

Volatility of the Stochastic Discount Factor, and the Distinction between Risk-Neutral and Objective Probability Measures

Gurdip Bakshi, Zhiwu Chen, Erik Hjalmarsson*

October 5, 2004

*Bakshi is at Department of Finance, Smith School of Business, University of Maryland, College Park, MD 20742, Tel: 301-405-2261, email: gbakshi@rhsmith.umd.edu, and website: www.rhsmith.umd.edu/finance/gbakshi/; Chen at Yale School of Management, 135 Prospect Street, New Haven, CT 06520, Tel: (203) 432-5948, email: zhiwu.chen@yale.edu and website: www.som.yale.edu/faculty/zc2; and Hjalmarsson is at Department of Economics, Yale University, 28 Hillhouse, New Haven, CT 06510. We thank Steve Ross, Rene Stulz, Mark Lowenstein, Dilip Madan, and Nengjiu Ju for conversations with them on this topic.

Volatility of the Stochastic Discount Factor, and the Distinction between Risk-Neutral and Objective Probability Measures

Abstract

This paper derives a measure that characterizes the distance between the risk-neutral and the objective probability measures for any candidate asset pricing model. We formally show that the distance metric is equal to the volatility of the stochastic discount factor. This theoretical result gives an alternative interpretation to the Hansen-Jagannathan bounds: they provide a lower bound for the distance between the objective and the risk-neutral probability measures. Our empirical application provides support for the notion that the crash of 1987 has widened the wedge between the risk-neutral and the objective probability measures.

JEL Classification: G10

Keywords: Risk-neutral measures, objective probability measures, volatility of the stochastic discount factor, no-arbitrage, Hansen-Jagannathan bounds.

1 Introduction

Risk-neutral probability measure and its objective counterpart share the same zero-probability events and are mathematically equivalent. The equivalence restriction, however, tells the financial economists little about how the risk-neutral and objective distributions are related to one other. Consider a two-state economy in which the true probability is some small $\eta > 0$ for the first state and $1 - \eta$ for the other. In this case, the risk-neutral measure may assign a probability of $1 - \eta$ to the first state and η to the second. Even though the risk-neutral and the objective probability measures are equivalent, the two measures are probabilistically a world apart. Consequently, any conclusions drawn about the true future distribution, but based on an estimated risk-neutral distribution, are bound to be flawed. The work of Rubinstein (1994) and Jackwerth and Rubinstein (1996) shows, for instance, that the market-index return distributions are far left-skewed under the risk-neutral measure but essentially symmetric under the objective probability measure.

In theory, how different can a given risk-neutral distribution be from its objective counterpart? How can one formally gauge their difference and what are the sources of the dichotomy? Given the increasingly popular use of risk-neutral distributions in financial modeling, it is important to form a better understanding of the relationship between risk-neutral and objective probability distributions. The intent of this paper is to provide guidance on this issue under fairly general assumptions about the underlying economy. Specifically, we propose an economically sensible and technically sound distance-metric that captures the dichotomy between the risk-neutral and objective probability measures. Moreover, we show that the derived metric has a natural interpretation in standard financial theory.

The starting point for our analysis relies on a result by Harrison and Kreps (1979) that, absent of arbitrage opportunities, there is a one-to-one correspondence between risk-neutral probability measures and positive stochastic discount factors. Thus, every admissible risk-neutral probability measure is defined by a unique stochastic discount factor, and every stochastic discount factor can be represented by a unique risk-neutral measure. Armed with this result and the Radon-Nikodym theorem, we show that the mean-square distance between the risk-neutral and the objective probability distribution, in a sense to be made precise, is, in fact, equal to the volatility of the stochastic discount factor, up to a constant of proportionality. This theoretical result shows how the distance between risk-neutral and objective probabilities can be expressed in terms of the properties of the underlying

valuation standards of the economy. It also gives us an alternative interpretation of the Hansen and Jagannathan (1991) bounds, namely that they provide a lower bound for the distance between the objective and the risk-neutral probability measures. Like Hansen and Jagannathan (1991), we can estimate the volatility of the stochastic discount factor and the distance between the risk-neutral and the objective probability distributions under plausible parametric assumptions.

In an empirical exercise, we use the methods of Hansen and Jagannathan (1991) to infer lower bounds, from observed security prices, on the volatility of the stochastic discount factor or, equivalently, on the distance between the risk-neutral and the objective probability distributions. Using data on monthly S&P 500 index returns from 1926-2003, our methods find evidence that there is a significant increase in the distance between the risk-neutral and the objective probability distributions after the crash in 1987. We also, however, find some evidence that this effect is dissipating in recent years. The theoretical and empirical results contribute to the existing literature on the market-index and the index option markets in two novel ways. First, the methods employed in our paper are completely model-free in the sense that they do not rely on any specific model for the price dynamics or investor preferences, apart from the absence of arbitrage. Second, we appeal to a single return series, rather than the cross-section of option prices, to demonstrate the dichotomy between risk-neutral measures and its objective counterpart. The caveat is, of course, that we can only estimate a lower bound on the distance, and not obtain a point estimate. In general, we cannot establish whether the actual distance shifted when the lower bound shifted.

The rest of this paper is organized as follows. Section 2 formalizes the idea behind the distance metric and derives the main theorem. Section 3 provides additional economic interpretation and presents parametric examples. Section 4 is devoted to an empirical illustration. The final section 5 concludes.

2 A Distance Metric for Arbitrage-Free Economies

Consider a frictionless economy endowed with a probability space $(\Omega, \mathbf{F}, \mathbf{p})$, where Ω is the state space, \mathbf{F} the corresponding sigma-field, and p the associated true, or objective, probability measure. Assume that this measure p is shared by all market participants. Asset payoffs in this economy are modeled by the linear space \mathcal{L}^2 of square-integrable random

variables. With this formal structure, we can measure differences between securities by the standard mean-square norm:

$$\|x\| \equiv \sqrt{E(x^2)} \quad \text{for any } x \in \mathcal{L}^2, \quad (1)$$

where $E(\cdot)$ is the expectation with respect to p .

