# THE SHARPE RATIO AND PREFERENCES: A PARAMETRIC APPROACH

**MARTIN LETTAU**
*Federal Reserve Bank of New York
and
CEPR*

**HARALD UHLIG**
*Center for Economic Research, Tilburg University,
and
CEPR*

We use a log-normal framework to examine the effect of preferences on the market price for risk, that is, the Sharpe ratio. In our framework, the Sharpe ratio can be calculated directly from the elasticity of the stochastic discount factor with respect to consumption innovations as well as the volatility of consumption innovations. This can be understood as an analytical shortcut to the calculation of the Hansen–Jagannathan volatility bounds, and therefore provides a convenient tool for theorists searching for models capable of explaining asset-pricing facts. To illustrate the usefulness of our approach, we examine several popular preference specifications, such as CRRA, various types of habit formation, and the recursive preferences of Epstein–Zin–Weil. Furthermore, we show how the models with idiosyncratic consumption shocks can be studied.

**Keywords:** Sharpe Ratio, Hansen–Jagannathan Bounds, Volatility Bounds, Equity Premium

## 1. INTRODUCTION

The literature generated by Mehra and Prescott's (1985) observation of the equity premium puzzle is enormous. Kocherlakota (1996) provides an excellent summary. He surveys a variety of plausible explanations but concludes that no satisfactory solution has been found yet. One possible avenue proposed by many is the search for the "right" preferences. Mehra and Prescott (1985) considered time-separable CRRA preferences, but subsequent authors have suggested alternatives. Two leading candidates are *habit formation,* following Constantinides (1990), Abel (1990), and many others, and *recursive preferences*, as in Epstein and Zin (1989, 1991) and Weil (1991).

Hansen and Jagannathan (1991) have provided a popular framework, linking preferences and the process of consumption to the market price of risk (or the Sharpe ratio). We use a lognormal framework to study this link. Here, the Sharpe ratio can be calculated directly from the elasticity of the stochastic discount factor with respect to consumption innovations as well as the volatility of consumption innovations. This can be understood as an analytical shortcut to the calculation of the Hansen–Jagannathan volatility bounds, summarizing much of the information of these bounds in a single and easy-to-compute parameter. The paper thereby provides a convenient tool for theorists searching for models capable of explaining asset-pricing facts. We compute the elasticity of the stochastic discount factor to consumption innovations for a number of popular preference specifications and relate it to standard preference parameters such as risk aversion and the elasticity of intertemporal substitution.

In the case of time-separable CRRA preferences, the stochastic-discount-factor elasticity is equal to the parameter of relative risk aversion (RRA) and the inverse of the elasticity of intertemporal substitution (EIS). These links are broken for more general preferences. For preferences with habit formation, the stochastic-discount-factor elasticity is still the inverse of the EIS but is decoupled from the RRA. In other words, there is no direct link between risk aversion and risk premia once a habit is included in preferences. Risk premia depend only on the EIS. For recursive Epstein–Zin–Weil preferences, the picture is exactly reversed. The stochastic-discount-factor elasticity is equal to RRA and delinked from the EIS. These examples show that, in general, neither RRA nor EIS completely characterizes the Sharpe ratio.

Our analysis is based on two important assumptions. First, we assume that all relevant random variables (such as innovations in log consumption and log asset returns) are conditionally lognormal. Second, we assume that the stochastic discount factor associated with the preferences under consideration (henceforth we will call this object PSDF, for preference-based stochastic discount factor) can be approximated in a loglinear form. Of course, neither assumption is innocuous. Lognormality has a long tradition in asset pricing, going back to at least Hansen and Singleton (1982), but it is restrictive—as shown, for example, by Rubinstein (1976) and Ferson (1983). However, to obtain analytical expressions as in this paper, some parametric assumptions must be made. Note that the lognormal framework has proved to be successful in empirical implementations of volatility bounds; see, for example, Cecchetti et al. (1994) and Balduzzi and Kallal (1997). Bansal and Lehmann (1997) compute growth-optimal bounds on the mean of the stochastic factors using lognormal asset returns. Concerning the second assumption of loglinearity, note that the PSDF for the benchmark case of time-separable CRRA preferences is linear in log consumption growth. However, the PSDF for some more general preferences has to be approximated (e.g., by habit formation). The accuracy of the approximation generally will depend on specific values of the habit parameters.

Our log-normal Sharpe ratio framework can be used as a diagnostic tool for preferences in the following way: The standard deviation of consumption innovations is

about 0.56% in quarterly postwar data. Using the S&P 500 Index as a proxy for the true market portfolio yields an observed Sharpe ratio of about 0.27 (=2%/7.5%) per quarter. Our model implies that the elasticity of the PSDF with respect to innovations in consumption must be around −50 to generate a Sharpe ratio of 0.27. For time-separable CRRA preferences, this means that a coefficient of relative risk aversion of 50 is required. This is, of course, just a restatement of the original Mehra–Prescott observation. The interesting implication is that the observation that PSDF elasticity must be around −50 holds for *any* preference specification to which our framework applies, not only for time-separable CRRA preferences. Hence the search for preferences that generate a high enough Sharpe ratio can be reduced to a search for preferences with a PSDF elasticity of −50.

Our framework is a parametric version of Hansen and Jagannathan's (henceforth HJ) (1991) volatility bounds. They identify conditions of the first two moments of the PSDF under which a set of preferences is consistent with asset returns. They show that the HJ diagram for excess return has an isomorphic representation in the mean–standard deviation return diagram. In particular, their PSDF conditions (using excess asset returns) can be rephrased in terms of the Sharpe ratio. Our lognormal Sharpe ratio framework is a parametric version of the HJ bounds for excess returns, and we thus obtain a closed-form expression for the Sharpe ratio. This provides a convenient tool for obtaining additional insights about the connection among preferences and the Sharpe ratio.[1]

Finally, we discuss a second difficulty in explaining the equity premium observation. Suppose that some preference specification generates a Sharpe ratio of 0.27. The equity premium is equal to or smaller than this ratio multiplied by the volatility of excess returns. Dividend processes used in the theoretical literature are often not volatile enough to match the observed excess return volatility, and fail to explain the excess volatility observation of Shiller (1981). In summary, explaining the equity premium observation puzzle requires explaining the Sharpe ratio puzzle by choosing appropriate preferences, and explaining the return volatility puzzle by choosing the appropriate dividend process.

