

Asset Pricing Through the Lens of the Hansen-Jagannathan Bound

Christopher Otrok and B. Ravikumar

Stochastic discount factor (SDF) models are the dominant framework for modern asset pricing. The Hansen-Jagannathan bound is a characterization of the admissible set of SDFs, given a vector of asset returns. The admissible set provides (i) a test of the asset-pricing model and (ii) information on how to modify the SDF to be consistent with asset returns, neither of which requires solving the model. In this article we use the Hansen-Jagannathan bound to examine asset-pricing implications and to test specific asset-pricing models using bootstrap experiments. (JEL G1, C15, E44)

Federal Reserve Bank of St. Louis *Review*, Third Quarter 2020, 102(3), pp. 255-69.
<https://doi.org/10.20955/r.102.255-69>

1 INTRODUCTION

An asset-pricing model is typically defined by its stochastic discount factor (SDF). For instance, Mehra and Prescott (1985) used constant-relative-risk-aversion (CRRA) preferences and the SDF in their model was a function of consumption growth. The validity of an SDF is determined by its ability to match the observed asset returns. An early test of an asset-pricing model with CRRA preferences was the Hansen and Singleton (1982) J-test. For U.S. stock and bond returns data, this test typically rejects the model. The J-test tells us whether or not an asset-pricing model has statistically significant pricing errors. It does not provide information on how to modify the SDF to improve the fit. Hansen and Jagannathan (1991) derive a volatility bound (HJ bound) that is based on necessary conditions that an asset-pricing model must satisfy. The HJ bound characterizes the admissible set of SDFs that is consistent with the observed asset returns.

The HJ bound exploits two conditions: (i) the intertemporal Euler equation that connects the price of an asset to the covariance of the asset's payoff with the SDF and (ii) the implication from linear pricing that the SDF be a linear function of payoffs. The asset-pricing model is said to be consistent with the data if the volatility of the proposed SDF (evaluated at the mean SDF) is greater than the volatility implied by the HJ bound. The HJ bound is a lower bound

Christopher Otrok is the Sam B. Cook Professor of Economics at the University of Missouri–Columbia and a research fellow at the Federal Reserve Bank of St. Louis. B. Ravikumar is an economist, a senior vice president, and the deputy director of research at the Federal Reserve Bank of St. Louis.

© 2020, Federal Reserve Bank of St. Louis. The views expressed in this article are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

and, hence, is a necessary but not sufficient condition that an asset-pricing model must satisfy. In other words, the HJ bound provides a “test” of an asset-pricing model based solely on necessary conditions implied by the model.

The HJ bound approach in a sense works backward: Instead of writing down a model, solving it, and then testing it, the HJ bound asks what a valid SDF should look like in the mean-variance space. The HJ bound approach has several advantages. First, the bound is model-free; that is, it is constructed using only observed asset returns. Second, one does not need to solve the nonlinear asset-pricing model. Specifically, there is no need to find a partial equilibrium or a general equilibrium solution to the model. Third, there is no limit on the number of assets used in the construction of the bound. Fourth, the bound is informative on how to modify the SDF in order to be consistent with the data.

In this article we provide a derivation of the HJ bound and then apply the bound to examine a few popular SDFs. The results provide an illustration of the equity premium puzzle. We then check the robustness of the resolutions of the puzzle with a bootstrap experiment. Our bootstrap results indicate that minor variations in asset return moments and consumption moments can yield large variations in the distance between an SDF’s volatility and the HJ bound. We conclude with some implications for business cycle models.

2 THE HANSEN-JAGANNATHAN BOUND

For frictionless asset-pricing models, Hansen and Jagannathan (1991) showed that the volatility of the SDF that satisfies the representative consumer’s Euler equation must exceed a lower bound that is a function of only asset returns. The derivation of the HJ bound is presented here purely for completeness. (In the appendix, we derive the Sharpe-ratio version of the HJ Bound; see also Ljungqvist and Sargent, 2018.)

