the P measure. Thus $F_{t,t+\tau}$ determines the conditional density of the forward measure with respect to the original measure. The implied one-period transition between date t and date $t+1$ is given by

$$E(F_{t,t+\tau} | \mathcal{F}_{t+1}) = \left(\frac{S_{t+1}}{S_t} \right) \left(\frac{E[S_{t+\tau}/S_{t+1} | \mathcal{F}_{t+1}]}{E[S_{t+\tau}/S_t | \mathcal{F}_t]} \right) = \left(\frac{S_{t+1}}{S_t} \right) \frac{[\mathbb{Q}_{\tau-1} \mathbf{1}](X_{t+1})}{[\mathbb{Q}_\tau \mathbf{1}](X_t)}$$

which by the Law of Iterated Expectations has expectation equal to unity conditioned on date t information. Using our previous calculations, taking limits as the investment horizon τ becomes arbitrarily large, the limiting transition distribution is determined by the random variable:

$$\left(\frac{S_{t+1}}{S_t}\right) R_{t,t+1}^\infty = \frac{\widehat{H}_{t+1}}{\widehat{H}_t} \quad (20)$$

which reveals the martingale increment in the stochastic discount factor. This also shows that the limiting one-period transition constructed from the forward measure coincides with the Perron–Frobenius transition probability. [Qin and Linetsky \(2014a\)](#) characterize this limiting behavior under more general circumstances without relying on a Markov structure.

When the right-hand side of (20) is exactly one, the one-period stochastic discount factor is the inverse of the limiting holding-period return. This link was first noted by [Kazemi \(1992\)](#). More generally, it follows from this formula that

$$E \left[\log R_{t,t+1}^\infty \mid X_t = x \right] \leq E \left[\log S_t - \log S_{t+1} \mid X_t = x \right], \quad (21)$$

since Jensen’s Inequality informs us that

$$E \left[\log \widehat{H}_{t+1} - \log \widehat{H}_t \mid X_t = x \right] \leq 0.$$

[Bansal and Lehmann \(1997\)](#) featured the maximal growth portfolio, that is the portfolio of returns that attains the right-hand side of (21). When $\log \widehat{H}_{t+1} - \log \widehat{H}_t$ is identically zero, $R_{t,t+1}^\infty$ coincides with the return on the maximal growth portfolio. As [Bansal and Lehmann \(1994\)](#) noted, the [Kazemi \(1992\)](#) result omits permanent components to the stochastic discount factor process. Nevertheless the one-period stochastic discount factor can still be inferred from that maximal growth portfolio provided that an econometrician has a sufficiently rich set of data on returns. Following [Alvarez and Jermann \(2005\)](#), we will use formula (20) in Section 9.3 when we discuss empirical methods and evidence for assessing the magnitude of the martingale component to the stochastic discount factor process.

Now suppose that we change measures and perform the calculations under \widehat{P} using stochastic discount factor:

$$\widehat{S}_t = S_t \frac{\widehat{H}_0}{\widehat{H}_t}.$$

In this case

$$\widehat{E} \left[\log R_{t,t+1}^\infty \mid X_t = x \right] = \widehat{E} \left[\log \widehat{S}_t - \log \widehat{S}_{t+1} \mid X_t = x \right].$$

To interpret this finding, consider any one-period positive return $R_{t,t+1}$. Since

$$\widehat{E} \left[\left(\frac{\widehat{S}_{t+1}}{\widehat{S}_t} \right) R_{t,t+1} \mid \mathcal{F}_t \right] = 1,$$

applying Jensen's Inequality,

$$\widehat{E} [\log R_{t+1} \mid \mathcal{F}_t] \leq E \left[\log \widehat{S}_t - \log \widehat{S}_{t+1} \mid \mathcal{F}_t \right]$$

By construction the martingale component of the stochastic discount factor under the \widehat{P} probability measure is degenerate. As a consequence, the inverse of the holding-period return on a long-term bond $R_{t,t+1}^\infty$ coincides with $\widehat{S}_{t+1}/\widehat{S}_t$ as in [Kazemi \(1992\)](#)'s model.

6 A quantitative example

We now show that a well-known structural model of asset pricing proposed by [Bansal and Yaron \(2004\)](#) implies a prominent martingale component. The model features growth-rate predictability and stochastic volatility in the aggregate consumption process. We utilize a continuous-time Brownian information specification described in [Hansen et al. \(2007\)](#) that is calibrated to the consumption dynamics postulated in [Bansal and Yaron \(2004\)](#).

We compare the implications of using the probability measure associated with the Perron–Frobenius extraction with the original probability measure and with the risk neutral measure. In this example the Perron–Frobenius extraction gives a very similar probability measure to the risk neutral measure but substantially different from the original probability measure. More generally, our aim in this section is to show that the differences in probability measures could be substantial, rather than giving a definitive conclusion that they are. The latter conclusion would necessitate a confrontation with direct statistical evidence, an aspect that we discuss in [Section 9](#).

Assume the date- t bivariate state vector takes the form $X_t = (X_{1,t}, X_{2,t})'$. In this model, $X_{1,t}$ represents predictable components in the growth rate of the multiplicative functional, and $X_{2,t}$ captures the contribution of stochastic volatility. The dynamics of X specified in [\(15\)](#) have parameters $\mu(x)$ and $\sigma(x)$ given by

$$\mu(x) = \bar{\mu}(x - \iota) \quad \sigma(x) = \sqrt{x_2} \bar{\sigma} \tag{22}$$

where

$$\bar{\mu} = \begin{bmatrix} \bar{\mu}_{11} & \bar{\mu}_{12} \\ 0 & \bar{\mu}_{22} \end{bmatrix} \quad \bar{\sigma} = \begin{bmatrix} \bar{\sigma}_1 \\ \bar{\sigma}_2 \end{bmatrix}. \tag{23}$$

The parameters $\bar{\sigma}_1$ and $\bar{\sigma}_2$ are 1×3 row vectors. The vector ι is the vector of means of the state variables in a stationary distribution.

All multiplicative functionals M that we consider satisfy [Restriction 3.2](#) with parameters $\beta(x)$ and $\alpha(x)$ such that:

$$\beta(x) = \bar{\beta}_0 + \bar{\beta}_1 \cdot (x - \iota) \quad \alpha_x = \sqrt{x_2} \bar{\alpha}. \tag{24}$$

For instance, the aggregate consumption process C is a multiplicative functional parameterized by $(\beta_c(x), \alpha_c(x))$.

Appendix D provides details on the calculations that follow. We use three uncorrelated shocks in our parameterization. The *direct consumption* shock is the component of the Brownian motion W that is a direct innovation to the consumption process $\log C_t$. The *growth rate* shock is the Brownian component that serves as the innovation to the growth rate $X_{1,t}$, while the *volatility* shock is the innovation to the volatility process $X_{2,t}$.

We endow the representative investor with the recursive homothetic preferences featured in Example 2.2. We impose a unitary elasticity of substitution for convenience and thus use the continuous-time counterpart to (9). The continuous-time version of these preferences was developed in Duffie and Epstein (1992) and Schroder and Skiadas (1999).

The stochastic discount factor solves

$$d \log S_t = -\delta dt - d \log C_t + d \log H_t^* \quad (25)$$

where H^* is the continuous-time counterpart to the martingale from equation (11).¹⁶ The H^* component is constructed as follows. Let V be the forward-looking continuation value process for the homogeneous-of-degree-one utility aggregator, and construct $V^{1-\gamma}$ where γ is the risk aversion parameter. The logarithm of the continuation value for the assumed consumption process is an additively separable function of $\log C$ and X . The Brownian increment for the martingale \bar{S} in the stochastic discount factor evolution coincides with the Brownian increment for $V^{1-\gamma}$. The stochastic discount factor inherits the functional form (24) with parameters (β_s, α_s) derived in Appendix D. Since the consumption process C is modeled using a permanent shock, it also contains a martingale component, and thus H^* is not the martingale arising from the Perron–Frobenius problem.