Many of the issues raised in this paper are also mentioned in the excellent survey article by Cochrane (1997). He also stresses the importance of the Sharpe ratio as a general measure risk–return trade-off in a given model. His back-of-the-envelope calculations about the required level of risk aversion in order to achieve a reasonable Sharpe ratio are similar to those presented in this paper. Our paper complements his work and provides additional insights into the dependence of the Sharpe ratio on preferences.

The remainder of the paper is organized as follows: Section 2 shows the decomposition of risk premia into the Sharpe ratio and dividends. Section 3 studies some preference specifications that have been proposed in the literature using our decomposition. Section 4 shows that applying volatility bounds in incomplete-markets settings can be misleading, and Section 5 concludes.

## 2. RISK PREMIA, HANSEN–JAGANNATHAN BOUNDS, AND THE SHARPE RATIO

The starting point of our analysis is the Lucas (1978) asset-pricing equation

$$E_t [M_{t+1} R_{t+1}] = 1, \tag{1}$$

where $E_t[\cdot]$ is the conditional expectation at time $t$; $M_{t+1}$ is a function of past, present, and future consumption of the agent, and $R_{t+1}$ is the gross return on some asset from period $t$ to $t + 1$ given by

$$R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t}, \tag{2}$$

where $P_t$ denotes the price in period $t$ after paying dividends $D_t$. Let $R^{rp}$ denote risk premia, that is, returns in excess of the risk-free rate $R^f$. $M_{t+1}$ is the *preference-based stochastic discount factor*, which one might think of as a generalized version of the intertemporal marginal rate of substitution. In fact, it coincides with the intertemporal marginal rate of substitution for the standard case of a time-separable utility function.

The asset pricing equation (1) implies bounds on the first and second conditional moments of returns and the PSDF:

$$\frac{E_t\left[R_{t+1}^{rp}\right]}{\sigma_t\left[R_{t+1}^{rp}\right]} = -\text{Corr}\left[R_{t+1}^{rp}, M_{t+1}\right] \frac{\sigma_t[M_{t+1}]}{E_t[M_{t+1}]}$$

$$\leq \frac{\sigma_t[M_{t+1}]}{E_t[M_{t+1}]} \equiv SR_t^{\max}. \tag{3}$$

This inequality says that the ratio of the expected return to the expected standard deviation of any asset is bounded from above by the ratio of the conditional standard deviation of the PSDF and its conditional mean. In other words, the maximal Sharpe ratio implied by a given PSDF is equal to the standard deviation of the PSDF divided by the conditional mean. Building on Shiller (1982), Hansen and Jagannathan (1991) provide a comprehensive analysis of this volatility bound, allowing for many risky assets and no riskless asset and derive implications of the restriction that the PSDF must be positive. They plot the volatility bounds in diagrams with $\sigma_t[M_{t+1}]$ and $E_t[M_{t+1}]$ on the axes.

In this paper, we provide a parametric version of the HJ bounds. For our purposes, it turns out that it is more convenient to recast the volatility bounds in the familiar mean–standard deviation diagram of returns. Since the HJ bounds contain the same information as the mean–standard deviation frontier, these two approaches are equivalent.[2] Next, we make a parametric assumption to derive a closed-form expression for $\sigma_t[M_{t+1}]/E_t[M_{t+1}]$ in (3). After obtaining such an expression, we use observed data on asset returns to compute $E_t[R_{t+1}^{rp}]/\sigma_t[R_{t+1}^{rp}]$ and check whether the maximum Sharpe ratio is consistent with actual returns.

To proceed with our analysis, it is useful to utilize a loglinear lognormal framework; see, for example, Hansen and Singleton (1982). Although this distributional assumption is not without controversy, it represents a good starting point to obtain more analytical insights into the mechanism of volatility bounds and risk premia. Assume that consumption $C_t$ is one-dimensional and let lowercase letters denote the logarithm of a variable; for example, $c_t = \log C_t$. We can decompose log consumption into its conditional expectations and its innovation

$$c_{t+1} = E_t c_{t+1} + \epsilon_{t+1}. \tag{4}$$

For the rest of the paper, we make the following assumptions.

Assumption 1. All relevant random variables are lognormaly distributed. For example, innovations in log consumption are normal: $\epsilon_t \sim N(0, \sigma_{\epsilon,t}^2)$.

Note that the variance is allowed to change over time; this allows for GARCH, stochastic volatility, or square-root processes, for example.

Assumption 2. The innovation to the log of the PSDF only depends on the innovation to the log of consumption:

$$m_{t+1} - E_t m_{t+1} = \eta_{m\epsilon,t+1} \epsilon_{t+1}, \tag{5}$$

where $m_{t+1} = \log M_{t+1}$.

This assumption follows if preferences are time separable and only depend on consumption. The assumption also follows for nonseparable preferences, if consumption is a martingale. If preferences depend only on consumption, but are nonseparable, and consumption is not a martingale, equation (5) generally does not hold. Instead, one needs to calculate

$$m_{t+1} - E_t[m_{t+1}] = \sum_{j=1}^{\infty} \eta_{m\epsilon,j,t+1} v_{t+1,j}$$

where

$$v_{t+1,j} = E_{t+1}[c_{t+j}] - E_t[c_{t+j}],$$

is the news at date $t + 1$ about consumption at date $t + j$ and $\eta_{m\epsilon,j,t+1}$ are the corresponding elasticities of the stochastic discount factor. Working with this expression is very cumbersome, and so , we stick to (5) above. The assumption also rules out cases in which the PSDF depends, for example on innovations to leisure or on a multidimensional vector of consumption innovations. Nonetheless, the assumption is in line with much of the CCAPM literature. Likewise, some restriction like the martingale assumption for consumption is also necessary for the computation of HJ bounds when preferences are not time separable; see Cochrane and Hansen (1992). So, Assumption 2 applies broadly enough to be useful.