Let R denote the $n \times 1$ (gross) return vector of risky assets. Consider an SDF m that prices the n assets according to

$$E_t(R_{t+1}m_{t+1}) = \iota,$$

where E_t is the expectation operator conditional on information in period t and ι is an $n \times 1$ vector of 1s. This is the standard Euler condition, which equates the expected marginal cost and marginal benefit of delaying consumption one period. For example, in the case of time-separable preferences, m_{t+1} is the ratio of the marginal utility of future consumption to the marginal utility of current consumption. The unconditional version of the Euler equation is then

$$(1) \quad E(Rm) = \iota.$$

Note that if there is a risk-free asset, then its gross return is $\frac{1}{E(m)}$. In the absence of a risk-free asset, we cannot pin down the mean of the SDF using return data.

Suppose we compute the least-squares projection of the SDF onto the linear space spanned by a constant and contemporaneous returns. The projection is of the form

$$(2) \quad m = m_v + \varepsilon,$$

where

$$(3) \quad m_v = v + (R - E(R))' \beta,$$

$\beta \in \mathfrak{R}^n$, $v = E(m) = E(m_v)$, and ε is orthogonal to the constant as well as contemporaneous returns. This implies $E(\varepsilon) = 0$ and $E(R\varepsilon) = 0$. Together with the Euler equation (1), this implies $E(Rm) = E(Rm_v) = \iota$. Then

$$\text{var}(m) = \text{var}(m_v) + \text{var}(\varepsilon) + 2\text{cov}(m_v, \varepsilon).$$

By construction of m_v , the projection error ε is orthogonal to m_v , so $E(m_v \varepsilon) = 0$. Thus,

$$\text{var}(m) = \text{var}(m_v) + \text{var}(\varepsilon) \geq \text{var}(m_v),$$

meaning that a lower bound on the variance of a model's SDF m is the variance of m_v . To find this lower bound, we need to know $\text{var}(m_v)$.

From (3) it is easy to see that $\text{var}(m_v) = \beta' \Omega \beta$, where Ω is the variance-covariance matrix of asset returns. Since (2) and (3) describe a linear least-squares projection, we can estimate the projection coefficient β via OLS as $\beta = \Omega^{-1} \text{cov}(R, m)$. Rewriting $\text{cov}(R, m)$, we have $\beta = \Omega^{-1}(E(Rm) - E(m)E(R))$. Since the model implies $E(Rm) = \iota$, we can solve for β with

$$\beta = \Omega^{-1}(\iota - E(m)E(R)).$$

Thus, we can write $\text{var}(m_v) = (\iota - E(m)E(R))' \Omega^{-1} (\iota - E(m)E(R))$. In terms of standard deviations, we can write the lower bound as

$$(4) \quad \text{std}(m) \geq \left\{ (\iota - E(m)E(R))' \Omega^{-1} (\iota - E(m)E(R)) \right\}^{\frac{1}{2}}.$$

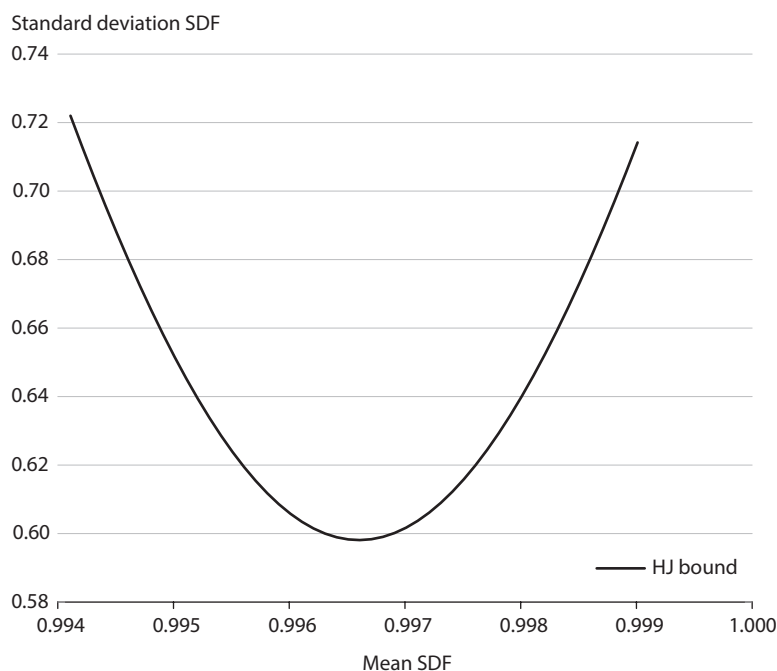
The right-hand side is the HJ bound. Note that the lower bound on the standard deviation of a model's SDF is a function of the mean of the model's SDF; so, it would seem like the lower bound depends on the model. However, we can generate a lower-bound *frontier* by picking different means. It is easy to see that the bound is a quadratic function of the mean SDF. A necessary condition for an SDF with mean $E(m)$ to be consistent with asset-return data is that it satisfies the inequality (4).