To clarify the terminology, suppose that there is another probability measure p^* defined on (Ω, \mathbf{F}) . We say p^* is probabilistically equivalent to p if p^* and p share exactly the same null events; that is, for any event $A \in \mathbf{F}$, $p^*(A) = 0$ if and only if $p(A) = 0$. Two probability measures can therefore be equivalent and yet assign diametrically different (positive) probability masses to the same events - as long as they agree on the zero-probability events.

Assume that there are N traded assets, with payoffs $x_n \in \mathcal{L}^2$ and prices $q_n \in \mathfrak{R}$, for each $n = 1, \dots, N$. These N securities can include stocks, bonds, equity options, and so on. For convenience, let $X = (x_1, \dots, x_N)'$ and $q = (q_1, \dots, q_N)'$. In addition, assume that one of the N assets is risk-free and offers an interest rate of r_0 . The existence of a risk-free asset is not necessary for our results but provides for an easier interpretation and exposition.¹ Given these securities, the set of all marketed payoffs is

$$\mathcal{M} \equiv \{x \in \mathcal{L}^2 : \exists \alpha \in \mathfrak{R}^N \text{ such that } \sum_{i=1}^N \alpha_i \cdot x_i = x\}. \quad (2)$$

The associated cost function $\pi[x]$ for each marketed payoff $x \in \mathcal{M}$ is as follows:

$$\pi[x] \equiv \min_{\alpha \in \mathfrak{R}^N} \alpha'q, \quad \text{subject to } \alpha'X = x. \quad (3)$$

That is, $\pi[x]$ is the minimum cost of obtaining x . As is standard in the literature, we take the security payoff-price pair (X, q) as given and characterize the pricing rules implied by these securities.

The security market (\mathcal{M}, π) is said to be free of arbitrage if, for every $x \in \mathcal{M}$ such that $x \geq 0$ with probability one and $\|x\| > 0$, $\pi(x) > 0$. Absent of arbitrage, therefore, all non-negative payoffs, which are positive with positive probability, must have positive

¹In our theoretical characterizations, we only rely on the fact that $E(d^*) = \frac{1}{1+r_0}$, where d^* is the stochastic discount factor (defined in theorem 1). In the absence of a risk-free asset, all our results carry through by simply replacing $\frac{1}{1+r_0}$ by $E(d^*)$ and a slight change in interpretation.

prices. The following well known result, due to Ross (1978) and Harrison and Kreps (1979), forms the basis of our subsequent analysis:

Theorem 1 *Suppose that in the frictionless economy there is at least one limited-liability security $x \in \mathcal{M}$ such that $x > 0$ with positive probability. Then, the following statements are equivalent:*

1. *The security market is free of arbitrage.*
2. *There exists a stochastic discount factor $d^* \in \mathcal{L}^2$ such that $d^* > 0$ with probability one and*

$$\pi[x] = E(xd^*), \text{ for each } x \in \mathcal{M}. \quad (4)$$

3. *There exist an equivalent probability measure p^* such that*

$$\pi[x] = \frac{1}{1+r_0} E^*(x), \text{ for each } x \in \mathcal{M}, \quad (5)$$

where $E^*(\cdot)$ stands for the expectation with respect to measure p^* .

The stochastic discount factor d^* satisfying (4) is also referred to as Arrow-Debreu state-price density, or the pricing operator. It determines how the future state-by-state payoffs are to be converted into today's price. For convenience, let D^* be the set of all d^* satisfying (4):

$$D^* \equiv \{d^* \in \mathcal{L}^2 : d^* > 0 \text{ almost surely and } E(xd^*) = \pi[x] \text{ for each } x \in \mathcal{M}\}. \quad (6)$$

Equation (5) says that under the equivalent measure p^* one can value all assets as if the economy was risk-neutral. For this reason, we often refer to p^* as a risk-neutral probability measure. Let P^* collect all risk-neutral equivalent probability measures p^* satisfying (5).

To introduce a standard distance metric for probability measures, assume first that Ω is countable and let $\{A_i \in \mathbf{F} : i = 1, 2, \dots\}$ be a partition of Ω , such that (i) all events A_i are disjoint, (ii) $\bigcup_{i=1,2,\dots} A_i = \Omega$, and (iii) for each i , A_i contains no subevent that is also in \mathbf{F} . In other words, $\{A_i \in \mathbf{F} : i = 1, 2, \dots\}$ represents the finest partition of Ω contained

in \mathbf{F} . Define the absolute distance between p and p^* by

$$\delta[p, p^*] \equiv \sum_{i=1,2,\dots} |p^*(A_i) - p(A_i)|, \quad (7)$$

which equals zero if and only if p and p^* assign the same probability mass to every given event $A_i \in \mathbf{F}$. If the risk-neutral and objective probability distributions are different, $\delta[p, p^*]$ will simply gauge the closeness of the two probability functions. Alternatively, for a continuous state space Ω , we define

$$\delta[p, p^*] \equiv \int_{\Omega} |dp^*(x) - dp(x)|. \quad (8)$$

In the analysis below, it will generally be advantageous to scale $\delta[p, p^*]$ by $\frac{1}{(1+r_0)}$ and to this end define $\delta_1[p, p^*] \equiv \frac{1}{(1+r_0)} \delta[p, p^*]$.

The following Theorem makes clear the connection between the distance between p and p^* , and the stochastic discount factor, d^* .

Theorem 2 *In the arbitrage-free economy described above, the distance measure $\delta_1[p, p^*]$ satisfies*

$$\delta_1[p, p^*] = E \left| d^* - \frac{1}{1+r_0} \right| \quad (9)$$

$$\leq \left\| d^* - \frac{1}{1+r_0} \right\| \equiv \delta_2[p, p^*]. \quad (10)$$

Further, since $E(d^*) = \frac{1}{1+r_0}$,

$$\delta_2[p, p^*] = \left\| d^* - \frac{1}{1+r_0} \right\| \quad (11)$$

$$= \left\| d^* - E(d^*) \right\| \quad (12)$$

$$= \sigma[d^*], \quad (13)$$

where $\sigma[d^*]$ stands for the standard deviation of the stochastic discount factor d^* .