The PSDF elasticity $\eta_{m\epsilon,t}$ is determined by preferences. For time-separable CRRA preferences, we have $m_{t+1} = -\gamma \Delta c_{t+1}$, where $\gamma$ is the coefficient of relative risk aversion. Hence $\eta_{m\epsilon,t} = -\gamma$.[3] In Section 3, we compute the decomposition (5)

for more general preferences that have been suggested in the literature (e.g., habit formation and Epstein–Zin–Weil nonrecursive preferences). In many of these cases the connection between risk aversion and $\eta_{m\epsilon}$ breaks down. In principle we allow for time-variable elasticities; see Campbell and Cochrane (2000) for an example in which this feature is important. For convenience of notation, we shall soon drop the time subscripts for all elasticities and variances and simply point out when time variation is important.

Let $r_{t+1} = \log R_{t+1}$ be the continuously compounded return of some asset. We decompose its news into a component that is proportional to the consumption news $\epsilon_{t+1}$ and an orthogonal component $\omega_{t+1}$:

$$r_{t+1} - E_t r_{t+1} = \eta_{r\epsilon}\epsilon_{t+1} + \omega_{t+1}, \tag{6}$$

where $\omega_{t+1}$ is normally distributed with $E_t[\epsilon_{t+1}\omega_{t+1}] = 0$ and $\sigma_\omega^2 = E_t\omega_{t+1}^2$.

Using the normality assumption, we can use the log version of the Lucas equation (1) to express the expected return of an asset as[4]

$$E_t r_{t+1} + \eta_{r\epsilon}^2 \sigma_\epsilon^2/2 = -E_t m_{t+1} - \eta_{m\epsilon}^2 \sigma_\epsilon^2/2. \tag{7}$$

For the log-risk-free rate $r_{t+1}^f$, we set $\eta_{r\epsilon} = \sigma_\omega^2 = 0$. Conditional risk premia $r_{t+1}^{rp}$ can be computed from (7):

$$E_t r_{t+1}^{rp} \equiv E_t r_{t+1} - r_{t+1}^f + \eta_{r\epsilon}^2 \sigma_\epsilon^2/2 = -\text{Cov}_t(r_{t+1}, m_{t+1}) \tag{8}$$

$$= -\eta_{m\epsilon}\eta_{r\epsilon}\sigma_\epsilon^2, \tag{9}$$

which is the consumption CAPM (CCAPM) written in logs.[5]

Under our lognormal assumption we find the following closed-form solution for the maximal Sharpe ratio. Since the PSDF is assumed to be lognormal, the volatility bound (3) implies

$$\text{SR}_t^{\max} = \left(e^{\sigma_{m,t}^2} - 1\right)^{1/2} \tag{10}$$

$$\approx -\eta_{m\epsilon}\sigma_\epsilon, \tag{11}$$

where $\sigma_{m,t}^2$ denotes the conditional variance of the log PSDF.[6] We can therefore rewrite expected risk premia as

$$E_t r_{t+1}^{rp} = \text{SR}_t^{\max}\eta_{r\epsilon}\sigma_\epsilon. \tag{12}$$

This expressions shows that preferences affect risk premia *only* through the Sharpe ratio. Given the Sharpe ratio, the expected risk premium of an asset only depends on the elasticity of the asset return with respect to the consumption innovation and the conditional standard deviation of consumption.

As explained previously, the Sharpe ratio can be used as a diagnostic test for preferences similar to the HJ bounds. Using the S&P500 as a proxy for the true

market portfolio, the Sharpe ratio is 0.27 ($=2\%/7.49\%$) in quarterly postwar data.[7] In Appendix A.1, we report a more detailed account of the Sharpe ratio in postwar data, including subperiods and a variety of other assets. Our benchmark value 0.27 appears to be a reasonable lower bound. It is also in line with other studies, for example, that by MacKinlay (1995), who advocates a quarterly Sharpe ratio of 0.3.[8] The standard deviation of consumption innovation $\sigma_\epsilon$ is 0.56% (using postwar Citibase data). This implies, according to (11), that the elasticity of the PSDF with respect to consumption innovations $-\eta_{m\epsilon}$ must be 48. For time-separable CRRA preferences, this translates into $\gamma = 48$ for the coefficient of relative risk aversion.[9] This is, of course, nothing more than a restatement of equity premium puzzle and the inability of CRRA preferences to enter the HJ bounds for low risk aversion. The advantage of this framework is that (11) holds for many different types of preferences. In other words, for a specific preference specification to be consistent with a Sharpe ratio of 0.27, its implied $-\eta_{m\epsilon}$ has to be (at least) 48. The following section shows how $-\eta_{m\epsilon}$ can be computed for a variety of different preferences.[10]