Computing the HJ bound frontier is straightforward. First, we calculate the sample mean of gross returns to use as a proxy for $E(R)$. Second, we calculate the variance-covariance matrix of the gross returns. Third, we choose a set of values for $E(m)$. For each value we compute the right-hand side of (4) to trace out a bound frontier.

Figure 1 illustrates the HJ bound using two asset returns from 1959:Q2 to 2019:Q2: the return on a 3-month Treasury bill and the return on the S&P 500. Both returns are transformed into real returns using the price deflator for personal consumption expenditures. (We use this deflator because when we conduct model evaluations later, we will be using personal

Figure 1

HJ Bound Frontier



NOTE: The figure depicts the set of admissible SDFs in mean-standard deviation space implied by stock and bond returns from 1959:Q2-2019:Q2.

consumption data.) The horizontal axis is $E(m)$, and the vertical axis is the HJ bound. The frontier is U-shaped, implying that SDFs with means far from the one associated with the least volatility will need to have higher volatility to satisfy the bound.¹

3 EQUITY PREMIUM PUZZLE


In this section, we use the HJ bound to illustrate the equity premium puzzle. To “test” a model using the HJ bound, we need the SDF implied by the model. As an example, suppose we want to check whether the Mehra and Prescott (1985) model is consistent with asset-return data. The preferences in their model are described by

$$E_0 \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma}, \quad \sigma > 0,$$

where E_0 is the conditional expectation given information at time 0, c_t is the representative agent’s consumption at time t , $\beta \in (0,1)$ is the subjective discount factor, and σ is the coefficient of relative risk aversion. (The preferences are assumed to be logarithmic when $\sigma = 1$.) The SDF for these preferences is given by

Figure 2

Mehra and Prescott (1985)



NOTE: Sample: 1959:Q2-2019:Q2. The HJ bound frontier in this figure is the same as in Figure 1. It also has the means and standard deviations of the Mehra-Prescott SDF for the range of the risk-aversion parameter (σ), from 1 to 10.

$$m_{t+1} = \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\sigma}.$$


We can compute the time series of m_{t+1} using consumption data and parameter values for β and σ . Different values of β and σ imply different means, $E(m)$, and different volatilities, $\text{std}(m)$. The question is whether there are any empirically plausible β and σ such that the pair $(E(m), \text{std}(m))$ is inside the frontier. That is, given the mean of the model's SDF, the test is whether $\text{std}(m)$ satisfies the bound in (4).

Figure 2 plots the same bound as in Figure 1 for the stock and bonds returns data. Figure 2 also plots the pairs $(E(m), \text{std}(m))$ using quarterly nondurables and service consumption data from 1959:Q2 to 2019:Q2, for $\beta = 0.99$ and values of σ from 1 to 10. These values of β and σ are in the range investigated by Mehra and Prescott. The volatility for $\sigma = 1$ is the right most "x." As risk aversion is increased, the x's move to the left, but the increases in volatility are small.

For no value of σ is the bound satisfied. In fact, the volatilities of the SDF are far below the bound. We conclude that the model with this parameterization is rejected. A natural question is whether or not there exists a parameterization of the model that satisfies the bound.

Figure 3

High Risk Aversion in Mehra and Prescott (1985)



NOTE: Sample: 1959:Q2-2019:Q2. The figure illustrates the means and standard deviations of the Mehra-Prescott SDF for a wider range of the risk-aversion parameter (σ), relative to Figure 2, from 1 to 460.

To do so we increase risk aversion and find that the Mehra-Prescott model generates enough volatility to satisfy the bound when $\sigma = 460$; see Figure 3.