Proof: Each stochastic discount factor $d^* \in D^*$ corresponds to the Radon-Nikodym derivative which defines a risk-neutral measure $p^* \in P^*$. Furthermore, this is a unique one-to-one correspondence between D^* and P^* . To see this, first suppose $d^* \in D^*$. Let $\rho \equiv \frac{d^*}{E(d^*)}$,

which implies $\rho > 0$ almost surely, since $d^* > 0$ almost surely. Next, define a measure p^* :

$$p^*(A) \equiv E(1_A) = \int_A \rho dp, \quad (14)$$

for every event $A \in \mathbf{F}$, where 1_A equals 1 if event A occurs and zero otherwise. This p^* is a valid probability measure because (i) $E(\rho) = 1$ and (ii) $p^*(A) \geq 0$ for each $A \in \mathbf{F}$. It is equivalent to the objective probability p because $\rho > 0$ almost surely. Substituting ρ into (4), we have

$$\pi[x] = E(x d^*) = E(d^*) \int_{\Omega} x \frac{d^*}{E(d^*)} dp \quad (15)$$

$$= E(d^*) \int_{\Omega} x dp^* = \frac{1}{1+r_0} E^*(x), \quad (16)$$

which gives (5), where we used the fact that for the risk-free asset, $E(d^*) = \frac{1}{1+r_0}$. Therefore, d^* defines a unique risk-neutral measure p^* .

Now suppose that p^* is a risk-neutral measure satisfying (5). Then, by the Radon-Nikodym theorem, there must exist some $\rho \in \mathcal{L}^2$ such that $\rho > 0$ almost surely and $\frac{dp^*}{dp} = \rho$. Consequently, rewrite (5) as follows:

$$\pi[x] = \frac{1}{1+r_0} \int_{\Omega} x dp^* = \frac{1}{1+r_0} \int_{\Omega} x \rho dp = E(x d^*), \quad (17)$$

where $d^* \equiv \frac{1}{1+r_0} \rho$ and $d^* > 0$ almost surely. This means that $d^* \in D^*$. Thus, behind each risk-neutral probability measure $p^* \in P^*$ there is a corresponding stochastic discount factor $d^* \in D^*$.

Fix a risk-neutral probability $p^* \in P^*$, and let $d^* \in D^*$ be the corresponding stochastic discount factor defining p^* . Then,

$$dp^* = \frac{d^*}{E(d^*)} dp, \quad (18)$$

as $\frac{d^*}{E(d^*)}$ is the Radon-Nikodym derivative of p^* with respect to p . Substituting this into (7), and realizing that any random variable defined in this economy takes the same value

in every state contained in a given A_i defined above, we obtain

$$\delta[p, p^*] = E \left| \frac{d^*}{E(d^*)} - 1 \right| = (1 + r_0) E \left| d^* - \frac{1}{1 + r_0} \right|, \quad (19)$$

where we again relied on $E(d^*) = \frac{1}{1+r_0}$. Since $1 + r_0$ represents a scaling factor, the absolute distance between p and p^* is completely determined by the absolute distance between d^* and $\frac{1}{1+r_0}$,

$$\delta_1[p, p^*] = E \left| d^* - \frac{1}{1 + r_0} \right|. \quad (20)$$

The results for continuous state spaces follow in the same manner. By Hölder's inequality, equation (20) implies

$$0 \leq \delta_1[p, p^*] \leq \left\| d^* - \frac{1}{1 + r_0} \right\| \equiv \delta_2[p, p^*] \quad (21)$$

which gives the desired result. \square

In a true risk-neutral economy (where every investor is risk-neutral under the objective probability), $\frac{1}{1+r_0}$ is the corresponding discount factor, which is non-stochastic. Therefore, one can think of $\delta_1[p, p^*]$ as measuring the absolute closeness of d^* to the true risk-neutral discount factor. The closer the valuation rule represented by d^* is to true risk-neutral valuation, the closer the risk-neutral probability p^* is to its objective counterpart p .

Since both the absolute difference and the mean-square difference are valid distance measures, $\delta_1[p, p^*] = \delta_2[p, p^*] = 0$ if and only if the stochastic discount factor d^* equals the true risk-neutral discount factor $\frac{1}{1+r_0}$, which in turn holds if and only if the objective and the risk-neutral probability measures, p and p^* , are identical (provided $r_0 > 0$). Generally, when p^* differs from p , the mean-square distance $\delta_2[p, p^*]$ provides an upper bound on $\delta_1[p, p^*]$. For most cases, the closeness rankings of a given set of equivalent probability measures should be similar according to either $\delta_1[p, p^*]$ or $\delta_2[p, p^*]$.

The key theoretical result of the paper is the equivalence of $\delta_2[p, p^*]$ and $\sigma[d^*]$. It is thus the volatility level of the defining stochastic discount factor that determines the distance between a given risk-neutral distribution and its objective counterpart. This observation yields another interpretation to the empirical asset pricing literature, where the search has been for a sufficiently volatile intertemporal-marginal-rate-of-substitution model so as to explain the high historical equity premium levels. From the angle just discussed, that

search is effectively equivalent to looking for a risk-neutral distribution that is sufficiently different from the underlying objective distribution.