For time-separable CRRA preferences, higher risk aversion is associated with a more volatile PSDF causing an upward shift in the HJ diagram.[11] The corresponding picture is shown in Figure 1, which displays the capital market line (CML) for CRRA preferences. The slope of the CML is just $\gamma\sigma_\epsilon$. Fixing the risk-free rate at its postwar average of 0.24%, Figure 1 shows that the CML is not steep enough for $\gamma = 25$ to include the S&P500 Index. As we increase $\gamma$, the CML becomes steeper and the Sharpe ratio increases. For $\gamma = 50$, the CML includes the S&P point. The increase in the slope of the CML corresponds to an upward shift in the HJ diagram. Both approaches can be used to assess whether a certain preference specification is consistent with asset returns. The advantage of the HJ bounds is that, in addition to being nonparametric, they also incorporate information about the mean of the PSDF. Our lognormal model yields instead a simple closed-form expression for the Sharpe ratio. Instead of computing and plotting the HJ bounds
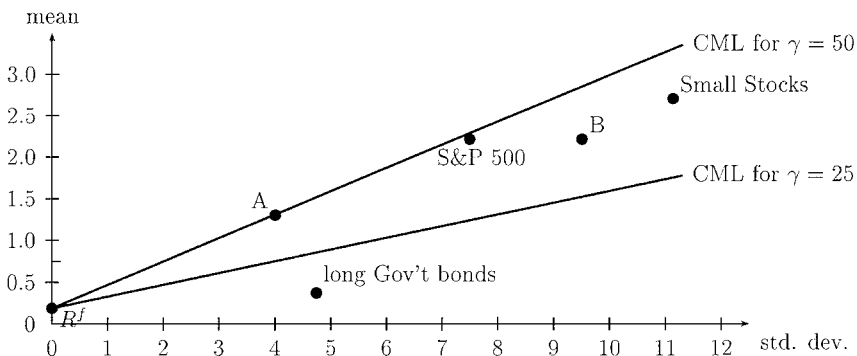


**FIGURE 1.** Conditional mean–standard deviation frontier implied by time-separable CRRA preferences with $\gamma = 25$ and $\gamma = 50$; $\sigma_\epsilon$ is set to 0.56%; the slope of the CML is given by $\gamma\sigma_\epsilon$.

for different preferences (and parameters), all that is needed for the Sharpe ratio is the preference parameter $\eta_{m\epsilon}$.

Using the lognormal framework just laid out, we can decompose the equity premium puzzle into two distinct pieces as follows[12]:

1. Explaining the Sharpe ratio of at least 0.27 [see equation (11)]. This requires finding a preference specification that generates $-\eta_{m\epsilon} = 48$ when using aggregate consumption data.
2. Explaining why the component of equity returns that is perfectly correlated with consumption has a volatility of 7.49% [equation (12)]. This requires an $\eta_{r\epsilon}$ of 13.4 for equity [again using aggregate consumption and (6)].

In other words, the equity premium puzzle decomposes into a purely preference-based Sharpe ratio puzzle (part 1) and a purely asset-based volatility puzzle (part 2). We have already seen that solving the Sharpe ratio puzzle with standard CRRA preferences requires high levels of risk aversion. A simple calculation shows that the volatility part of the equity premium is indeed also puzzling.

Suppose we have found preferences that produce a Sharpe ratio of 0.27 as in the data. Consider a one-period claim with $D_t = C_t^\xi$. In the usual Lucas tree-economy, $\xi = 1$; in a more general setting, $\xi$ can be interpreted as a leverage ratio [as in Campbell (1986)]. To generate the same return volatility as in the data requires a $\xi$ of 13.4 even given a Sharpe ratio of 0.27. Alternatively, consider an infinitely lived consumption claim. It is well known that all term premia are zero when log consumption follows a random walk. In this case, a long-run claim to consumption will carry the same risk premium as a one-period claim. If consumption growth is autocorrelated, then this is no longer true. In postwar data, the autocorrelation of quarterly consumption growth is 0.16. With this value, an infinite claim to consumption implies that $\eta_{r\epsilon} = 1.2$. Hence, the risk premium will be 20% higher than that of the one-period claim. Using the same argument as before, a $\xi$ of about 11 is required to obtain a 2% risk premium for the infinite claim. In other words, dividends have to be much more volatile than aggregate consumption even if the Sharpe ratio is 0.27.

Figure 1 also shows that both parts of the puzzle are important. Suppose a model produces an asset denoted as "A" in the figure. Since the asset lies on the CML, this model passes the HJ-bounds test. However, the risk premium implied by the model is too low because asset A is not volatile enough. In contrast, consider asset B. This asset has the same equity premium as the S&P500 Index, but the underlying model will fail the HJ-bounds test since the asset lies beneath the CML.

## 3. ALTERNATIVE PREFERENCES

In this section, we evaluate some preference specifications that have been studied in the literature using our Sharpe ratio diagnostic tool. We consider the standard time-separable CRRA case, various types of habit formation [Constantinides (1990)], "catching-up-with-the-Joneses" [Abel (1990)], and nonexpected utility [Epstein and Zin (1991) and Weil (1991)] to see whether they are capable of producing

a realistic risk-return trade-off (i.e., a Sharpe ratio of 0.27) for realistic parameter values. Recall that the only preference parameter that affects risk premia and the Sharpe ratio is $\eta_{m\epsilon}$. Hence, even if different preferences look very different at first sight, they can imply *exactly* the same risk premia, for a given a consumption stream, if the implied $\eta_{m\epsilon}$ coincide. Another implication is that neither risk aversion nor intertemporal substitution determines risk premia. As we will see, for certain preferences (e.g., CRRA and Epstein–Zin–Weil) there will be a link between $\eta_{m\epsilon}$ and risk aversion and hence risk aversion can be said to determine risk premia. However, for other preferences such as habit formation, $\eta_{m\epsilon}$ is linked to the elasticity of intertemporal substitution and not to risk aversion. For those preferences, risk aversion does not play a direct role for risk premia. Hence the deep parameter that controls risk premia for any general set of preferences is $\eta_{m\epsilon}$.