For the Perron–Frobenius probability extraction, we find a solution $(\hat{e}, \hat{\eta})$ to the Perron–Frobenius Problem 4.1 such that

$$S_t = \exp(\hat{\eta}t) \frac{\hat{e}(X_0)}{\hat{e}(X_t)} \frac{\hat{H}_t}{\hat{H}_0} \doteq \hat{S}_t \frac{\hat{H}_t}{\hat{H}_0}$$

and \hat{H} implies a probability measure \hat{P} that satisfies Restriction 4.2.¹⁷

We show in the appendix that the martingale \hat{H} associated with \hat{P} takes the form

$$\frac{d\hat{H}_t}{\hat{H}_t} = \sqrt{X_{2t}} \hat{\alpha}_h \cdot dW_t$$

where $\hat{\alpha}_h$ is a vector that depends on the parameters of the model. This implies that we can write

¹⁶A rather different motivation for the martingale H^* comes from literature on robustness concerns and asset pricing. For instance, see Anderson et al. (2003). In this case H^* is an endogenously determined probability adjustment for potential model misspecification. Other models of ambiguity aversion based on max-min utility also induce a martingale component to the stochastic discount factor.

¹⁷The ergodicity and stationarity of X under the recovered measure \hat{P} , as well as the existence of a solution for the recursive utility stochastic discount factor, can always be checked for given parameters by a direct calculation. For instance, see the calculations in Borovička et al. (2014) for details.




Figure 1: Stationary densities for the state vector $X = (X_1, X_2)'$ under the correctly specified probability measure P (left panel) and the recovered probability measure \hat{P} (right panel). The dashed line in the right panel corresponds to the outermost contour for the distribution under the risk neutral probability measure \bar{P} . The parameterization of the model is $\bar{\beta}_{c,0} = 0.0015$, $\bar{\beta}_{c,1} = 1$, $\bar{\beta}_{c,2} = 0$, $\bar{\mu}_{11} = -0.021$, $\bar{\mu}_{12} = \bar{\mu}_{21} = 0$, $\bar{\mu}_{22} = -0.013$, $\bar{\alpha}_c = [0.0078 \ 0 \ 0]'$, $\bar{\sigma}_1 = [0 \ 0.00034 \ 0]$, $\bar{\sigma}_2 = [0 \ 0 \ -0.038]$, $\iota_1 = 0$, $\iota_2 = 1$, $\delta = 0.002$, $\gamma = 10$. Parameters are calibrated to monthly frequency.

the joint dynamics of the state vector $X = (X_1, X_2)'$ as

$$\begin{aligned} dX_{1t} &= [\hat{\mu}_{11}(X_{1t} - \hat{\iota}_1) + \hat{\mu}_{12}(X_{2t} - \hat{\iota}_2)] dt + \sqrt{X_{2t}} \bar{\sigma}_1 d\hat{W}_t \\ dX_{2t} &= \hat{\mu}_{22}(X_{2t} - \hat{\iota}_2) dt + \sqrt{X_{2t}} \bar{\sigma}_2 d\hat{W}_t \end{aligned}$$

which has the same structure as (22)–(23) with a new set of coefficients $\hat{\mu}_{ij}$ derived in the Appendix. The process \hat{W} is a Brownian motion under \hat{P} .

6.1 Forecasts with alternative probability measures

Structural macro-finance models allow us to construct predictions about the future distribution of macroeconomic quantities and financial cash flows. Probability measures extracted from asset market data can be used to forecast the future state of the macroeconomy and play a role in the discussion of public policy. In this section we compare forecasts under the alternative distributions.

Figure 1 plots the joint stationary distribution of the state vector X both under P (left panel) and under \hat{P} (right panel). While the distribution in the left panel is the true distribution that is consistent with time series evidence, the distribution in the right panel is the one expected to be observed by a hypothetical investor with beliefs given by \hat{P} . The distribution under \hat{P} exhibits a lower mean growth rate X_1 and a higher conditional volatility X_2 than the distribution under P . Moreover, the adverse states are correlated; low mean growth rate states are more likely to




Figure 2: Yields and holding period returns under the true and recovered probability measure. The top row plots the annualized yields on cash flows corresponding to the aggregate consumption process (left panel) and on bonds (right panel) with different maturities. The blue bands with solid lines correspond to the distribution under P , while the red bands with dashed lines to the distribution under \hat{P} . The lines represent quartiles of the distribution. The parameterization is as in Figure 1.

occur jointly with high volatility states. Bidder and Smith (2013) document similar distortions in a model with robustness concerns using the martingale \hat{S} from equation (25).

The black dashed line in the right panel of Figure 1 gives the outermost contour line for the joint density under the risk neutral dynamics. The distribution under the risk neutral probability is remarkably similar to the \hat{P} state probabilities and both are very different from the physical probabilities.

The similarity between the probability measures \hat{P} and \bar{P} emerges because the martingale component is known to dominate the behavior of the stochastic discount factor. See Hansen (2012) and Backus et al. (2014) for evidence to this effect. Consider the extreme case in which the stochastic discount factor implies that the Perron–Frobenius eigenfunction is constant and the associated martingale implies that under the probability measure \hat{P} the process X is ergodic. In this case $\hat{P} = \bar{P}$, the short-term interest rate is constant over time and the term structure is flat. While these term structure implications are not literally true for our parameterized recursive utility model, the martingale component is sufficiently dominant to imply that risk adjustments embedded \hat{P} and \bar{P} are very similar.

6.2 Asset pricing implications

The probability measures P and \hat{P} have substantially different implications for yields and holding period returns. In Section 5, we have shown that under \hat{P} , yields on risky cash flows in excess of the riskless benchmark converge to zero as the maturity of these cash flows increases.

The left panel in Figure 2 plots the yields (19) on a payoff that equals aggregate consumption at t as a function of t . The blue solid lines depict the quartiles of the yield distribution $y_t[C](x)$

corresponding to the stationary distribution of $X_0 = x$ computed under P . The red dashed lines represent the yields $\hat{y}_t[C](x)$ inferred by an investor who uses the recovered measure \hat{P} to compute expected payoffs; but the distribution of these yields is plotted under the correct probability measure P for the current state $X_0 = x$.

Because the consumption process is negatively correlated with the martingale \hat{H} , the yields computed under \hat{P} are downward biased relative to P . This is necessarily true by construction for long maturities as we showed in Section 5.2, but is also true throughout the term structure. By construction, the probability measure associated with the Perron–Frobenius martingale component eliminates risk adjustments associated with the cash flows from aggregate consumption at long horizons. For this example this long-term risk neutral measure accounts for virtually the whole risk premium (in excess of the maturity-matched bond) associated with the cash flows from aggregate consumption at all investment horizons.

7 Fundamental identification question

We now turn to an identification question. Suppose we observe Arrow prices for alternative realizations of the Markov state. Can we recover subjective beliefs? As we have already observed, asset prices as depicted by equation (12) depend simultaneously on stochastic discount factor processes and on investor beliefs. A stochastic discount factor process is thus only well defined for a given probability. If we happen to misspecify investor beliefs, this misspecification can be offset by altering the stochastic discount factor. This ability to offset a belief distortion poses a fundamental challenge to the identification of subjective beliefs. In this section we formalize this identification problem, and we consider potential restrictions on the stochastic discount factors that can solve this challenge.