The high value of σ is unreasonable for two reasons. First, it implies an extreme aversion to risk. Second, it implies a high risk-free rate of 36 percent annually.² Figure 3 thus demonstrates the risk-free rate puzzle as well: The level of risk aversion that matches the observed equity premium comes at the cost of unreasonable values for the risk-free rate.

In sum, the Mehra-Prescott model of asset pricing is rejected for reasonable parameterizations of risk aversion. One needs implausibly high values of risk aversion to generate sufficient volatility to satisfy the bound. Moving forward, we need to find an SDF that generates higher volatility without high risk aversion.

Note that the above evaluation of the Mehra-Prescott model did not require us to solve the model or compute *equilibrium* asset returns. The test involved merely checking whether a necessary implication of the model was satisfied. We learned the same lessons that emerge from a full solution of the model. Another alternative to testing models using just the first-order conditions would be to estimate the Euler equation via GMM (generalized method of moments) as in Hansen and Singleton (1982) and then apply a J-test to the overidentifying restrictions. As is well known, this would lead to a statistical rejection of the Mehra-Prescott model. It

would not, however, provide any guidance as to why the model was rejected and what to do to fix the model.

4 RESOLUTIONS

A similar procedure can be applied to two other popular asset-pricing models. Both are based on relaxing separability in the utility function: in one case state separability and in the other case time separability. Both add just one parameter to the Mehra-Prescott model, and both increase the volatility of the SDF.

Epstein and Zin (1991) and Weil (1989) generalize the time-separable preferences to allow for an independent parameterization of attitudes toward risk and intertemporal substitution. Following Weil (1989), these state-nonseparable preferences have a recursive representation:

$$V_t = U[c_t, E_t V_{t+1}],$$

where V is a von-Neumann-Morgenstern utility index and

$$U[c, V] = \frac{\left\{ (1-\beta)c^{1-\rho} + \beta \left[1 + (1-\beta)(1-\sigma)V \right]^{\left(\frac{1-\rho}{1-\sigma}\right)} \right\}^{\left(\frac{1-\sigma}{1-\rho}\right)} - 1}{(1-\beta)(1-\sigma)}.$$

The elasticity of intertemporal substitution is $1/\rho$, and σ is the coefficient of relative risk aversion. As shown by Weil (1989), the SDF for these preferences simplifies to

$$\left[\beta \left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \right]^{\left(\frac{1-\sigma}{1-\rho}\right)} [R_{t+1}]^{\left(\frac{1-\sigma}{1-\rho}\right)-1},$$

where R_{t+1} is the return on the market portfolio.


Constantinides (1990) models consumers as habitual, in that levels of consumption in adjacent periods are complementary. That is, the time-nonseparable preferences of consumers (in a discrete-time, one-lag version of Constantinides, 1990) are given by

$$U_0 = E_0 \sum_{t=0}^{\infty} \beta^t \frac{[c_t - \delta c_{t-1}]^{1-\sigma}}{1-\sigma},$$

where $\delta > 0$. The representative agent's SDF is given by

$$m_{t+1} = \beta \frac{(c_{t+1} - \delta c_t)^{-\sigma} + \beta \delta E_{t+1} (c_{t+2} - \delta c_{t+1})^{-\sigma}}{(c_t - \delta c_{t-1})^{-\sigma} + \beta \delta E_t (c_{t+1} - \delta c_t)^{-\sigma}}.$$

Figure 4 plots the bound and the SDF volatilities for the two models. Again, the HJ bound is the solid curve, as in Figure 1; the “x” represents the habit SDF, while the “o” the Epstein-Zin SDF. For state-nonseparable preferences, the parameters are $\beta = 0.99$, $\rho = 0.9$, and $\sigma = 1.7$. For time-nonseparable preferences, the parameters are $\beta = 0.99$, $\delta = 0.8$, and $\sigma = 1.61$.³ Both

Figure 4**Epstein-Zin and Habit-Formation Models**

NOTE: Sample: 1959:Q2-2019:Q2. The figure illustrates the means and standard deviations of the SDFs for the Epstein-Zin and habit-formation preferences that satisfy the HJ bound. The HJ bound is the same as in Figure 1.

models satisfy the HJ bound by increasing the volatility of the SDF. In the case of the Epstein-Zin model, wealth, which is volatile, is part of the SDF. In the case of the time-nonseparable models, consumption growth is operated on by a difference operator, which increases volatility when raised to moderate powers.