### 3.1. Time-Separable CRRA

Consumers maximize the following utility function:

$$U_t = E_t \sum_{j=0}^{\infty} \beta^j \frac{C_{t+j}^{1-\gamma} - 1}{1 - \gamma}. \tag{13}$$

The parameter $\gamma$ measures RRA. Note that the inverse of $\gamma$ equals the elasticity of intertemporal substitution (EIS): EIS $= 1/\gamma$. The PSDF is given by

$$M_{t+1} = \beta \left( \frac{C_{t+1}}{C_t} \right)^{-\gamma}, \tag{14}$$

or, in logs,

$$m_{t+1} = -\gamma \Delta c_{t+1} + \log\beta; \tag{15}$$

hence

$$-\eta_{m\epsilon} = \gamma. \tag{16}$$

In this case, $-\eta_{m\epsilon}$ is equal to the coefficient of relative risk aversion and the inverse of the EIS. To generate a Sharpe ratio of 0.27, time-separable CRRA preferences require high risk aversion and low EIS.

### 3.2. Habit Formation

A number of papers have studied the effect of habit formation on asset prices. Prominent examples include Abel (1990), Constantinides (1990), and, more recently, Campbell and Cochrane (2000). In this subsection, we analyze how different habit formulations affect the Sharpe ratio. We stress that we do not study the effect of habit formation on the risk-free rate. It is well known that many habit models imply a volatile risk-free rate that is inconsistent with the data.[13] Here, we concentrate on the effect on risk premia. As an alternative to standard models of habit formation, Abel (1990) proposes a simplified specification in which the habit

depends not on own lagged consumption but on the other agents' consumption, in Abel's words, "Catching-up-with-the-Joneses." The implications for risk premia can be quite different for both models, as we will see.

Let $X_t$ be the habit stock as of period $t$ and let $X_t$ be some function of past (and potentially current) consumption $C_t$, $C_{t-1}$, .... Two aspects are worth pointing out. First, we do not specify any particular functional form for the habit. Some authors have suggested slow-moving habits [e.g., Constantinides (1990), Heaton (1995), Boldrin et al. (1997), Campbell and Cochrane (2000)] whereas others make habit depend only on a single lag of consumption [Dunn and Singleton (1986), Abel (1990), Ferson and Constantinides (1991)]. Our general model includes all examples as special cases. Second, we allow consumption in the *current* period to affect the *current* habit stock. We will see the role of this assumption below.

Another distinguishing feature is how the habit stock affects utility. We consider three different cases. First, Abel (1990) suggests a *ratio* model:

$$U_t = E_t \sum_{j=0}^{\infty} \beta^j \frac{(C_{t+j}/X_{t+j})^{1-\gamma} - 1}{1 - \gamma}. \tag{17}$$

Alternatively, utility may depend on the *difference* between consumption and habit, along the lines of Constantinides (1990) model:

$$U_t = E_t \sum_{j=0}^{\infty} \beta^j \frac{(C_{t+j} - X_{t+j})^{1-\gamma} - 1}{1 - \gamma}. \tag{18}$$

Campbell and Cochrane (2000) have introduced a variant of the difference model. Instead of writing utility in terms of the habit stock $X_t$, they propose to work with the surplus ratio $S_t = (C_t - X_t)/C_t$ and

$$U_t = E_t \sum_{j=0}^{\infty} \beta^j \frac{(C_{t+j} S_{t+j})^{1-\gamma} - 1}{1 - \gamma}. \tag{19}$$

Strictly speaking, the Campbell–Cochrane model is only a variant of the difference model (18), but we consider it separately because it has received a lot of attention.

The three models are similar, but they display subtle differences in their implications for asset prices. We illustrate these using $\eta_{m\epsilon}$, the elasticity of the PSDF.

*"Catching-up-with-the-Joneses."* First, consider the simplified case of "Catching-up-with-the-Joneses." In equilibrium, agents treat the habit as exogenous and the PSDF does not reflect the effect of current consumption on future habits. For the ratio model, we have

$$M_{t+1} = \beta \left( \frac{C_{t+1}}{C_t} \right)^{-\gamma} \left( \frac{X_{t+1}}{X_t} \right)^{\gamma - 1}. \tag{20}$$

Computing the PSDF elasticity is straightforward:

$$-\eta_{m\epsilon} = \gamma + (1-\gamma)\eta_{x\epsilon}, \tag{21}$$

where $\eta_{x\epsilon}$ is the elasticity of the habit with respect to *current* consumption innovation. Note that $-\eta_{m\epsilon}$ is different from $\gamma$ only if the habit depends on the current innovation. However, if $\gamma > 1$, then a positive consumption innovation decreases $-\eta_{m\epsilon}$ for $\eta_{x\epsilon} > 0$. Consequently, a higher $\gamma$ is required to generate a high Sharpe ratio if the habit increases with current consumption.

The PSDF for the difference model is

$$M_{t+1} = \beta \left( \frac{C_{t+1}}{C_t} \right)^{-\gamma} \frac{(1 - X_{t+1}/C_{t+1})^{-\gamma}}{(1 - X_t/C_t)^{-\gamma}} \tag{22}$$

with the following approximation in logs:

$$m_{t+1} = -\gamma \Delta c_{t+1} + \frac{\gamma \bar{x}}{1 - \bar{x}} \left( \Delta x_{t+1} - \Delta c_{t+1} \right), \tag{23}$$

where $\bar{x}$ is the steady-state ratio of habit and consumption. The elasticity $\eta_{m\epsilon}$ is therefore

$$-\eta_{m\epsilon} = \gamma \frac{1 - \bar{x}\eta_{x\epsilon}}{1 - \bar{x}}. \tag{24}$$

If the elasticity of the habit with respect to the consumption innovation is unity, then $-\eta_{m\epsilon}$ equals $\gamma$ as in the time-separable case. However, if $\eta_{x\epsilon} < 1$, then $-\eta_{m\epsilon}$ is larger than in the time-separable case and increases with the steady-state ratio of habit and consumption.