In order to better understand the distance metric $\delta_2 [p, p^*]$, it is useful to recast it in fundamental economic terms. Pick any arbitrary $d^* \in D^*$ and let $\bar{\pi}$ be the extension of π to all of \mathcal{L}^2 . That is, for any $x \in \mathcal{L}^2$, $\bar{\pi} [x] \equiv E (x d^*)$. Now, consider the pricing error incurred by using the risk-neutral discount factor $\frac{1}{1+r_0}$, rather than the risk adjusted d^* . We have, for $x \in \mathcal{L}^2$,

$$\left| \bar{\pi} [x] - \frac{1}{1+r_0} E (x) \right| = \left| E (x d^*) - \frac{1}{1+r_0} E (x) \right| \quad (22)$$

$$\leq \left\| d^* - \frac{1}{1+r_0} \right\| \|x\|, \quad (23)$$

by the Cauchy-Schwarz inequality. It immediately follows that

$$\sup_{x \in \mathcal{L}^2, \|x\|=1} \left| \bar{\pi} [x] - \frac{1}{1+r_0} E (x) \right| = \left\| d^* - \frac{1}{1+r_0} \right\| = \delta_2 [p, p^*]. \quad (24)$$

The measure $\delta_2 [p, p^*]$ can thus be interpreted as the supremum over all *possible* pricing errors arising from using a naive risk-neutral valuation approach. The inequality in equation (23) is satisfied with equality for $x^* \equiv \frac{(d^* - 1/(1+r_0))}{\|d^* - 1/(1+r_0)\|}$. Thus, if $x^* \in \mathcal{M}$ the result in (24) would hold for $x \in \mathcal{M}$. Without further assumptions on d^* , however, there is no way to determine whether or not $x^* \in \mathcal{M}$. Equation (24) does, however, hold for all $d^* \in D^*$; we can therefore state the following result:

$$\sup_{x \in \mathcal{M}, \|x\|=1} \left| \pi [x] - \frac{1}{1+r_0} E (x) \right| \leq \min_{d' \in D^*} \left\| d' - \frac{1}{1+r_0} \right\| \leq \left\| d^* - \frac{1}{1+r_0} \right\|. \quad (25)$$

That is, the supremum of the *observed* pricing error, arising from using risk-neutral evaluation, provides a lower bound for the distance between the risk-neutral and the objective probability measures.

3 Examples and Economic Implications

In the examples below, we first use a finite-state-space case to illustrate the preceding discussion about measuring differences between risk-neutral and objective probability measures,

and then show the connection of this metric with some previous studies.

CASE 1 Consider an economy in which the set Ω contains S states of nature, each with a positive probability of occurrence. Asset n 's payoff and price are respectively $x_n \in \mathfrak{R}^S$ and $q_n \in \mathfrak{R}$. Here, the s -th entry $x_{n,s}$ stands for the payoff to asset n if state s occurs in the future. Let X be the $S \times N$ payoff matrix whose n -th column equals x_n , and $q = (q_1, \dots, q_N)'$. Then, there is no arbitrage if and only if there exists a discount factor vector $d^* \in \mathfrak{R}^S$ such that $d^* \gg 0$ and $\pi[x] = \sum_{s=1}^S x_s d_s^* p_s$ for every $x \in \mathcal{M}$, where $d^* \gg 0$ means “ $d_s^* > 0$ for every component s ,” and $\pi[x]$ and \mathcal{M} are as defined before. In this case, letting $p_s^* \equiv \frac{d_s^* p_s}{\sum_{s=1}^S d_s^* p_s}$ gives

$$\pi[x] = \frac{1}{1+r_0} \sum_{s=1}^S x_s p_s^*. \quad (26)$$

Therefore, this p^* is the risk-neutral probability measure corresponding to d^* . The absolute distance between this p^* and the objective measure p is:

$$\delta[p, p^*] = \sum_{s=1}^S |p_s^* - p_s| \quad (27)$$

$$= (1+r_0) \sum_{s=1}^S \left| d_s^* - \frac{1}{1+r_0} \right| p_s. \quad (28)$$

In a true risk-neutral world, the state- s price would be $d_s^0 \equiv \frac{1}{1+r_0}$, for every s . Thus, $\delta_1[p, p^*] = \sum_{s=1}^S \left| d_s^* - \frac{1}{1+r_0} \right| p_s$ reflects the average absolute state-price difference across the states and between d^* and d^0 . The mean-square metric $\delta_2[p, p^*]$ is similarly adapted. The closer the observed economy is to true risk-neutral valuation, the smaller the distance between the risk-neutral measure and its objective counterpart according to either $\delta_1[p, p^*]$ or $\delta_2[p, p^*]$.

CASE 2 The discussion can also be applied to give the mean-square distance measures in Chen and Knez (1995) and Hansen and Jagannathan (1997) a different meaning. Take any $d_1, d_2 \in \mathcal{L}^2$ such that $d_1 > 0$ and $d_2 > 0$ almost surely. Then, let p_1 and p_2 be, respectively, the probability measures defined by the Radon-Nikodym derivatives $\frac{d_1}{E(d_1)}$ and $\frac{d_2}{E(d_2)}$. Both p_1 and p_2 are equivalent to p , and to one another, because d_1 and d_2 are both positive almost surely. Following the same sequence of steps as given above, we have the absolute distance

between p_1 and p_2 :

$$\begin{aligned}
\delta[p_1, p_2] &\equiv \sum_{i=1,2,\dots} |p_1(A_i) - p_2(A_i)| \\
&= E \left| \frac{d_1}{E(d_1)} - \frac{d_2}{E(d_2)} \right| \\
&\leq \frac{1}{E(d_1)} \|d_1 - d_2\|, \tag{29}
\end{aligned}$$

provided $E(d_1) = E(d_2)$. Thus, the mean-square distance metrics studied in Chen and Knez (1995) and Hansen and Jagannathan (1997) also capture the closeness between the relevant stochastic discount factors.

3.1 Relation with Hansen-Jagannathan Bounds

If any asset-pricing model is to explain the observed structure of asset returns, it is necessary, according to Hansen and Jagannathan (1991), that the volatility of the stochastic discount factor satisfy the following relation:

$$\sigma [d^*] \geq \left((E(q) - E(d^*) E(X))' \Sigma_X^{-1} (E(q) - E(d^*) E(X)) \right)^{1/2}, \tag{30}$$

where X and q are the N vector of included payoffs and prices respectively, and Σ_X^{-1} is the variance-covariance matrix of X . Multiplying by $(1 + r_0)$ on both sides of (30), we achieve the bound on the distance between the risk-neutral and the objective probability measures, as in:

$$(1 + r_0) \delta_2 [p, p^*] = (1 + r_0) \sigma [d^*] \tag{31}$$

$$\geq \left((E(q) - E(d^*) E(X))' \Sigma_X^{-1} (E(q) - E(d^*) E(X)) \right)^{1/2}. \tag{32}$$

An alternative interpretation of the Hansen-Jagannathan bound is thus that it is the minimum distance between the risk-neutral and the objective probability measures, which is compatible with observed prices and returns.