4.1 Are the Resolutions Robust?

Our model evaluation shows that both Epstein-Zin and habit-formation models satisfy the HJ bound for apparently reasonable parameter values.⁴ The evaluation was simple: It compared just two points—the volatility of the SDF and the HJ bound at the mean of the same SDF. The evaluation does not account for sampling variability in (i) asset-return data and (ii) consumption data. The sampling variability in (i) and (ii) might affect our inference on the model since the HJ bound is affected by (i) and the SDF is affected by (ii).

We now conduct a bootstrap experiment to take into account the two sampling variabilities and check whether the resolutions are robust to changes in the data sample. We adopt a variant of the bootstrap procedure in Otrok, Ravikumar, and Whiteman (2004): They first find parameters of the asset-pricing models that satisfy the HJ bound for the whole post-WWII sample. Using these parameters, they then show that the models do not satisfy the bound for subsamples.

Our bootstrap experiment here, however, investigates how the HJ bound and the volatility of the SDF vary across artificial samples drawn from the full data set (1959:Q2-2019:Q2). To do this we compute the time series for the representative agent’s SDF for the two successful asset-pricing models in Section 3 using consumption growth data. We then use a bootstrap procedure to sample a vector of asset returns and the SDF. We bootstrap the entire vector—consumption, equity return, and bond return—so that the observed correlation properties between the two returns and the SDF are maintained in our experiment.

The bootstrap procedure is as follows:

- (i) Use the parameters from Section 3 and observed consumption growth data to get time series for the SDFs of the models. (The parameters are $\beta = 0.99$, $\rho = 0.9$, and $\sigma = 1.7$ for the Epstein-Zin SDF and $\beta = 0.99$, $\delta = 0.8$, and $\sigma = 1.61$ for the habit SDF.)
- (ii) Draw (with replacement) a time series of length 241 from the joint “empirical” distribution of the SDFs, equity returns, and T-bill returns. That is, for each period we draw a 3-tuple (SDF, R^{equity} , $R^{\text{T-bill}}$).
- (iii) Calculate the mean and volatility of the SDF.
- (iv) Calculate the HJ bound using the time series for equity and T-bill returns at the mean SDF.
- (v) Repeat steps ii-iv 1,000 times.

Figure 5 is a scatter plot of the distance between the HJ bound and SDF volatility, calculated as the SDF volatility minus the HJ bound, for each of the 1,000 bootstrap simulations. Panel A plots the habit-formation model, while Panel B plots the Epstein-Zin model. The striking feature of these figures is that the distance is almost always negative, implying that the models miss the bound in most simulations. In fact, the habit model misses in 96 percent of the simulations and the Epstein-Zin model misses in 95 percent of the simulations.


Burnside (1994) casts the distance to the bound in a GMM framework. He studies statistical measures of the distance between the HJ bound and SDF volatility in the *time-separable* model and argues that the over rejection is partly due to variations in the mean of the SDF. Cecchetti, Lam, and Mark (1994) also show that in the context of models with time-separable preferences and habit formation preferences, much of the variability in the distance is due to the uncertainty in estimating the mean of the SDF. To their point, even if we consider only the lowest possible bound from the bootstrap simulations, we will reject the models most of the time. The reason is that the mean of the SDF varies greatly across bootstrap samples. Since the bound itself rises rapidly for $E(m)$ different from 0.99, the distance to the bound becomes large and leads to a rejection.⁵

Statistically speaking, for the HJ bound to be a useful evaluation device, the test should not reject a true model. Specifically, suppose one uses observed consumption data to solve an asset-pricing model, that is, compute the *equilibrium* asset returns implied by the model. Then the test based on the distance between the HJ bound associated with the equilibrium returns and the volatility of the model SDF should not reject the true model. One can judge the test by simulating the true model many times and counting the number of times the the HJ bound is violated. Gregory and Smith (1992) conduct this exercise for time-separable preferences and conclude that the true model is rejected frequently.