Definition 7.1. *The pair (S, P) explains asset prices if equation (12) gives the date zero price of any bounded, \mathfrak{F}_t measurable claim Φ_t payable at any time $t \in T$.*

Consider now a multiplicative martingale H satisfying Restriction 3.2 and the associated probability measure P^H defined through (14). Similarly let S be a multiplicative functional satisfying Restriction 3.2 and initialized at $S_0 = 1$. We define:

$$S^H = S \frac{H_0}{H}. \quad (26)$$

The following proposition is immediate:

Proposition 7.2. *Suppose H is a martingale satisfying Restriction 3.2 with $E(H_0) = 1$ and S satisfies Restriction 3.2 with $S_0 = 1$. If the pair (S, P) explains asset prices then the pair (S^H, P^H) also explains asset prices. Moreover, S^H satisfies Restriction 3.2 and $S_0^H = 1$.*

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 typically must

change the stochastic discount factor to represent the same asset prices. Let H^1 and H^2 be two positive martingales that satisfy Restriction 3.2 with $E(H_0^1) = E(H_0^2) = 1$. Construct the corresponding stochastic discount factors using formula (26). Then we cannot distinguish the potential subjective probabilities implied by H^1 from those implied by H^2 from Arrow prices alone. There is a pervasive identification problem. To achieve identification of investor beliefs, 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 $\Pi_{\tau,t}$ for $\tau \leq t \in T$.

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 allow us to identify S .¹⁸

Alternatively, we may restrict the stochastic discount factor process further. For instance, risk-neutral pricing restricts the stochastic discount factor to be predetermined or locally predictable. Thus for a discrete-time specification:

$$\log S_{t+1} - \log S_t = \log[\bar{q}(X_t)]$$

where $\bar{q}(X_t)$ is the price of one-period discount bond. When this restriction is used, typically there is no claim that the resulting probability distribution is the same as that used by investors.

A different restriction imposes a special structure on S :

Restriction 7.3. *Let*

$$\tilde{S}_t = \exp(-\delta t) \frac{m(X_t)}{m(X_0)}$$

for some positive function m and some real number δ .

Ross (2015) proved an identification result under Restriction 7.3 when the dynamics of X are driven by a finite-state Markov chain as in Section 2. A strengthening of Restriction 7.3 is sufficient to guarantee that Arrow prices identify the stochastic discount factor and a probability distribution associated with that discount factor in the more general framework introduced in Section 3.

Proposition 7.4. *Suppose (S, P) explains asset prices. Let H be a positive multiplicative martingale such that (S^H, P^H) also explains asset prices and X is stationary and ergodic under P^H . If S^H also satisfies Restriction 7.3, then H is uniquely determined.*

Thus if \tilde{S} satisfies Restriction 7.3, it is the unique S^H identified by this proposition, and P^H identifies the subjective beliefs \tilde{P} of investors. The proof of this result follows directly from our previous analysis. Let

$$\hat{e} = \frac{1}{m}$$

¹⁸See 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.

and $\hat{\eta} = -\delta$. This choice of $(\hat{e}, \hat{\eta})$ solves the Perron–Frobenius Problem 4.1. The resulting martingale \hat{H} leaves the process X as stationary and ergodic. The conclusion follows directly from Assumption 3.1 and the uniqueness of the solution to Problem 4.1 established in Proposition 4.3.¹⁹

Recall from Section 4 that in general the \hat{P} probability measure absorbs the long-term risk risk adjustment. An immediate implication of Proposition 7.4 is that Restriction 7.3 makes the long-term risk-return tradeoff degenerate. Equivalently, the subjective discount factor process has only a degenerate martingale component since the candidate probability measure for depicting subjective beliefs is $\tilde{P} = \hat{P}$.

8 Additional state vector

Proposition 4.3 shows how to identify a martingale associated with the Perron–Frobenius problem. As we illustrated through examples, this martingale is, in general, non-trivial. As a consequence the “recovered” probability measure differs from the subjective probability measure. Perhaps the problem is that we limited the choice of eigenfunctions too much, by assuming they depend only on the state vector X_t , and not on Y_t . As we show, relaxing this restriction on the eigenfunction could allow for the subjective probability to correspond to one solution to the eigenfunction problem, but we may lose identification even if we impose ergodicity on X and stationary increments on Y . We also examine two ways by which one could approximate a process with stationary increments using highly persistent stationary processes. We argue that, in general, the lack of identification of the limit (stationary increments) process makes it likely that the highly persistent approximations have many *approximate* solutions to the Perron–Frobenius problem. This phenomenon makes the practical construction of solutions to the Perron–Frobenius problem challenging with highly persistent processes. On a more positive note, we show that when we specify *a priori* a multiplicative process that has the same martingale component as the stochastic discount factor, we may recover subjective beliefs from Arrow prices.

8.1 Perron–Frobenius revisited

We illustrate the consequences of enlarging the state space in discrete time. Similar considerations would hold in continuous time. Recall the joint Markov process:

$$\begin{aligned} X_{t+1} &= \phi_x(X_t, \Delta W_{t+1}) \\ Y_{t+1} - Y_t &= \phi_y(X_t, \Delta W_{t+1}). \end{aligned}$$

¹⁹In a continuous-time Brownian information setup, alternative conditions on the boundary behavior of the underlying Markov process may also uniquely identify a probability measure. These conditions 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 (2014) analyzes natural boundaries. While technical conditions like these may deliver a unique solution, they do not resolve the fundamental identification problem once we relax Restriction 7.3.

Up until now, we have only allowed for eigenfunctions that depend on X_t , but not Y_t . We now entertain eigenfunctions that depend on (X_t, Y_t) when solving the Perron–Frobenius problem. Thus we now solve:

$$E \left[\frac{S_{t+1}}{S_t} \varepsilon(X_{t+1}, Y_{t+1}) \mid X_t = x, Y_t = y \right] = \exp(\eta) \varepsilon(x, y)$$

Our previous solution remains a solution to this equation, but there may be many others. To see why, notice that

$$\exp(\zeta \cdot Y)$$

is a multiplicative functional for alternative choices of the vector ζ . For each choice ζ solve

$$E \left[\frac{S_{t+1}}{S_t} \exp[\zeta \cdot g(X_t, \Delta W_{t+1})] e_\zeta(X_{t+1}) \mid X_t = x \right] = \exp(\eta_\zeta) e_\zeta(x).$$

As a direct implication of this equation,

$$E \left[\frac{S_{t+1}}{S_t} \exp(\zeta \cdot Y_{t+1}) e_\zeta(X_{t+1}) \mid X_t = x, Y_t = y \right] = \exp(\eta_\zeta) \exp(\zeta \cdot y) e_\zeta(x).$$

and hence we set

$$\varepsilon(x, y) = \exp(\zeta \cdot y) e_\zeta(x).$$

For each ζ we could select a solution (should it exist) for which the implied probability measure leaves the state vector process X ergodic and thus Y has stationary and ergodic increments. But notice that we have constructed a family of solutions indexed by ζ . While we have not established existence, the approaches for doing so remain the same as $S \exp(\zeta \cdot Y)$ is itself a multiplicative functional. We have thus shown how augmenting the state vector to include the stationary increment process Y introduces more solutions to the eigenfunction problem raising the challenge of how to select a particular solution among this family of solutions.

This broader construction of an eigenfunction is of particular interest when the following restriction is satisfied.

Restriction 8.1. *Let*

$$\tilde{S}_t = \exp(-\tilde{\delta}t) \exp[\tilde{\zeta} \cdot (Y_t - Y_0)] \left[\frac{\tilde{m}(X_t)}{\tilde{m}(X_0)} \right]$$

for some choice of parameters $\tilde{\delta}$ and $\tilde{\zeta}$ and some positive function \tilde{m} of the state vector X_t .

For instance when $Y_t = \log C_t$, Assumption 8.1 requires a power utility function perhaps modified by a function of the state vector X_t , that could measure “habit persistence.” Given Restriction 8.1, one of the potential Perron–Frobenius eigenfunctions could be used to reveal the subjective probabilities, but we would still be left with the problem in identifying $\tilde{\zeta}$. Even with the specificity of Restriction 8.1, there is typically a continuum of solutions to the Perron–Frobenius problem. Changing ζ will typically alter the growth rates of the components of Y , thus prior restrictions or other information on the subjective rate of growth could be used in conjunction with Arrow prices in a productive way to achieve identification.