3.2 Equity Premium and the Distance Measure

Failure of the representative agent models to reconcile the per-capita consumption growth and stock return correlation is well-documented in the equity premium puzzle literature (Mehra and Prescott (1985, 2003) and Campbell (2003)). This framework nevertheless serves as a useful benchmark to illustrate some of the main ideas of this paper. The first order conditions of the model imply that the time- t asset price at time t must satisfy $\pi_t[x] = E_t\left(\beta \frac{u'[c_{t+1}]}{u'[c_t]} x_{t+1}\right)$, where c_t is per-capita consumption at time t , β is a subjective discount factor, and $E_t(\cdot)$ is the expectations operator conditioning on the information available at time t . The stochastic discount factor in this endowment economy is,

$$d_{t+1}^* = \beta \frac{u'[c_{t+1}]}{u'[c_t]}. \quad (33)$$

Suppose $u[c] = \frac{1}{1-\gamma} c^{1-\gamma}$, and the objective density, p , for per-capita consumption growth is gaussian, that is, $Y_{t+1} \equiv \log\left(\frac{c_{t+1}}{c_t}\right) = \log(\mu_c) - \frac{1}{2}\sigma_c^2 + \epsilon_{c,t+1}$, where $\epsilon_{c,t+1} \sim \mathcal{N}(0, \sigma_c^2)$. Under these assumptions, $d_{t+1}^*[Y] = \beta e^{-\gamma Y_{t+1}}$, and therefore,

$$\text{Var}_t(d_{t+1}^*) = \beta^2 \mu_c^{-2\gamma} e^{(1+\gamma)\gamma\sigma_c^2} (e^{\gamma^2\sigma_c^2} - 1), \quad (34)$$

and, thus, the distance metric is:

$$\delta_2[p, p^*] = \sigma[d^*] = \beta \mu_c^{-\gamma} e^{(1+\gamma)\gamma\sigma_c^2/2} \sqrt{e^{\gamma^2\sigma_c^2} - 1}. \quad (35)$$

Furthermore, the risk-free return, r_0 , satisfies $\frac{1}{1+r_0} = E_t(d_{t+1}^*) = \beta \mu_c^{-\gamma} e^{(1+\gamma)\gamma\sigma_c^2/2}$ and is endogenous. Therefore, consider the transformed distance measure:

$$(1+r_0) \delta_2[p, p^*] = \sqrt{e^{\gamma^2\sigma_c^2} - 1}, \quad (36)$$

which shows that the source of the dichotomy between the risk-neutral and objective probability measure is related to both the level of risk aversion, γ , and the variance of the consumption-growth.

Figure 1 plots the distance between p and p^* as a function of γ , for several different values of σ_c . It is noteworthy that varying σ_c has a substantial impact on the distance measure in (36). Thus, choosing a suitably volatile consumption-growth proxy can lead to

a more reasonable estimate of γ and alleviate the equity premium puzzle, as also shown by Ait-Sahalia, Parker, and Yogo (2004).

In our gaussian setting, it is easy to see that the risk-neutral density, p^* , for Y_{t+1} inherits the gaussian structure of p , but with a lower mean and the same variance σ_c^2 . Here a more volatile stochastic discount factor produces more pronounced downward consumption growth fears in the risk-neutral mean while leaving all remaining higher-moments intact. Thus, in this economy, the source of the distance is risk-adjustment associated with the first-moment of the pricing distribution. Although not shown here to preserve focus, for a broader set of stochastic discount factors and/or p -distributions outside of the gaussian class (say, in the family of Carr, Geman, Madan, and Yor (2002)), a more volatile d^* may result in p^* -distribution higher-moments that differ from the objective distribution counterparts. Such asset pricing models are capable of generating left-skewed and peaked p^* -distributions even though the p -distribution is symmetric. Economic models where higher volatility of the stochastic discount factor can lead to conceptually distinct p^* and p measures in the tails are desirable for explaining options on the market-portfolio.

4 Measuring $\delta_2 [p, p^*]$ Across Time

Having established a measure of the distance between the risk-neutral and the objective probability measures, we now look for a way to estimate this distance. The metric we have been considering is: $\delta_2 [p, p^*] = \left\| d^* - \frac{1}{1+r_0} \right\| = \sigma [d^*]$. Unfortunately, as is evident from equation (25), we cannot in general obtain an estimate of $\delta_2 [p, p^*]$, but only a lower bound. Estimators of these lower bounds are, however, well-known and described in Hansen and Jagannathan (1991).

Under the assumption that d^* is an admissible discount factor, in the sense that $\pi [x] = E(xd^*)$ for all $x \in \mathcal{M}$, we achieve the bound on $\sigma [d^*]$ described in equation (30). We can strengthen the bound further by restricting d^* to be strictly positive. The procedure for finding this bound is more complicated and detailed in Hansen and Jagannathan (1991).

Moving to our empirical application, there have been several studies recovering the risk-neutral distribution implied by options on the S&P 500 index. One frequently highlighted finding is that the risk-neutral distribution before and after the 1987 crash seems quite different. On the other hand, there has been little evidence that the *objective* distribution of the

index changed over the same period of time (see Rubinstein (1994), Jackwerth (2000) and Aït-Sahalia et al. (2001)). Thus, if these observations are indeed true, the distance between the risk-neutral and the objective probability measures must have increased after 1987 for the market-index. The general explanation of this phenomenon would be that investors' risk-preferences changed with the crash in the market and the objective distributions have become more fat-tailed (Bakshi, Kapadia, and Madan (2003)).