Figure 5

Bootstrap Simulations

A. Habit-formation model



B. Epstein-Zin model



NOTE: Sample: 1959:Q2-2019:Q2. The figure plots the distance between the standard deviation of the SDF and the HJ bound evaluated at the mean of the SDF for each draw of the bootstrap. The parameters for the Epstein-Zin SDF are $\beta = 0.99$, $\rho = 0.9$, and $\sigma = 1.7$; the parameters for the habit SDF are $\beta = 0.99$, $\delta = 0.8$, and $\sigma = 1.61$.

A formal statistical evaluation involves calculating rejection rates based on critical values of the test statistic, as in Burnside (1994) and Cecchetti, Lam, and Mark (1994). However, Otrok, Ravikumar, and Whiteman (2002) show that tests based on the distance to the HJ bound are non-pivotal in finite samples: The finite-sample critical values depend upon the SDF parameters: risk aversion and the discount factor. Therefore, one has to calculate parameter-specific critical values for each point in the null hypothesis of interest. Nevertheless, for the case of time-separable preferences, Otrok, Ravikumar, and Whiteman (2002) show that the finite-sample distribution of the test statistic associated with the risk-neutral case is extreme. The critical values for the risk-neutral case deliver type-I errors no larger than intended, regardless of risk aversion or the discount factor. They also show that the maximal type-I error critical values for time-separable preferences are appropriate for habit formation as well as state nonseparable preferences. Their conclusion is that the HJ bound is indeed a useful statistical evaluation device, in that type-I errors can be controlled, while type-II error rates are acceptably small. Using their finite-sample critical values, they report evidence against time-separable preferences and mixed evidence for Epstein-Zin and habit preferences.

5 ASSET-PRICING IMPLICATIONS OF BUSINESS CYCLE MODELS

Our focus so far in this article has been on asset-pricing models and financial returns data typically used in the asset-pricing literature. The HJ bound is also useful for analyzing the asset-pricing implications of business cycle models. Such an approach is useful since a business cycle model is typically solved with a first-order approximation, which eliminates risk premia. Higher-order solutions are possible but costly for moderate-sized models. An early approach to using the HJ bound in the context of a business cycle model was Tallarini (2000). That paper first showed that risk-aversion per se did not affect the business cycle behavior of standard macroeconomic aggregates. It then showed that the SDF from that model did satisfy the HJ bound with sufficiently high risk aversion.


Typically, in business cycle models that study asset-pricing implications, asset return is measured by the S&P 500. Gomme, Ravikumar, and Rupert (2011) argue that business cycle theory does not necessarily imply using financial return to measure the return to capital. They construct the return to capital in the United States using NIPA statistics on capital income and capital stock. They show that while the mean return is roughly the same for equity returns and the return to capital, the NIPA return to capital is less volatile.

In Figure 6 we construct the HJ bound using the return to capital as in Gomme, Ravikumar, and Rupert (2011) and the return to the 3-month Treasury bill. We also plot the HJ bound from Figure 1 for comparison. The sample period here is 1959:Q2-2008:Q4. The bound for the return to capital is significantly higher. This follows from the fact that the return to capital is less volatile than the return to equity, which leads to a sharper set of restrictions on the set of admissible SDFs.

The implication for business cycle models is that the asset-pricing puzzle is in fact more challenging than the one we see with financial return data. As in Figure 2, the time-separable model does not generate enough volatility to satisfy the HJ bound with equity returns, so it

Figure 6

Return to Capital: Gomme, Ravikumar, and Rupert (2011)



NOTE: Sample: 1959:Q2-2008:Q4. The figure depicts the set of admissible SDFs in the mean-standard deviation space implied by stock and bond returns as in Figure 1. It also depicts the set of admissible SDFs implied by return to capital from Gomme, Ravikumar, and Rupert (2011) and bond return.

certainly will not satisfy the bound with return to capital. While it may be possible to find time- or state-nonseparable preferences that satisfy this bound, they will still suffer from the stability problem we documented earlier with these resolutions. In addition, in the case of time-nonseparable preferences, Otrok (2001) shows that the data prefer only moderate amounts of habit formation, which will not generate much volatility in the SDF.

6 CONCLUSION

The Hansen-Jagannathan bound is a helpful tool for understanding asset-pricing implications. By characterizing the admissible set of SDFs, we can use the bound to test proposed SDFs. Further, we can understand what types of asset-return data will pose greater difficulty for an asset-pricing model. Lastly, the bound can be constructed with only the first and second moments of asset-return data, hence implementation of the bound requires no computing power beyond a spreadsheet program.