8.2 Stationary approximation

Highly persistent stationary processes are hard to distinguish from processes with stationary increments. While we find it unattractive to exclude a rich class of models that specify stochastic growth, a possible challenge to this view is to argue that model builders should focus instead on stationary Markov models that are highly persistent. In this subsection we suggest that such highly persistent Markov models will also present challenges in actual implementation.

We now study a sequence of models that reflect this challenge. For simplicity suppose that Y is a scalar process. The multivariate counterpart adds notation but not insight. Consider a sequence of stationary, ergodic Markov models indexed by j :

$$\begin{aligned} X_{t+1} &= \phi_x(X_t, \Delta W_{t+1}) \\ Y_{t+1} - Y_t &= \varphi_y(-\rho_j Y_t, X_t, \Delta W_{t+1}) \end{aligned}$$

where $\rho_j > 0$ converges to zero. In the limit

$$Y_{t+1} - Y_t = \varphi_y(0, X_t, \Delta W_{t+1}) = \phi_y(X_t, \Delta W_{t+1}).$$

For instance, φ_y could be affine in its first argument:

$$\varphi_y(-\rho_j Y_t, X_t, \Delta W_{t+1}) = -\rho_j Y_t + \phi_y(X_t, \Delta W_{t+1}).$$

We impose the following counterpart to Restriction 7.3.

Restriction 8.2. *Let*

$$\tilde{S}_t = \exp(-\tilde{\delta}t) \frac{\tilde{m}(X_t, Y_t)}{\tilde{m}(X_0, Y_0)}$$

for some positive function \tilde{m} and some real number $\tilde{\delta}$.

For each $j < \infty$, Proposition 7.4 guarantees that the Perron–Frobenius problem has a unique solution $\varepsilon(x, y) = 1/\tilde{m}(x, y)$ that preserves the ergodicity of (X, Y) . For the limit problem, Restriction 8.2 is not sufficient to guarantee that S is a multiplicative functional satisfying Restriction 3.2. Even after adding Restriction 8.1, as we have already argued, there is typically a continuum of solutions to the Perron–Frobenius problem for the limiting $\rho_\infty = 0$ problem. This makes approximation challenging as ρ_j becomes close to zero. The other solutions we described in subsection 8.1 above become “approximate” solutions when ρ_j is close to zero.

The formal sense of approximation will matter. Strictly speaking there is no long-term growth in Y along the sequence even though the limit process has growth. Consider the perspective of a researcher who uses standard statistical criteria of approximation that target transition dynamics. A highly persistent stochastic process will have growth episodes, and thus it is challenging to tell such a process from one for which ΔY has a positive unconditional mean using standard statistical criteria of approximation that target transition dynamics. To analyze this further would require a

more formal discussion of the approximation which is beyond the scope of this paper.

While this sequence excludes trend growth, the argument could be extended by including a trend line whereby

$$Y_t^* = Y_t + \nu t$$

and date t Arrow contracts are written terms of (Y_{t+1}^*, X_{t+1}) or $(Y_{t+1}^* - Y_t^*, X_{t+1})$. Then

$$Y_{t+1}^* - Y_t^* = Y_{t+1} - Y_t + \nu.$$

We suppose that our Markov specification of pricing is given in terms of (Y_{t+1}, X_t) . If ν is known to both the analyst and to investors inside the model, our previous analysis applies. Alternatively, ν could be revealed by the time inhomogeneity in the Arrow prices as a function of the current state. Either way, the computational challenges we describe previously would persist.

8.3 Structured recovery

We now explore a different generalization of Restriction 7.3 that also delivers a recovery of subjective probabilities.

Restriction 8.3. *The stochastic discount factor process \tilde{S} satisfies*

$$\tilde{S}_t = \exp(-\delta t) \left(\frac{Y_t^r}{Y_0^r} \right) \left[\frac{\tilde{m}(X_t)}{\tilde{m}(X_0)} \right]$$

for some pre-specified multiplicative functional Y^r satisfying Restriction 3.2.

The prespecified process Y^r captures the long-term risk return tradeoff, and the stochastic discount factor process \tilde{S} will have the same martingale component as Y^r . The reciprocal of a positive multiplicative functional is itself a multiplicative functional. We may restrict the extended Perron–Frobenius eigenfunction to be of the form

$$(y^r)^{-1} e(x).$$

Once we pre-specify Y^r , as in Proposition 7.4, \tilde{m} and δ can be inferred from the Arrow prices. Obtaining uniqueness in these circumstances amounts to postulating a process that contains the martingale component to consumption. Given knowledge of Y^r , we may again omit the more general vector y from the argument of the eigenfunction in the Perron–Frobenius calculation. The Arrow prices then reveal the probabilities.

Bansal and Lehmann (1997) and Hansen (2012) note that in many examples the multiplicative functional Y^r could come from a reference model with a direct interpretation. For instance, many models of habit persistence, both internal and external have a stochastic discount factor that can be depicted as

$$\frac{\tilde{S}_{t+1}}{\tilde{S}_t} = \exp(-\delta) \exp[-\gamma(\log C_{t+1} - \log C_t)] \frac{\tilde{m}(X_{t+1})}{\tilde{m}(X_t)}$$

where $1 - \gamma$ is used in the power specification for per period utility and the argument of this utility function depends on a private or social “habit stock”. For this example to be applicable we take

$$\log Y_t^r = -\gamma \log C_t$$

for a known value of γ . The function m of the Markov state captures the impact of the implied non-separability over time in preferences.²⁰ Alternatively, [Hansen and Scheinkman \(2014\)](#) suggest that m can be viewed as a transient misspecification of an underlying asset pricing model-based stochastic discount factor process Y^r .

9 Measuring the martingale component

In this section we consider methods for extracting evidence from asset market data about the magnitude of the martingale component in the stochastic discount factor process. We provide a unifying discussion of the literature and by so doing we add to the existing methods. This opens the door to new avenues for empirical research.

The existence of the martingale has two interpretations. Under one interpretation, the assumption of rational expectations allows us to assess the importance of long-term risk adjustments for stochastically growing cash flows. Under a second interpretation it measures the statistical discrepancy between subjective beliefs and the actual stochastic evolution of the state variables. For this second interpretation to be valid, we exclude a martingale component in the stochastic discount factor process of the subjective probability model.

9.1 Quantifying the martingale component

Statistical measures of discrepancy are often constructed using conveniently chosen convex functions. Consider functions ϕ_θ defined on the positive real numbers as:

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

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 $\phi_0(r) = r \log r$ and $\phi_{-1}(r) = -\log r$. 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

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

²⁰If the habit persistence model is fully parameterized, then given the probabilities one could solve for m as function of these parameters by computing the implied intertemporal marginal rates of substitution. This would lead to over-identifying restrictions.

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

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

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

with equality only when $\widehat{H}_{t+1}/\widehat{H}_t$ 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 $\widehat{H}_{t+1}/\widehat{H}_t$;
- ii) $\theta = 0$ in which case the implied measure of discrepancy is based on conditional relative entropy:

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

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

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

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

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

We consider two uses of these discrepancy measures.

9.2 Incomplete asset market data

Constructing the full range of Arrow prices for alternative states can be challenging, if not impossible, in practice. For this reason it is of interest to study empirically what can be learned about stochastic discount factor processes without using an explicit asset pricing model and without using a complete set of prices of Arrow securities.

We build on the approach initiated by Hansen and Jagannathan (1991) aimed at nonparametric characterizations of stochastic discount factors without using a full set of Arrow prices. While full identification is not possible, data from financial markets remain informative. We draw on the pedagogically useful characterization of Almeida and Garcia (2013) and Hansen (2014), but adapt it to misspecified beliefs along the lines suggested in Ghosh et al. (2012) and Hansen (2014). In so doing we build on a key insight of Kazemi (1992) and Alvarez and Jermann (2005).