Recovering risk-neutral distribution from option prices requires fairly strong assumptions on price dynamics and option pricing formulas. The approaches either use simple but flexible (nonparametric) models (say, Jackwerth (2000) and Aït-Sahalia et al. (2001)), or more complex but also more parametrically restricted models (say, Engle and Rosenberg (2002)). Since it has turned out to be a tough enterprise indeed to reconcile stock and options data (see Pan (2002) and Eraker (2004) for recent attempts), it seems reasonable to ask whether we can observe the same shift in the risk-neutral distribution using index data alone, rather than relying on option prices. Aït-Sahalia et al. (2001) do show how to estimate the risk-neutral distribution using returns data. Their method, however, restricts them to using a univariate diffusion model for the stock price dynamics.

In light of the above discussion, we therefore appeal to the measure of distance between the risk-neutral and the objective distribution derived in this paper to add some new insight to this question. The advantage of using our closeness measure, apart from the fact that no option data is required, is that it is completely model-free, in that we need to make no assumptions on returns distributions or the investors' preferences. The downside is that we can not actually get an estimate of the distance between the risk-neutral and objective probability measures, but only an estimate of the lower bound on this distance. The results should still be indicative, however.

In order to get an estimate of the lower bound of $\delta_2 [p, p^*]$, using the S&P 500 index, we use the bound defined by equation (32). While using the positivity constraint on the discount factor could further strengthen this bound, previous results indicate that the effect for stock data is small (see in particular, Hansen and Jagannathan (1991)). Monthly returns for the S&P 500 index and 30-day Treasury bonds are used, where the 30-day Treasury bond is considered as a proxy for the risk-free rate. The data span from 1926 to 2003, for a total of 936 observations, and is obtained from the CRSP database.

To infer the lower bound displayed in equation (32), we calculate the sample mean of the

realized S&P 500 and Treasury bond returns for ten year periods and use these as estimates of the expected return and the risk-free rate, respectively. The sample standard deviation of the S&P 500 returns are likewise calculated for the same ten year periods. Using the whole sample of 78 years, an average bound over the sample period is also calculated. Since the risk-free rate is not constant over the whole sample period, we again multiply $\delta_2 [p, p^*]$ by the risk-free rate to estimates that are comparable across time (as we did when comparing across different risk aversion parameters in Section 3.2). Since the data is on a monthly basis, the resulting estimates give bounds on the monthly volatility of the discount factor, and are thus not directly comparable to the results in Figure 1, which are based on annual parameter values.

The results are shown in Figure 2. The top part of the figure shows estimates of the lower bound on $(1 + r_0) \delta_2$ for each ten year period from 1926 and onward. The last estimate is calculated using eight years of data, between 1996 and 2003. We see that the bound moves a fair amount over time and that there is some indication that the bound became greater for the years after 1987, but the evidence is fairly weak.

In the middle graph, we explicitly take into consideration that there might have been a large change in 1987. Ten year periods are again used from 1926 to 1975, but then the next bound is based on eleven years of data between 1976 and 1986. The data from 1987 is excluded, such that the subsequent bound is based on data from 1988-97. The last bound is based on six years of data, between 1998 and 2003. The results are now more convincing, with a big jump in the bound after the crash in 1987.

The bottom graph shows estimates of the lower bounds using ten year overlapping estimates, from 1926 and onward. The last estimate is based on data for the years 1994-2003. These estimates tell the same story as those in the middle graph. There is also some evidence, both in the middle and bottom graphs, that the lower bound has decreased in the last few years, though the big drop noted in the middle graph should be interpreted with caution since it is based on only six years of data.

Since we can only obtain estimates of a lower *bound* on the distance between the risk-neutral and objective probability measures, rather than an estimate of the actual distance, we cannot say whether what we observe as shifts in the estimates, actually represents true shifts in the actual distance, or not. Our results, however, do indicate that there might have been a substantial increase in the distance between the risk-neutral and objective

probability measures after 1987.

The possibility that the distance between the risk-neutral and objective probability measures might have decreased again over the last few years is not something that has been observed before. If this is indeed true, it might indicate that the crash fears, that most likely caused a larger gap between the risk-neutral and objective distribution after 1987, have evaporated over time.

There are a few other things to note about the results in Figure 2. First, we cannot see a similar effect of the crash in 1929 as we do for the crash in 1987. Indeed, around 1930 we observe a very small lower bound. Second, there seems to be great changes in the lower bound over time, with fairly long periods of small and large bounds alternating.

5 Conclusions

In this paper, we derive a theoretically sound and intuitively appealing measure of the distance between the risk-neutral and the objective probability. The distance measure, which is based on the mean-square norm for square integrable random variables, turns out to be equivalent to the volatility of the stochastic discount factor. This follows directly from the Radon-Nikodym derivative linking the objective and risk-neutral measures, and its link to the stochastic discount factor.

The search in the equity premium literature for sufficiently volatile discount factors can thus be equivalently interpreted as a search for a risk-neutral distribution that is sufficiently different from its objective counterpart. Another implication of our results is that it provides an alternative interpretation of the Hansen-Jagannathan bounds. Hansen and Jagannathan (1991) show how to infer lower bounds on the volatility of the stochastic discount factor, given observed asset prices. Our study shows that we can also interpret these bounds as lower bounds on the distance between the risk-neutral and the objective probability measures.

As in the literature dealing with the volatility of the stochastic discount factor, we cannot directly estimate the distance measure, without any further assumptions about returns distributions and investor preferences. However, using the Hansen-Jagannathan bounds, we can provide lower bounds for our distance measure. Using this technique, we consider monthly S&P 500 index returns from 1926 to 2003 and estimate, over 10