For the Campbell–Cochrane surplus ratio model (19), we obtain

$$-\eta_{m\epsilon} = \gamma(1 + \eta_{s\epsilon}), \tag{25}$$

where $\eta_{s\epsilon}$ is the elasticity of the surplus ratio with respect to current consumption innovations.[14] In this version of "Catching-up-with-the-Joneses," $-\eta_{m\epsilon}$ equals $\gamma$ as long as the current surplus ratio does not depend on the current consumption shock. If the current surplus ratio increases in response to the current consumption shock, then $-\eta_{m\epsilon}$ increases, which in turn causes the Sharpe ratio to increase.

Notice that the elasticity of current habit with respect to current consumption innovations plays another role as stressed by Campbell and Cochrane (2000). If $\eta_{x\epsilon}$ (or $\eta_{s\epsilon}$ in their model) varies over time, then $-\eta_{m\epsilon}$ and the Sharpe ratio will also be time varying. Otherwise, $-\eta_{m\epsilon}$ is constant and so the Sharpe ratio is constant. Campbell and Cochrane argue strongly for a model with time variation in the Sharpe ratio in order to explain a variety of asset-pricing puzzles.

The foregoing analysis shows that the specific functional form of the habit plays no role whatsoever in the Sharpe ratio. We did not specify how the habit depends on past consumption. For the Sharpe ratio, it does not matter whether the habit is slow moving or depends on just one consumption lag. The only relevant feature (aside from the steady-state ratio of habit and consumption) is whether the habit

depends on *current* consumption. Of course, the functional form of the habit does affect particular asset prices, such as the risk-free rate.

How is $-\eta_{m\epsilon}$ related to more conventional preference parameters such as risk aversion and the elasticity of intertemporal substitution in the presence of an external habit? It is easy to check that $-\eta_{m\epsilon}$ is still the inverse of the EIS, as in the time-separable case. Risk aversion is defined as the curvature of the value function rather than simple consumption gambles. Computing risk aversion in general is quite complicated; see, for example, the discussion by Campbell and Cochrane (2000). Moreover, the coefficient of relative risk aversion is no longer identical to $-\eta_{m\epsilon}$. However, following the bulk of the literature, we can easily compute risk aversion along the nonstochastic steady state of the model. In this case, RRA is again equal to $-\eta_{m\epsilon}$, as for time-separable preferences.

*Habit.*    In internal habit models, consumers take the effect of current consumption on future habit explicitly into account. Let $u(C_t, X_t)$ denote the period-utility function. The PSDF can be written as

$$
M_{t+1} = \beta \frac{\dfrac{\partial u(C_{t+1}, X_{t+1})}{\partial C_{t+1}} + E_{t+1} \displaystyle\sum_{j=0}^{\infty} \beta^j \frac{\partial u(C_{t+1+j}, X_{t+1+j})}{\partial X_{t+1+j}} \frac{\partial X_{t+1+j}}{\partial C_{t+1}}}{\dfrac{\partial u(C_t, X_t)}{\partial C_t} + E_t \displaystyle\sum_{j=0}^{\infty} \beta^j \frac{\partial u(C_{t+j}, X_{t+j})}{\partial X_{t+j}} \frac{\partial X_{t+j}}{\partial C_t}}. \tag{26}
$$

It is obvious that the PSDF is much more complicated than in the external habit case. Although the PSDF elasticity can be written easily, the expression is too complex to be useful for any economic insight. Hence, we focus on a concrete example. Consider the difference model (18) combined with a linear habit $X_t = \theta C_{t-1} + \psi X_{t-1}$ as in Constantinides (1990). We evaluate all preference parameters along the steady state $\Delta c_t = g$. Messy but straightforward algebra yields the following PSDF elasticity:

$$
\eta_{m\epsilon} = -\frac{\gamma}{1 - \bar{x}} \frac{1 + \dfrac{\beta \theta^2 e^{-g(1+\gamma)}}{1 - \beta \psi^2 e^{-g(1+\gamma)}}}{1 - \dfrac{\beta \theta e^{-g\gamma}}{1 - \beta \psi e^{-g\gamma}}}. \tag{27}
$$

Note that, in contrast to the external habit model, $\eta_{m\epsilon}$ depends not only on the curvature parameter $\gamma$ and the discount factor $\beta$ but also on the habit parameters $\theta$ and $\psi$. Hence, the particular habit chosen does affect the Sharpe ratio and HJ bounds. This habit model differs from the "Catching-up-with-the-Joneses" model in another important dimension. Whereas the elasticity of intertemporal substitution is the inverse of $\eta_{m\epsilon}$ in both cases, for internal habit models the coefficient of relative risk aversion is no longer equal to $\eta_{m\epsilon}$ even on the steady state path. As shown by Constantinides (1990), the coefficient of relative risk aversion is now

**TABLE 1.** RRA versus $\eta_{m\epsilon}$ (Habit formation parameters: $\beta = 0.95$, $g = 0.44\%$, $\gamma = 1$, $\psi = 0$)[a]

| $\theta$ | RRA | $-\eta_{m\epsilon}$ |
|---|---|---|
| 0.1 | 1.005 | 1.238 |
| 0.3 | 1.011 | 2.159 |
| 0.5 | 1.046 | 4.667 |
| 0.7 | 1.106 | 14.27 |
| 0.8 | 1.182 | 32.36 |
| 0.9 | 1.410 | 113.99 |
| 0.95 | 1.863 | 336.56 |
| 0.99 | 5.427 | 2106.52 |

[a]This table shows RRA using (28) and $\eta_{m\epsilon}$ using (27) for different values of $\theta$. Parameters $\beta$ and $g$ are calibrated from postwar U.S. data; $\gamma$ is set to unity.

$$\text{RRA} = \frac{\gamma}{1 - \bar{x}\dfrac{e^{-\gamma g} - \beta\gamma}{e^{-\gamma g} - \beta(\theta + \psi)}}. \tag{28}$$

Agents with an internal habit react differently to changes in short-term consumption (as measured by the EIS) and changes in wealth (as measured by risk aversion). In general, agents with an external habit are more averse to the former than the latter.