We describe how to compute lower bounds for these discrepancy measures. We are led to the study of 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 for the holding period return on the long term bond

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

which implies that

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

As in [Alvarez and Jermann \(2005\)](#) and [Bakshi and Chabi-Yo \(2012\)](#), suppose that the limiting holding-period return $R_{t,t+1}^\infty$ can be well approximated. In this case, one could *test* directly for the absence of the martingale component by assessing whether

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

More generally, we express the pricing restrictions as

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

where \widehat{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:

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

Solve

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

subject to the linear constraints:

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

A strictly positive $\lambda_\theta(x)$ implies a nontrivial martingale component to the stochastic discount factor.²¹

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

$$\bar{\lambda}_\theta = \inf_{J_{t+1} > 0} E\phi_\theta(J_{t+1}) \tag{27}$$

²¹A simple computation shows that for the continuous-time diffusion case the discrepancies equal one-half the local variance of $\log \widehat{H}$, for all values of θ .

subject to:

$$\begin{aligned}
 & E[J_{t+1}] - 1 = 0 \\
 & E \left[J_{t+1} \left(\frac{1}{R_{t,t+1}^\infty} \right) (Y_{t+1})' - (Q_t)' \right] = 0.
 \end{aligned} \tag{28}$$

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 (27) 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 (28) is often quite tractable. See Hansen et al. (1995) and the literature on implementing generalized empirical likelihood methods for further discussion.²² Alvarez and Jermann (2005) apply these methods to produce the corresponding bounds for the martingale component of the stochastic discount factor process. While neither paper computes sharp bounds of the type characterized here, both Alvarez and Jermann (2005) and Bakshi and Chabi-Yo (2012) provide empirical evidence in support of a substantial martingale component to the stochastic discount factor process using a very similar approach.²³

9.3 Arrow prices reconsidered

In section 8.1, we provided an extended version of the Perron–Frobenius problem to allow for economies in which there were long-term risk adjustments in valuation. As a result of this extension, we described a parameterized family of solutions even when we observe the full array of Arrow prices. Recall that a multiplicative martingale is associated with each solution to the Perron–Frobenius problem. One possibility is to select the parameter for which the implied subjective beliefs are as close as possible to the actual data generating process using one of the measures of statistical discrepancy that we described in Section 9.1.

There are precursors to such an approach in the literature. For instance, Stutzer (1996) uses a $\theta = 0$ relative entropy measure of statistical discrepancy to study comparisons of risk neutral distributions and in effect Arrow prices to empirical counterparts formed from the actual data generation. Also, Chen et al. (2014) and Christensen (2014) suggest Perron–Frobenius methods for semiparametric identification of stochastic discount factor models under rational expectations, extending the approach of Hansen and Singleton (1982). Ait-Sahalia and Lo (2000) apply a formula

²²For 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 of 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.

²³Bakshi 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.

from [Breedon and Litzenberger \(1978\)](#) to infer risk neutral densities and actual densities using nonparametric statistical methods for low-dimensional Markov state environments.²⁴ Insights from this research may further help in crafting an approach to actual implementation that can be formally justified.

10 Conclusion

Perron–Frobenius theory applied to Arrow prices identifies a martingale component to the market-determined stochastic discount factor process. This martingale component defines a distorted probability measure that absorbs long-term risk adjustments, in the same spirit as the risk neutral probability measure absorbs one-period risk adjustments. We call this measure the long-term risk neutral measure.

One identifying assumption, featured by [Ross \(2015\)](#), assumes that this martingale component is identical to one under the subjective beliefs of the investors. In this case, the probability measure recovered by the Perron–Frobenius theory coincides with the subjective probability measure. If, however, the stochastic discount factor process includes a martingale component, then the use of the Perron–Frobenius eigenvalue and function recovers a long-term risk neutral pricing measure that is distorted by this martingale component.²⁵ By expanding the space of functions used in seeking a solution to the Perron–Frobenius problem, we avoid assuming that long-term risk return tradeoffs are degenerate under subjective beliefs, but we inherit an identification problem. There is typically a whole family of solutions with little guidance as to which solution should be selected by an analyst.

Many structural models of asset pricing that are motivated by empirical evidence have non-trivial martingale components in the stochastic discount factors. These martingales characterize what probability is actually recovered by the application of Perron–Frobenius Theory. We illustrated this outcome in one example with long-run risk components to the macroeconomy and with investors that have non-separable recursive preferences. We also provided a unifying discussion of the empirical literature that derives non-parametric bounds for the magnitude of the martingale component and that finds a quantitatively large role for this martingale component for valuation. Finally, we suggested ways how to further expand the set of testable implications in this literature and to combine information about subjective beliefs from Arrow prices with the observed time series evolution.

In our previous work we showed how Perron–Frobenius Theory helps us understand risk-return tradeoffs. The probability measure identified by Perron–Frobenius Theory absorbs the long-term

²⁴See also [Garcia et al. \(2010\)](#) for a recent survey of econometric methods for the study of options prices and [Gagliardini et al. \(2011\)](#) for a discussion of the econometric impact of combining a rich set of options prices for a limited number of dates with time series evidence over many periods. This latter work also imposes rational expectations.

²⁵Recent working papers by [Qin and Linetsky \(2014a,b\)](#) provide additional results in the continuous state space environment, with explicit connections to previous results in [Hansen and Scheinkman \(2009\)](#) and [Hansen and Scheinkman \(2014\)](#).

risk adjustments. Its naive use can distort the risk-return tradeoff in unintended ways. One might argue, however, that the dynamics under the Perron–Frobenius probability measure are of interest precisely because this measure adjusts for the long-term riskiness of the macroeconomy. While we see value in using this probability measure prospectively, our analysis makes clear that the resulting forecasts are slanted in a particular but substantively interesting way.

Finally, 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 Multiplicative functional

The construct of a multiplicative functional is used elsewhere in the probability literature and in [Hansen and Scheinkman \(2009\)](#). Formally a *multiplicative functional* is a process M that is adapted to \mathfrak{F} , initialized at $M_0 = 1$ and, with a slight abuse of notation:

$$M_t(Y) = M_\tau(Y)M_{t-\tau}(\theta_\tau(Y)). \quad (29)$$

In formula (29), θ_τ is the shift operator that moves the time subscript of Y by τ , that is, $(\theta_\tau(Y))_s = Y_{\tau+s}$.

We generalize this construct by building an *extended multiplicative functional*. The process M is an extended multiplicative functional if M_0 is a strictly positive (Borel measurable) function of X_0 and $\{M_t/M_0 : t \in T\}$ is a multiplicative functional. This allows the process M to be initialized at M_0 different from unity. In this paper we drop the use of the term *extended* for pedagogical convenience.

B Perron–Frobenius Theory

Proof of Proposition 4.3. Suppose there are two solutions $\hat{\eta}, \hat{e}$ and $\check{\eta}, \check{e}$. Hence:

$$\exp(\hat{\eta}t) \frac{\hat{H}_t \hat{e}(X_0)}{\hat{H}_0 \hat{e}(X_t)} = \exp(\check{\eta}t) \frac{\check{H}_t \check{e}(X_0)}{\check{H}_0 \check{e}(X_t)}$$

or if $k(x_t) = \frac{\check{e}(X_t)}{\hat{e}(X_t)} > 0$, and $\eta = \check{\eta} - \hat{\eta}$

$$\exp(-\eta t) \frac{\hat{H}_t k(X_t)}{\hat{H}_0 k(X_0)} = \frac{\check{H}_t}{\check{H}_0},$$

Calculating expected values in both sides and using the fact that \check{H} is a martingale we obtain that for $H = \hat{H}$

$$E^H [k(X_t) | X_0 = x] = \exp(\eta t)k(x) \quad (30)$$

In what follows we consider the discrete-time case. The continuous-time case uses an identical approach, with the obvious changes. First note that for a bounded function f the Law of Large Numbers implies that

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N E^H [f(X_t) | X_0 = x] = E^H [f(X_0)] \quad (31)$$

for $H = \hat{H}$ and $H = \check{H}$. Consider three cases. First suppose that $\eta < 0$. Set

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

Since $\eta < 0$, the right-hand side of (30) converges to zero for each x as $t \rightarrow \infty$. Thus, for $H = \hat{H}$

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

Thus we have established a contradiction.