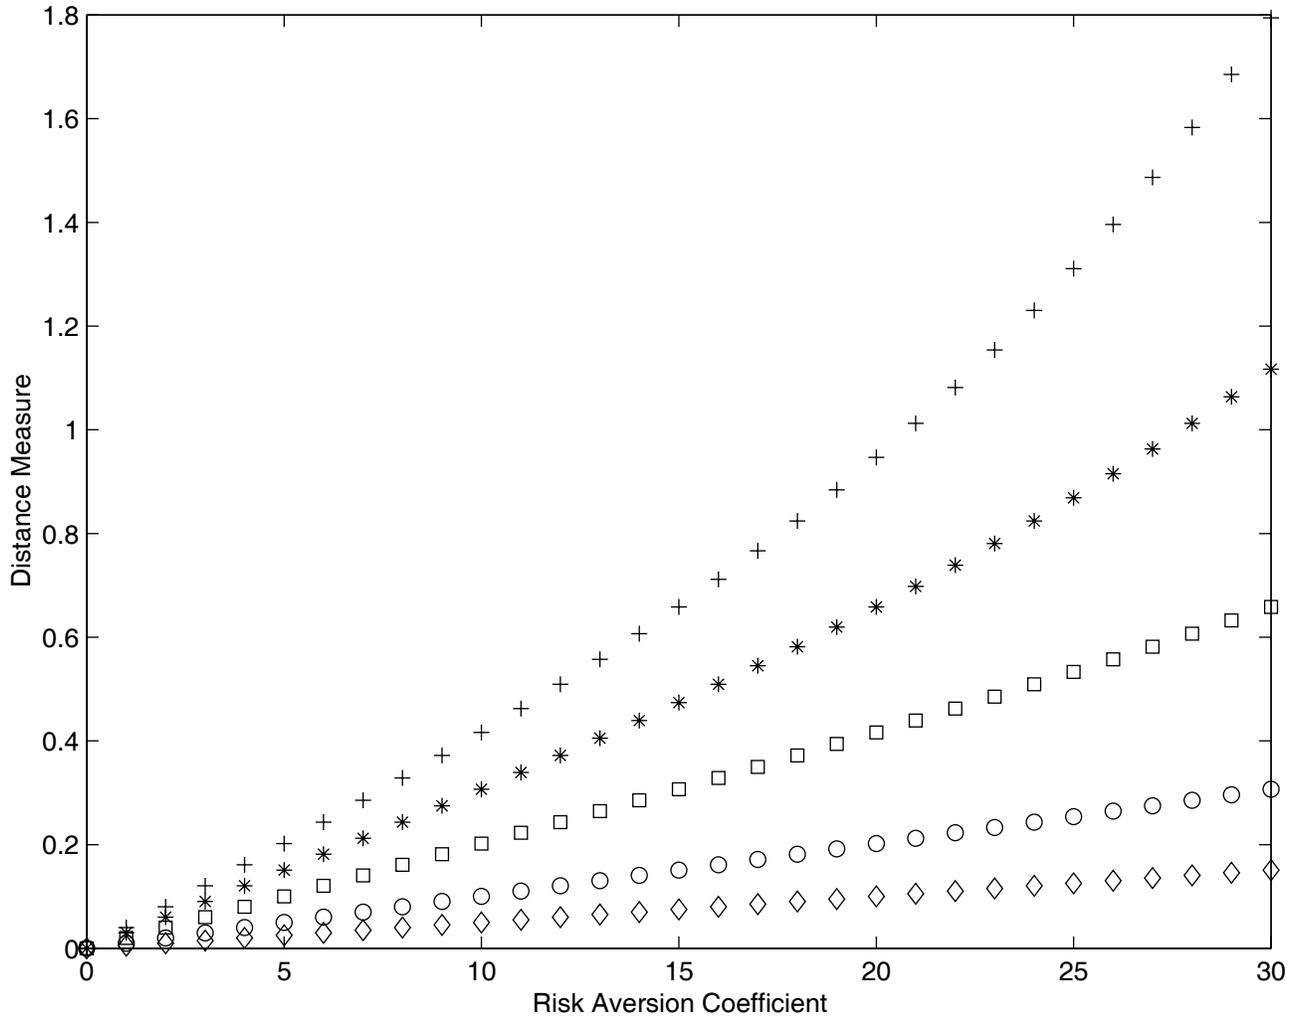
year periods, lower bounds for the distance between the risk-neutral and the objective probability measure. We find evidence of a relatively large change in this bound after the crash in 1987, which is in line with the existing literature that has found similar evidence using S&P 500 index options. There is some evidence that the bound has again decreased in recent years, implying that investors might have overcome some of the crash fears that most likely increased the distance between the risk-neutral and the objective probability measures after 1987. This paper proposes a measure that characterizes the distance between the risk-neutral and the objective probability measures for any candidate asset pricing model, and then studies its behavior in the U.S. equity market.

References

- Aït-Sahalia, Y., J. Parker, and M. Yogo, 2004. Luxury goods and the equity premium, *Journal of Finance* (forthcoming).
- Aït-Sahalia, Y., Y. Wang, and F. Yared, 2001. Do option markets correctly price the probabilities of movement of the underlying asset?, *Journal of Econometrics* 102, 67-110.
- Bates, D., 2000. Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics* 94, 181-238.
- Bliss, R., and N. Panigirtzoglou, 2003. Option-Implied risk aversion estimates, *Journal of Finance* (forthcoming).
- Bakshi, G., N. Kapadia, and D. Madan, 2003. Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16 (1), 101-143.
- Breeden, D., and R. Litzenberger, 1978. Prices of state-contingent claims implicit in option prices, *Journal of Business* 51, 621-651.
- Campbell, J., 2003. Consumption-based asset pricing, in: Constantinides, G., Harris M., and Stulz R. eds., *Handbook of the Economics of Finance*, Vol. 1B (North-Holland, Amsterdam) 803-888.
- Carr, P., H. Geman, D. Madan, and M. Yor, 2002, The fine structure of asset returns: an empirical investigation, *Journal of Business* 75, No. 2, 305-332.
- Chen, Z., and P. Knez, 1995. Measurement of market integration and arbitrage, *Review of Financial Studies* 8, 287-325.
- Cochrane, J., 2001. *Asset Pricing*, Princeton University Press, Princeton.
- Engle, R., and J. Rosenberg, 2002. Empirical pricing kernels, *Journal of Financial Economics* 64, 341-372.
- Eraker, B., 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option Prices, *Journal of Finance* 59, 1367-1403.
- Eraker, B., M. Johannes, and N. Polson, 2003, The impact of jumps in volatility and returns, *Journal of Finance* 58, No. 3, 1269-1300.