The HJ bound uses the means, variances, and contemporaneous correlations of asset-return data to construct the lower bound that an SDF must satisfy. Otrok, Ravikumar, and Whiteman (2007) develop a volatility bound that uses serial correlation properties of the

return data as well. This generalization allows for an evaluation of whether models fail to match the data in the long run, at business cycle frequencies, etc. That is, the generalized bound can help identify the frequencies at which a model violates the necessary conditions. A business cycle model that violates the bound at business cycle frequencies might be unacceptable, but violations at other frequencies might not be a cause for concern. The generalization involves projecting the SDF onto the space of current, past, and future returns. Because the projection is onto a larger space than that for the HJ bound, the generalized bound is tighter than the HJ bound. They find that the state-nonseparable SDF satisfies the bound at business cycle frequencies, while the time-nonseparable SDF does poorly at those frequencies. Open questions for future research are whether the resolutions at those frequencies are stable and whether these SDFs can satisfy the generalized bound at business cycle frequencies when the bound is constructed with Gomme, Ravikumar, and Rupert (2011) capital-return data. ■

APPENDIX: HJ BOUND AND THE SHARPE RATIO

The Sharpe ratio is the mean excess return on an asset (relative to the risk-free rate) divided by the standard deviation of that asset's return. The Sharpe ratio measures how the market views risk: A higher Sharpe ratio implies that the market demands a higher return for a given level of risk. The connection between the HJ bound and the Sharpe ratio presented here follows Cochrane (2001).

Consider the unconditional Euler equation used to price assets:

$$1 = E(R^{\text{equity}} m).$$

For the sake of exposition we will assume that the only asset is equity, with return R^{equity} , and there is a risk-free rate such that $R^f = \frac{1}{E(m)}$. We can write the right-hand side of the equation as

$$E(R^{\text{equity}} m) = E(m)E(R^{\text{equity}}) + \rho_{R^{\text{equity}}, m} \sigma_{R^{\text{equity}}} \sigma_m,$$

where $\rho_{R^{\text{equity}}, m}$ is the correlation of equity returns with m and σ s represent standard deviations. Next, divide through by $E(m)$ to get

$$\frac{1}{E(m)} = E(R^{\text{equity}}) + \frac{\rho_{R^{\text{equity}}, m} \sigma_{R^{\text{equity}}} \sigma_m}{E(m)}.$$

Replacing $\frac{1}{E(m)}$ with R^f , dividing by $\sigma_{R^{\text{equity}}}$, and rearranging terms yields

$$\frac{E(R^{\text{equity}}) - R^f}{\sigma_{R^{\text{equity}}}} = -\frac{\rho_{R^{\text{equity}}, m} \sigma_m}{E(m)}.$$

Since $-1 \leq \rho_{R^{\text{equity}}, m} \leq 1$, we have the inequality

$$\frac{E(R^{\text{equity}}) - r^f}{\sigma_{R^{\text{equity}}}} \leq \frac{\sigma_m}{E(m)}.$$

The left-hand side is the Sharpe ratio. For given $E(m)$, a higher Sharpe ratio implies that the lower bound on SDF volatility is higher.

NOTES

- ¹ Matlab code and data used for this and all subsequent examples in this article can be found on Christopher Otrok's REPEC webpage: <https://ideas.repec.org/e/pot2.html>.
- ² For the value of σ that satisfies the bound, $Em = 0.9258$, or a quarterly return of $1/Em = 1.0801$.
- ³ Note that σ is not the coefficient of risk aversion in the habit model, though it is proportional to various measures of risk aversion. See Boldrin, Christiano, and Fisher (1997).
- ⁴ Since we can dismiss the Mehra-Prescott model for achieving the bound with only unreasonable amounts of risk aversion, we will focus on only these two models in this section.
- ⁵ The value 0.99 is in the lower part of the bound in Figure 1.