Table 1 shows RRA and $\eta_{m\epsilon}$ for $\psi = 0$ and different values of $\theta$. The curvature in the period utility function $\gamma$ is set to unity. Increasing $\theta$ increases $\eta_{m\epsilon}$ much more than RRA. The reason is that habit-formation consumers are so averse to gambles in wealth because they can slowly adjust consumption in response to a decrease in wealth. In contrast, they are very reluctant to suddenly adjust consumption because this will require them to consume more in the future as well. This distinction has been discussed in some depth by Boldrin et al. (1995). Recall that a $-\eta_{m\epsilon}$ of around 50 is needed to obtain the point estimate of the Sharpe ratio of 0.27. Table 1 shows that this can be achieved by setting $\theta$ equal to about 0.82 with $\gamma = 1$. However, this implies that the proportion of the habit is around 82% of total consumption. Alternatively, one can increase $\gamma$ and select a lower $\theta$. Note that risk aversion does not increase to unrealistic values in either case.

## 3.3. Non-expected Utility

Building on Kreps and Porteus (1978), Epstein and Zin (1991) formulate a preference specification that allows the distinction between RRA and the elasticity of intertemporal substitution. The objective function is defined recursively as

$$U_t = \left[ (1 - \beta)C_t^{1-\sigma^{-1}} + \beta\left(E_t U_{t+1}^{1-\gamma}\right)^{\frac{1-\sigma^{-1}}{1-\gamma}} \right]^{\frac{1}{(1-\sigma^{-1})}}. \tag{29}$$

The notation is from Campbell (1993); $\gamma$ is the coefficient of RRA and $\sigma$ is the EIS. See Epstein and Zin (1991), Giovannini and Weil (1989), and Campbell (1993) for a more detailed description of these preferences.

It is straightforward to calculate the preference-based stochastic discount factor for pricing assets $M_{t+1}$ for these preferences. However, $M_{t+1}$ involves "unobservables" such as $U_{t+1}$. To make the asset-pricing equation econometrically more useful, Epstein and Zin (1989, 1991) and essentially the entire literature following it made three additional assumptions: that the agent is representative, that he owns assets (the returns to which he can either consume or reinvest), and that there is no other source of income. Using these obviously highly restrictive assumptions, Epstein and Zin (1991) show that the PSDF can be rewritten as

$$M_{t+1} = \beta^{\frac{1-\gamma}{1-\sigma^{-1}}} \left( \frac{C_{t+1}}{C_t} \right)^{-\frac{1-\gamma}{\sigma-1}} R_{m,t+1}^{\frac{\sigma^{-1}-\gamma}{1-\sigma^{-1}}}, \tag{30}$$

where $R_m$ is the gross return on invested wealth (i.e., the market portfolio). In logs, the PSDF can then be rewritten as

$$m_{t+1} = \frac{1-\gamma}{1-\sigma^{-1}} \log \beta - \frac{1-\gamma}{\sigma-1} \left( \Delta c_{t+1} + \frac{1-\sigma\gamma}{1-\gamma} r_{m,t+1} \right). \tag{31}$$

It seems that an EIS close to unity can generate large $\eta_{m\epsilon}$ irrespective of RRA. However, as shown by Campbell (1993), the model implies a relationship between unexpected consumption and unexpected return on the market portfolio:

$$c_{t+1} - E_t c_{t+1} = r_{m,t+1} - E_t r_{m,t+1} + (1-\sigma)(E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{m,t+1+j}, \tag{32}$$

where $\rho$ is a parameter that is related to the consumption/wealth ratio and is close to unity. Substituting out the market return in (31), we obtain

$$m_{t+1} = E_t m_{t+1} - \gamma \epsilon_{t+1} - (1-\sigma\gamma)\lambda_{m,t+1}, \tag{33}$$

where $\lambda_{m,t+1} = (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{m,t+1+j}$ represents the "news" of future expected returns of the market portfolio between period $t$ and $t+1$. Hence

$$\eta_{m\epsilon} = -\gamma - (1-\sigma\gamma) \frac{\partial \lambda_{m,t+1}}{\partial \epsilon_{t+1}}. \tag{34}$$

Equation (34) tells us that an EIS close to unity does *not* imply a high $\eta_{m\epsilon}$. The impact of the consumption innovation on the PSDF depends only on RRA, as in the CRRA case. The EIS affects the PSDF only through changes in expected returns of the market portfolio. If returns are i.i.d. or if $\gamma = 1/\sigma$, we obtain the same PSDF as in the time-separable CRRA case [see also Campbell (1993) for a more extensive discussion]. The effect of $\lambda_m$ in (34) is likely to be fairly small because returns are difficult to forecast. Therefore, the Epstein and Zin preferences require a high risk aversion to generate a high Sharpe ratio. In this regard,

they are only a small improvement over the standard time-separable CRRA preferences. Of course, Epstein–Zin preferences do allow for high risk aversion without simultaneously implying a low EIS.

### 3.4. Summary and Outlook

Table 2 summarizes the results in this section. It reports the coefficient of relative risk aversion, the elasticity of intertemporal substitution, and the elasticity of the PSDF with respect to consumption innovations $-\eta_{m\epsilon}$. Recall that $-\eta_{m\epsilon}$ determines the Sharpe ratio and hence the HJ bounds.

Risk aversion, intertemporal substitution, and the PSDF elasticity are closely linked for time-separable CRRA and "Joneses" habit cases. To pass the HJ-bounds or Sharpe ratio tests, these preferences require high risk aversion and a low EIS. For internal habit formation and Epstein–Zin–Weil recursive preferences, the link between RRA, EIS, and $-\eta_{m\epsilon}$ is broken. In the habit case, $-\eta_{m\epsilon}$ equals the inverse of the EIS, while $-\eta_{m\epsilon}$ coincides with risk aversion for Epstein–Zin–Weil preferences. Habit formation allows $-\eta_{m\epsilon}$ to be high while keeping risk aversion low. The converse holds in the Epstein–Zin–Weil case, $-\eta_{m\epsilon}$ can be high without restricting the EIS. Of course, risk aversion will be high as well.