- Hansen, L., and R. Jagannathan, 1991. Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225-262.
- Hansen, L., and R. Jagannathan, 1997. Assessing specification errors in stochastic discount factor models, *Journal of Finance* 52, 557-590.
- Harrison, J., and D. Kreps, 1979. Martingales and arbitrage in multiperiod securities markets, *Journal of Economic Theory* 20, 381-408.
- Jackwerth, J., 2000. Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* 13, 433-451.
- Jackwerth, J., and M. Rubinstein, 1996. Recovering probability distribution from option prices, *Journal of Finance* 51, 1611-1631.
- Mehra, R., and E. Prescott, 1985. The equity premium: a puzzle, *Journal of Monetary Economy* 15, 145-161.
- Mehra, R., and E. Prescott, 2003. The equity premium in retrospect, in: Constantinides, G., Harris M., and Stulz R. eds., *Handbook of the Economics of Finance*, Vol. 1B (North-Holland, Amsterdam) 889-938.
- Pan, J., 2002. The jump-risk premia implicit in options: evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3-50.
- Ross, S., 1978. A simple approach to the valuation of risky streams, *Journal of Business* 51, 453-475.
- Rubinstein, M., 1994. Implied binomial trees, *Journal of Finance* 49, 771-818.

FIGURE 1: Distance Measure in Power Utility Economies



Note to Figure 1: This figure plots the distance measure between p and p^* in power utility economies. We assume that the objective density, p , for per-capita consumption growth is gaussian, that is, $Y_{t+1} \equiv \log\left(\frac{c_{t+1}}{c_t}\right) = \log(\mu_c) - \frac{1}{2}\sigma_c^2 + \epsilon_{c,t+1}$, where $\epsilon_{c,t+1} \sim \mathcal{N}(0, \sigma_c^2)$. Thus, $d_{t+1}^*[Y] = \beta e^{-\gamma Y_{t+1}}$, and hence $(1 + r_0) \delta_2 [p, p^*] = \sqrt{e^{\gamma^2 \sigma_c^2} - 1}$. The \diamond -curve, \circ -curve, \square -curve, the \star -curve, and $+$ -curve respectively represent $(1 + r_0) \delta_2 [p, p^*]$ when σ_c takes a value of 0.5%, 1%, 2%, 3%, and 4%.

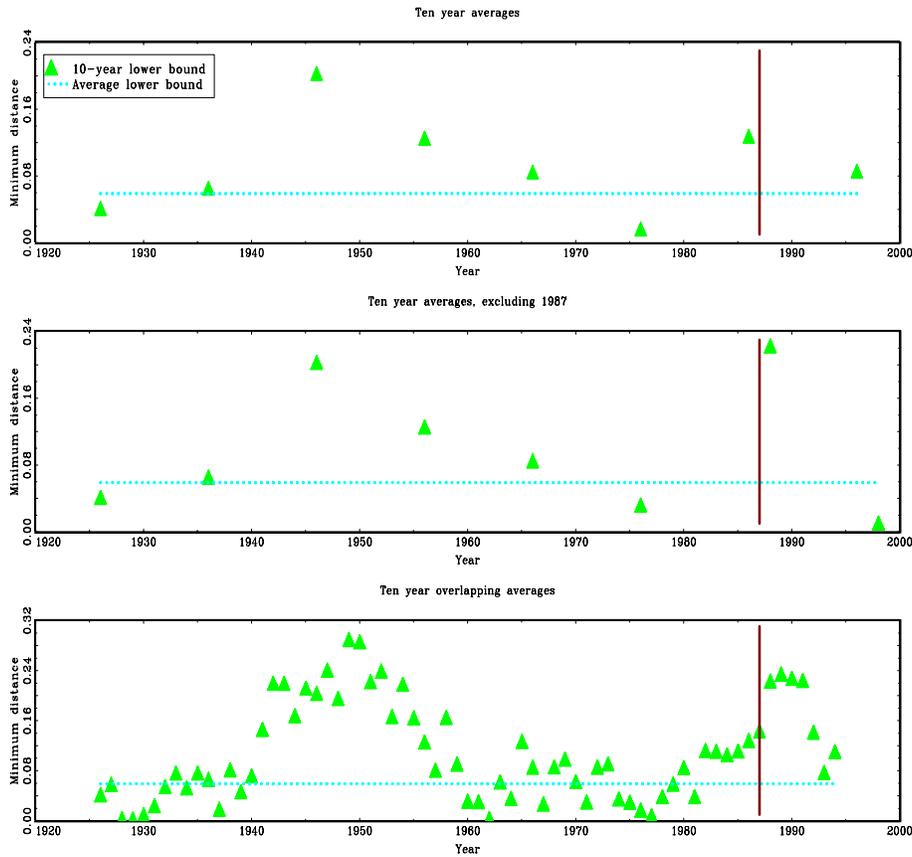


Figure 2: Dynamic variation in $(1 + r_0) \delta_2 [p, p^*]$.

Estimated bounds of the minimum distance between the risk-neutral and the objective distribution, using monthly returns on the S&P 500 and 30 day Treasury bonds. Each estimate is calculated using returns over the next ten years for a total of 120 observations. That is, the estimate plotted for 1926 uses data for 1926-1935, and so forth. The dotted line in each graph is the estimate of the average lower bound during the whole sample period, and the vertical line indicates the crash year, 1987. Three different estimation schemes are considered. (1) The top graph pays no special attention to the crash in 1987 and shows estimates for non-overlapping ten year periods beginning in 1926. (2) The middle graph is based on the hypothesis that investor behavior changed after the crash in 1987. As in the top graph, non-overlapping ten year averages are calculated, starting in 1926, but to get an estimate of the bounds right before and after 1987, an eleven year period from 1976 to 1986 is used, while the next bound is based on data from 1988 to 1997. (3) The bottom graph shows estimates for ten year *overlapping* periods, starting in 1926. In the top two graphs, the last estimate is based on less than ten years of data; eight years in the top graph and six years in the middle graph.