REFERENCES

- Boldrin, M.; Christiano, L.J. and Fisher, J.D.M. "Habit Persistence and Asset Returns in an Exchange Economy." *Macroeconomic Dynamics*, 1997, 1, pp. 312-32; <https://doi.org/10.1017/S1365100597003027>.
- Burnside, C. "Hansen-Jagannathan Bounds as Classical Tests of Asset Pricing Models." *Journal of Business and Economic Statistics*, 1994, 12, pp. 57-79; <https://doi.org/10.1080/07350015.1994.10509991>.
- Cecchetti, S.G.; Lam, P-S. and Mark, N.C. "Testing Volatility Restrictions on Intertemporal Marginal Rates of Substitution Implied by Euler Equations and Asset Returns." *Journal of Finance*, 1994, 49, pp. 123-52; <https://doi.org/10.1111/j.1540-6261.1994.tb04423.x>.
- Cochrane, J.H. *Asset Pricing*. Princeton University Press, 2001.
- Constantinides, G.C. "Habit Formation: A Resolution of the Equity Premium Puzzle." *Journal of Political Economy*, 1990, 98, pp. 519-43; <https://doi.org/10.1086/261693>.
- Epstein, L.G. and Zin, S.E. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis." *Journal of Political Economy*, 1991, 99(2), pp. 263-86; <https://doi.org/10.1086/261750>.
- Gomme, P.; Ravikumar, B. and Rupert, P. "The Return to Capital and the Business Cycle." *Review of Economic Dynamics*, 2011, 14(2), pp. 262-78; <https://doi.org/10.1016/j.red.2010.11.004>.
- Gregory, A.W. and Smith, G.W. "Sampling Variability in Hansen-Jagannathan Bounds." *Economics Letters*, 38, pp. 263-67; [https://doi.org/10.1016/0165-1765\(92\)90068-A](https://doi.org/10.1016/0165-1765(92)90068-A).
- Hansen, L.P. and Jagannathan, R. "Implications of Security Market Data for Models of Dynamic Economies." *Journal of Political Economy*, 1991, 99, pp. 225-62; <https://doi.org/10.1086/261749>.
- Hansen, L.P. and Singleton, K.J. "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models." *Econometrica*, 1982, 50(5), pp. 1269-86; <https://doi.org/10.2307/1911873>.
- Ljungqvist, L. and Sargent, T.J. *Recursive Macroeconomic Theory*. MIT Press, 2018.
- Mehra, R. and Prescott, E.C. "The Equity Premium: A Puzzle." *Journal of Monetary Economics*, 1985, 15, pp. 145-61; [https://doi.org/10.1016/0304-3932\(85\)90061-3](https://doi.org/10.1016/0304-3932(85)90061-3).
- Otrok, C.; Ravikumar, B. and Whiteman, C.H. "Evaluating Asset-Pricing Models Using the Hansen-Jagannathan Bound: A Monte Carlo Investigation." *Journal of Applied Econometrics*, 2002, 17, pp. 149-174; <https://doi.org/10.1002/jae.640>.
- Otrok, C.; Ravikumar, B. and Whiteman, C.H. "Stochastic Discount Factor Models and the Equity Premium Puzzle." Unpublished manuscript, 2004; <https://ideas.repec.org/p/pramprapa/22938.html>.
- Otrok, C.; Ravikumar, B. and Whiteman, C.H. "A Generalized Volatility Bound for Dynamic Economies." *Journal of Monetary Economics*, 2007, 54(8), pp. 2269-90; <https://doi.org/10.1016/j.jmoneco.2007.06.028>.
- Otrok, C. "On Measuring the Welfare Cost of Business Cycles." *Journal of Monetary Economics*, 2001, 47(1), pp. 61-92; [https://doi.org/10.1016/S0304-3932\(00\)00052-0](https://doi.org/10.1016/S0304-3932(00)00052-0).
- Tallarini, T. "Risk-Sensitive Real Business Cycles." *Journal of Monetary Economics*, 2000, 45(3), pp. 507-32; [https://doi.org/10.1016/S0304-3932\(00\)00012-X](https://doi.org/10.1016/S0304-3932(00)00012-X).
- Weil, P. "The Equity Premium Puzzle and the Risk-Free Rate Puzzle." *Journal of Monetary Economics*, 1989, 24, pp. 401-21; [https://doi.org/10.1016/0304-3932\(89\)90028-7](https://doi.org/10.1016/0304-3932(89)90028-7).