Next suppose that $\eta > 0$. Note that for $H = \check{H}$

$$E^H \left[\frac{1}{k(X_t)} | X_0 = x \right] = \exp(-\eta t) \frac{1}{k(x)}. \quad (32)$$

Form

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

Since $\eta > 0$, the right-hand side of (32) converges to zero for each x as $t \rightarrow \infty$. Thus,

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

We have again established a contradiction.

Finally, suppose $\eta = 0$. Then again for $H = \check{H}$,

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

for all x . From (31):

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

for $k^n = \min\{1/k, n\}$. The same equality applies to the limit as $n \rightarrow \infty$ 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 $\frac{\widehat{H}_t}{\widehat{H}_0} = \frac{\check{H}_t}{\check{H}_0}$ with probability one. \square

C Long-term valuation limits

First we verify that the approximation results described in Section 5.1 hold under weaker conditions than Restriction 5.1. We assume instead:

$$\liminf_{t \rightarrow \infty} \widehat{E} [\psi(X_t) | X_0 = x] > 0 \quad (33)$$

and that

$$\liminf_{t \rightarrow \infty} \widehat{E} \left[\frac{\psi(X_t)}{\widehat{e}(X_t)} | X_0 = x \right] > 0.$$

Notice that these assumptions follows if for every bounded function f

$$\lim_{t \rightarrow \infty} \widehat{E} [f(X_t) | X_0 = x] = \widehat{E} [f(X_0)] \quad (34)$$

almost surely. [Meyn and Tweedie \(2009\)](#) establish in Theorem 13.3.3 on page 327 that (34) holds for bounded functions f , provided X is aperiodic and positive Harris recurrent under the measure \widehat{P} .

Furthermore, we assume that

$$\widehat{E} \left[\frac{\psi(X_t)}{\widehat{e}(X_t)} \right] < \infty, \quad \widehat{E} [\psi(X_t)] < +\infty,$$

and that X satisfies Restriction 4.2 under \widehat{P} . Then,

$$\begin{aligned} \frac{1}{N} \log \widehat{E} [\psi(X_N) | X_0 = x] &\leq \frac{1}{N} \log \left(\widehat{E} \left[\sum_{t=1}^N \psi(X_t) | X_0 = x \right] \right) \\ &= \frac{1}{N} \log N + \frac{1}{N} \log \left(\widehat{E} \left[\frac{1}{N} \sum_{t=1}^N \psi(X_t) | X_0 = x \right] \right). \end{aligned}$$

Result (33) implies that the lim inf on the left-hand side converges to zero. The lim sup of the left-hand side also converges to zero. To verify this, note that the Law of Large Numbers and the resulting almost sure convergence extend to the time-series averages of conditional expectations,

$$\lim_{N \rightarrow \infty} \widehat{E} \left[\frac{1}{N} \sum_{t=1}^N \psi(X_t) | X_0 = x \right] = \widehat{E} [\psi(X_0)]$$

almost surely. Hence both terms on the right-hand side of (35) converge to zero with probability one. Consequently, lim sup on the left-hand side does as well. Given both the lim inf and lim sup on the left-hand side of (35) to zero the left-hand side must converge to zero almost surely.

The same logic implies that

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \widehat{E} \left[\frac{\psi(X_N)}{\widehat{e}(X_N)} | X_0 = x \right] = 0,$$

with probability one.²⁶

Next introduce stochastic growth into the analysis as in Section 5.2. First note that the multiplicative functional SG satisfies Restriction 3.2. Let η^* denote the Perron–Frobenius eigenvalue using SG in place of S when solving Problem 4.1. This entails solving:

$$E [S_t G_t e^*(X_t) | X_0 = x] = \exp(\eta^* t) e^*(x)$$

and selecting the eigenvalue-eigenfunction pair such that the implied martingale induces stationarity and ergodicity. Imitating our calculation of yields on securities with stationary payoffs

$$E \left(S_t G_t \left[\frac{\widehat{e}(X_t)}{\widehat{e}(X_0)} \right] | X_0 = x \right) = \exp(t\eta^*) E^* \left[\frac{\widehat{e}(X_t) e^*(X_0)}{\widehat{e}(X_0) e^*(X_t)} | X_0 = x \right]$$

where E^* is constructed using the martingale H^* with increments:

$$\frac{H_t^*}{H_0^*} = \exp(-\eta^* t) S_t G_t \left[\frac{e^*(X_t)}{e^*(X_0)} \right].$$

Similarly,

$$E(S_t G_t | X_0 = x) = \exp(t\eta^*) E^* \left(\left[\frac{e^*(X_0)}{e^*(X_t)} \right] | X_0 = x \right)$$

Suppose that H^* induces a probability measure under which Restriction 5.1 is satisfied. In addition suppose

²⁶For the continuous-time case, sample the Markov process at integer points in time. The process will remain stationary under \widehat{P} but not necessarily ergodic. The Law of Large Numbers still applies but with a limit that is the expectation conditioned on invariant events. The previous argument with these modifications establishes the limiting behavior.

that

$$E^* \left[\frac{\widehat{e}(X_t)}{e^*(X_t)} \right] < \infty, \quad E^* \left[\frac{1}{e^*(X_t)} \right] < \infty.$$

From Restriction 5.1 we obtain,

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{1}{t} E^* \left[\frac{\widehat{e}(X_t)}{e^*(X_t)} \mid X_0 = x \right] &= 0 \\ \lim_{t \rightarrow \infty} \frac{1}{t} \log E^* \left[\frac{1}{e^*(X_t)} \mid X_0 = x \right] &= 0. \end{aligned}$$

Thus

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{1}{t} \log E \left(S_t G_t \left[\frac{\widehat{e}(X_t)}{\widehat{e}(X_0)} \right] \mid X_0 = x \right) &= \\ = \eta^* + \lim_{t \rightarrow \infty} \frac{1}{t} [\log e^*(x) - \log \widehat{e}(x)] + \lim_{t \rightarrow \infty} \frac{1}{t} E^* \left[\frac{\widehat{e}(X_t)}{e^*(X_t)} \mid X_0 = x \right] &= \eta^*, \end{aligned}$$

and similarly

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{1}{t} \log E (S_t G_t \mid X_0 = x) &= \\ = \eta^* + \lim_{t \rightarrow \infty} \frac{1}{t} \log e^*(x) + \lim_{t \rightarrow \infty} \frac{1}{t} \log E^* \left[\frac{1}{e^*(X_t)} \mid X_0 = x \right] &= \eta^*. \end{aligned}$$

Hence

$$\lim_{t \rightarrow \infty} y_t[G](x) = -\eta + \eta^* - \eta^* = -\eta$$

D Derivations for the model with predictable consumption dynamics

In this section, we provide derivations for the model analyzed in Section 6. We will focus on the analysis of the Perron–Frobenius problem. A complete analysis of the model can be found in Hansen (2012) and in the appendix in Borovička et al. (2014).

D.1 Martingale decomposition

We solve the Perron–Frobenius problem

$$E [S_t \widehat{e}(X_t) \mid X_0 = x] = \exp(\widehat{\eta}t) \widehat{e}(x)$$

where S is a multiplicative functional parameterized by coefficients $(\beta_s(x), \alpha_s(x))$. Since the problem holds for every t , then it also holds in the limit (as long as it exists)

$$\lim_{t \searrow 0} \frac{1}{t} [E [S_t \widehat{e}(X_t) \mid X_0 = x] - \exp(\widehat{\eta}t) \widehat{e}(x)] = 0.$$

The limit yields the partial differential equation

$$\mathbb{S}e = \eta e$$

where, in the general Brownian information setup given by equations (15), the infinitesimal generator \mathbb{S} is given by

$$\mathbb{S}\widehat{e} = \left(\beta_s + \frac{1}{2} |\alpha_s|^2 \right) \widehat{e} + \widehat{e}_x \cdot (\mu + \sigma \alpha_s) + \frac{1}{2} \text{tr} [\widehat{e}_{xx} \sigma \sigma'].$$

Equation (D.1) is therefore a second-order partial differential equation, and we are looking for a solution in the form of a number $\widehat{\eta}$ and a strictly positive function \widehat{e} . Hansen and Scheinkman (2009) show that if there are multiple such solutions, then only the solution associated with the lowest value of $\widehat{\eta}$ can generate ergodic dynamics under the implied change of measure.

In the case of the long-run risk model introduced in Section 6 and parameterized by (22)–(24), we can guess

$$\widehat{e}(x) = \exp(\bar{e}_1 x_1 + \bar{e}_2 x_2)$$

which leads to the system of equations

$$\begin{aligned} \eta &= \bar{\beta}_{s,0} - \bar{\beta}_{s,11} \iota_1 - \bar{\beta}_{s,12} \iota_2 - \bar{e}_1 (\bar{\mu}_{11} \iota_1 + \bar{\mu}_{12} \iota_2) - \bar{e}_2 \bar{\mu}_{22} \iota_2 \\ 0 &= \bar{\beta}_{s,11} + \bar{\mu}_{11} \bar{e}_1 \\ 0 &= \bar{\beta}_{s,12} + \frac{1}{2} |\bar{\alpha}_s|^2 + \bar{e}_1 (\bar{\mu}_{12} + \bar{\sigma}_1 \bar{\alpha}_s) + \frac{1}{2} \bar{e}_1^2 |\bar{\sigma}_1|^2 + \bar{e}_2 (\bar{\mu}_{22} + \bar{\sigma}_2 \bar{\alpha}_s + \bar{e}_1 \bar{\sigma}_1 \bar{\sigma}_2') + \frac{1}{2} (\bar{e}_2)^2 |\bar{\sigma}_2|^2 \end{aligned}$$

The coefficients $\bar{\beta}_{s,\cdot}$ and $\bar{\alpha}_s$ will be determined below. The last equation is a quadratic equation for \bar{e}_2 and we choose the solution for \bar{e}_2 that leads to the smaller value of $\widehat{\eta}$. From the decomposition

$$S_t = \exp(\widehat{\eta} t) \frac{\widehat{e}(X_0) \widehat{H}_t}{\widehat{e}(X_t) \widehat{H}_0}$$

we can extract the martingale \widehat{H} :

$$d \log \widehat{H}_t = d \log S_t - \widehat{\eta} dt + d \log \widehat{e}(X_t)$$

and thus

$$\frac{d\widehat{H}_t}{\widehat{H}_t} = \sqrt{X_{2t}} (\bar{\alpha}_s + \bar{\sigma}_1' \bar{e}_1 + \bar{\sigma}_2' \bar{e}_2) \cdot dW_t \doteq \sqrt{X_{2t}} \widehat{\alpha}_h \cdot dW_t$$

This martingale implies a change of measure such that \widehat{W} defined as

$$d\widehat{W}_t = -\sqrt{X_{2t}} \widehat{\alpha}_h dt + dW_t$$

is a Brownian motion under the new measure \widehat{P} . Under the change of measure implied by \widehat{H} , we can write the joint dynamics of the model as

$$\begin{aligned} dX_{1t} &= [\widehat{\mu}_{11} (X_{1t} - \widehat{\iota}_1) + \widehat{\mu}_{12} (X_{2t} - \widehat{\iota}_2)] dt + \sqrt{X_{2t}} \bar{\sigma}_1 d\widehat{W}_t \\ dX_{2t} &= \widehat{\mu}_{22} (X_{2t} - \widehat{\iota}_2) dt + \sqrt{X_{2t}} \bar{\sigma}_2 d\widehat{W}_t \end{aligned}$$

where

$$\begin{aligned} \widehat{\mu}_{11} &= \bar{\mu}_{11} & \widehat{\mu}_{12} &= \bar{\mu}_{12} + \bar{\sigma}_1 \widehat{\alpha}_h & \widehat{\mu}_{22} &= \bar{\mu}_{22} + \bar{\sigma}_2 \widehat{\alpha}_h \\ \widehat{\iota}_2 &= \frac{\bar{\mu}_{22}}{\widehat{\mu}_{22}} \iota_2 & \widehat{\iota}_1 &= \iota_1 + (\bar{\mu}_{11})^{-1} (\bar{\mu}_{12} \iota_2 - \widehat{\mu}_{12} \widehat{\iota}_2). \end{aligned}$$

Similarly, every multiplicative functional M with parameters given by (24) can be rewritten as

$$d \log M_t = \left[\widehat{\beta}_0 + \widehat{\beta}_{11} (X_{1t} - \widehat{\iota}_1) + \widehat{\beta}_{12} (X_{2t} - \widehat{\iota}_2) \right] dt + \sqrt{X_{2t}} \widehat{\alpha} \cdot d\widehat{W}_t$$

where

$$\begin{aligned} \widehat{\beta}_0 &= \bar{\beta}_0 + \bar{\beta}_{11} (\widehat{\iota}_1 - \iota_1) + \bar{\beta}_{12} (\widehat{\iota}_2 - \iota_2) + (\bar{\alpha} \cdot \widehat{\alpha}_h) \widehat{\iota}_2 \\ \widehat{\beta}_{11} &= \bar{\beta}_{11} \quad \widehat{\beta}_{12} = \bar{\beta}_{12} + \bar{\alpha} \cdot \widehat{\alpha}_h. \end{aligned}$$

D.2 Value function and stochastic discount factor for recursive utility

We choose a convenient choice for representing continuous values. Similar to the discussion in [Schroder and Skiadas \(1999\)](#), we use the counterpart to discounted expected logarithmic utility.

$$d \log V_t = \mu_{v,t} dt + \sigma_{v,t} \cdot dW_t$$

The local evolution satisfies:

$$\mu_{v,t} = \delta \log V_t - \delta \log C_t - \frac{1-\gamma}{2} |\sigma_{v,t}|^2 \quad (36)$$

When $\gamma = 1$ this collapses to the discounted expected utility recursion.

Let

$$\log V_t = \log C_t + v(X_t)$$

and guess that

$$v(x) = \bar{v}_0 + \bar{v}_1 \cdot x_1 + \bar{v}_2 x_2$$

We may compute $\mu_{v,t}$ by applying the infinitesimal generator to $\log C + v(X)$. In addition,

$$\sigma_{v,t} = \alpha_c(X_t) + \sigma(X_t)' \frac{\partial}{\partial x} v(X_t).$$

Substituting into (36) leads to a set of algebraic equations

$$\begin{aligned} \delta \bar{v}_0 &= \bar{\beta}_{c,0} - \iota_1 (\bar{\beta}_{c,1} + \bar{\mu}_{11} \bar{v}_1) - \iota_2 (\bar{\beta}_{c,2} + \bar{\mu}_{12} \bar{v}_1 + \bar{\mu}_{22} \bar{v}_2) \\ \delta \bar{v}_1 &= \bar{\beta}_{c,1} + \bar{\mu}_{11} \bar{v}_1 \\ \delta \bar{v}_2 &= \bar{\beta}_{c,2} + \bar{\mu}_{12} \bar{v}_1 + \bar{\mu}_{22} \bar{v}_2 + \frac{1}{2} (1-\gamma) |\bar{\alpha}_c + \bar{\sigma}'_1 \bar{v}_1 + \bar{\sigma}'_2 \bar{v}_2|^2 \end{aligned}$$

which can be solved for the coefficients \bar{v}_i . The third equation is a quadratic equation for \bar{v}_2 that has a real solution if and only if

$$\begin{aligned} D &= \left[\bar{\mu}_{22} - \delta + (1-\gamma) (\bar{\alpha}_c + \bar{\sigma}'_1 \bar{v}_1)' \bar{\sigma}'_2 \right]^2 - \\ &\quad - 2(1-\gamma) |\bar{\sigma}_2|^2 \left(\bar{\beta}_{c,2} + \bar{\mu}_{12} \bar{v}_1 + \frac{1}{2} (1-\gamma) |\bar{\alpha}_c + \bar{\sigma}'_1 \bar{v}_1|^2 \right) \geq 0. \end{aligned}$$

In particular, the solution will typically not exist for large values of γ . If the solution exists, it is given by

$$\bar{v}_2 = \frac{-[\bar{\mu}_{22} - \delta + (1-\gamma) (\bar{\alpha}_c + \bar{\sigma}'_1 \bar{v}_1)' \bar{\sigma}'_2] \pm \sqrt{D}}{(1-\gamma) |\bar{\sigma}'_2|^2}$$

The solution with the minus sign is the one that interests us.

The resulting stochastic discount factor has two components. One component is the intertemporal marginal rate of substitution for discounted log utility and the other is a martingale constructed from the continuation value

$$d \log S_t = -\delta dt - d \log C_t + d \log H_t^*$$

where H^* is a martingale given by

$$\frac{dH_t^*}{H_t^*} = \sqrt{X_{2,t}} (1 - \gamma) (\bar{\alpha}_c + \bar{\sigma}'_1 \bar{v}_1 + \bar{\sigma}'_2 \bar{v}_2)' dW_t.$$

This determines the coefficients $(\beta_s(x), \alpha_s(x))$ of the stochastic discount factor.

When we choose the ‘minus’ solution in equation (D.2), then \hat{S} implies a change of measure that preserves ergodicity. Notice that while H^* is a martingale, it is distinct from \hat{H}_t/\hat{H}_0 as long as the consumption process itself contains a nontrivial martingale component.

D.3 Conditional expectations of multiplicative functionals

In order to compute asset prices and their expected returns, we need to compute conditional expectations of multiplicative functionals M parameterized by (22)–(24). These conditional expectations are given by

$$E[M_t | X_0 = x] = \exp[\theta_0(t) + \theta_1(t) \cdot x_1 + \theta_2(t)x_2],$$

where parameters $\theta_i(t)$ satisfy a system of ordinary differential equations derived in Hansen (2012) and in the appendix of Borovička et al. (2014).

References

- Aït-Sahalia, Yacine and Andrew Lo. 2000. Nonparametric risk management and implied risk aversion. *Journal of Econometrics* 94 (1-2):9–51.
- Almeida, Caio and René Garcia. 2013. Robust Economic Implications of Nonlinear Pricing Kernels.
- Alvarez, Fernando and Urban J. Jermann. 2005. Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth. *Econometrica* 73 (6):1977–2016.
- Anderson, Evan W., Lars Peter Hansen, and Thomas J. Sargent. 2003. A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection. *Journal of the European Economic Association* 1 (1):68–123.
- 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. 1994. Addendum to ‘Growth-Optimal Portfolio Restrictions on Asset Pricing Models’. Unpublished manuscript.
- . 1997. Growth-Optimal Portfolio Restrictions on Asset Pricing Models. *Macroeconomic Dynamics* 1 (2):333–354.
- Bansal, Ravi and Amir Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59 (4):1481–1509.
- Bidder, Rhys and Matthew E. Smith. 2013. Doubts and Variability: A Robust Perspective on Exotic Consumption Series. Federal Reserve Bank of San Francisco Working Paper 2013–28.
- Borovička, Jaroslav, Lars Peter Hansen, Mark Hendricks, and José A. Scheinkman. 2011. Risk Price Dynamics. *Journal of Financial Econometrics* 9 (1):3–65.
- Borovička, Jaroslav, Lars Peter Hansen, and José A. Scheinkman. 2014. Shock Elasticities and Impulse Responses. *Mathematics and Financial Economics* 8 (4):333–354.
- Breedon, Douglas T. and Robert H. Litzenberger. 1978. Prices of State-Contingent Claims Implicit in Option Prices. *The Journal of Business* 51 (4):621–51.
- Breiman, Leo. 1982. *Probability*. Volume 7 of Classics in Applied Mathematics. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM).

- Carr, Peter and Jiming Yu. 2012. Risk, Return and Ross Recovery. *The Journal of Derivatives* 20 (1):38–59.
- Chen, Xiaohong, Victor Chernozhukov, Sokbae Lee, and Whitney K. Newey. 2014. Local Identification of Nonparametric and Semiparametric Models. *Econometrica* 82 (2):785–809.
- Christensen, Timothy M. 2014. Nonparametric Identification of Positive Eigenfunctions. *Econometric Theory* FirstView:1–21.
- 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 Larry G. Epstein. 1992. Asset Pricing with Stochastic Differential Utility. *The Review of Financial Studies* 5 (3):411–436.
- 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.
- Gagliardini, Patrick, Christian Gourieroux, and Eric Renault. 2011. Efficient Derivative Pricing by the Extended Method of Moments. *Econometrica* 79 (4):1181–1232.
- Garcia, René, Eric Ghysels, and Eric Renault. 2010. The Econometrics of Option Pricing. In *Handbook of Financial Econometrics: Tools and Techniques*, vol. 1 of *Handbooks in Finance*, edited by Lars Peter Hansen and Yacine Aït-Sahalia, chap. 9, 479 – 552. San Diego: North-Holland.
- Ghosh, 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. Nobel Lecture: Uncertainty Outside and Inside Economic Models. *Journal of Political Economy* 122 (5):945–987.
- 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, John C. Heaton, Junghoon Lee, and Nikolai Roussanov. 2007. Intertemporal Substitution and Risk Aversion. In *Handbook of Econometrics*, vol 6A, edited by James J. Heckman and Edward E. Leamer, 3967–4056. Amsterdam: North-Holland.
- Hansen, Lars Peter, John C. Heaton, and Nan Li. 2008. Consumption Strikes Back? Measuring Long-Run Risk. *Journal of Political Economy* 116:260–302.
- Harrison, J. Michael and David M. Kreps. 1979. Martingales and Arbitrage in Multiperiod Securities Markets. *Journal of Economic Theory* 20 (3):381–408.
- Hilscher, Jens, Alon Raviv, and Ricardo Reis. 2014. Inflating Away the Public Debt? Columbia University.
- Kazemi, Hossein B. 1992. An Intertemporal Model of Asset Prices in a Markov Economy with a Limiting Stationary Distribution. *Review of Financial Studies* 5 (1):85–104.
- Kreps, David M. and Evan L. Porteus. 1978. Temporal Resolution of Uncertainty and Dynamic Choice Theory. *Econometrica* 46 (1):185–200.
- Meyn, Sean P. and Richard L. Tweedie. 1993. Stability of Markovian Processes II: Continuous-Time Processes and Sampled Chains. *Advances in Applied Probability* 25 (3):487–517.
- . 2009. *Markov Chains and Stochastic Stability*. Cambridge, United Kingdom: Cambridge University Press.
- Park, Hyungbin. 2014. Ross Recovery with Recurrent and Transient Processes.
- Qin, Likuan and Vadim Linetsky. 2014a. Long Term Risk: A Martingale Approach. Mimeo, Northwestern University.
- . 2014b. Positive Eigenfunctions of Markovian Pricing Operators: Hansen–Scheinkman Factorization and Ross Recovery. Mimeo, Northwestern University.
- Ross, Stephen A. 1978. A Simple Approach to the Valuation of Risky Streams. *The Journal of Business* 51 (3):453–75.

- Ross, Stephen A. 2015. The Recovery Theorem. *Journal of Finance* 70 (2):615–648.
- Schroder, Mark and Costis Skiadas. 1999. Optimal Consumption and Portfolio Selection with Stochastic Differential Utility. *Journal of Economic Theory* 89 (1):68–126.
- Stutzer, Michael. 1996. A Simple Nonparametric Approach to Derivative Security Valuation. *The Journal of Finance* 51 (5):1633–1652.
- Walden, Johan. 2014. Recovery with Unbounded Diffusion Processes.