The table also shows that none of the preferences considered here is able to generate a high Sharpe ratio while keeping risk aversion fairly low *and* the EIS fairly high, as would be desirable. Most researchers have objections against high levels of risk aversion. It is hard to estimate risk aversion directly, but Barsky, et al. (1995) present survey evidence. They find that most individuals are very risk averse, with an average risk aversion across individuals of about 4. As an illustration of the implications of an RRA of 50, consider an agent who is faced with a 50/50 gamble of gaining or losing 10% of her total wealth. An agent whose RRA is unity would be willing to pay 0.5% of her wealth to avoid that gamble, an agent with RRA of 50 would be willing to pay 8.7% of her wealth.

On the other hand, an elasticity of intertemporal substitution that is close to zero is hard to reconcile with models in which agents can smooth consumption. A low EIS implies that agents prefer a very smooth consumption path; if they have an opportunity to smooth income shocks, then they will do so. This can lead to counterfactually smooth consumption paths. In the case of habit formation, $-\eta_{m\epsilon} = 50$ implies an EIS of 0.02. Consumers with such a low EIS are extremely reluctant to adjust consumption in response to changes in interest rates. Lettau (2000) and Lettau and Uhlig (2000) demonstrate these effects in the context of a real-business-cycle model in which consumption can react endogenously in response to technology shocks. Increasing risk aversion or introducing habit formation will increase $-\eta_{m\epsilon}$ but result also in smoother consumption ($\sigma_\epsilon$ is lower). Hence, the effect on the Sharpe ratio (and HJ bounds) is much smaller than in exchange economies in which consumption is endogenous. The foregoing discussion uses the *observed* standard deviation of consumption innovations rather than one that is implied by a specific model.

TABLE 2. Alternative preferences (parameters: $\beta = 0.95$, $g = 0.44\%$, $\gamma = 1$, $\psi = 0$)

| Preferences | RRA | 1/EIS | $-\eta_{m\epsilon}$ |
|---|---|---|---|
| Time-separable CRRA | $\gamma$ | $\gamma$ | $\gamma$ |
| "Joneses" ratio | $\gamma + (1-\gamma)\eta_{x\epsilon}$ | $\gamma + (1-\gamma)\eta_{x\epsilon}$ | $\gamma + (1-\gamma)\eta_{x\epsilon}$ |
| "Joneses" difference | $\gamma \dfrac{1-\bar{x}\eta_{x\epsilon}}{1-\bar{x}}$ | $\gamma \dfrac{1-\bar{x}\eta_{x\epsilon}}{1-\bar{x}}$ | $\gamma \dfrac{1-\bar{x}\eta_{x\epsilon}}{1-\bar{x}}$ |
| "Joneses" surplus | $\gamma(1+\eta_{x\epsilon})$ | $\gamma(1+\eta_{x\epsilon})$ | $\gamma(1+\eta_{x\epsilon})$ |
| Habit difference | $\dfrac{\gamma}{1-\bar{x}}\dfrac{e^{-\gamma g}-\beta\gamma}{e^{-\gamma g}-\beta(\theta+\psi)}$ | $-\dfrac{\gamma}{1-\bar{x}}\dfrac{1+\dfrac{\beta\theta^2 e^{-g(1+\gamma)}}{1-\beta\psi^2 e^{-g(1+\gamma)}}}{1-\dfrac{\beta\theta e^{-g\gamma}}{1-\beta\psi e^{-g\gamma}}}$ | $-\dfrac{\gamma}{1-\bar{x}}\dfrac{1+\dfrac{\beta\theta^2 e^{-g(1+\gamma)}}{1-\beta\psi^2 e^{-g(1+\gamma)}}}{1-\dfrac{\beta\theta e^{-g\gamma}}{1-\beta\psi e^{-g\gamma}}}$ |
| Epstein–Zin–Weil | $\gamma$ | $\sigma$ | $\gamma$ |

*Note:* The table shows the coefficient of relative risk aversion (RRA), the inverse of the intertemporal elasticity of substitution (1/EIS), and the elasticity of the PSDF with respect to consumption innovations $\eta_{m\epsilon}$ for preferences in Section 3 at steady state. For Epstein–Zin preferences, returns on the market portfolio are assumed to be i.i.d. The parameters are defined in the respective subsections.

## 4. INCOMPLETE MARKETS

Recently, researchers have expanded the standard complete markets model to allow for idiosyncratic risk for the individual consumer. Constantinides and Duffie (1996), Den Haan (1996), and Heaton and Lucas (1996) have studied versions of these models as a possible avenue to explain high-risk premia. In this section, we use the Sharpe ratio to evaluate the potential of models with idiosyncratic shocks. We demonstrate that some care is required in computing HJ bounds and Sharpe ratios in models with idiosyncratic risk. Here, we present only a simple example; see Lettau (1998) for a more complete analysis. In particular, we abstract from frictions such as borrowing constraints. The simple derivation here is in the spirit of Constantinides and Duffie (1996). In their model, agents cannot insure against their idiosyncratic labor risk, but the PSDF depends only on aggregate variables.

Decompose log consumption of consumer $i$, $c_{t+1}^i$, as follows:

$$c_{t+1}^i = E_t c_{t+1}^i + \epsilon_{t+1} + \epsilon_{t+1}^i, \tag{35}$$

where $\epsilon_{t+1}$ represents the aggregate shock common to each consumer. The term $\epsilon_{t+1}^i$ denotes the idiosyncratic shock of consumer $i$ and is orthogonal to all aggregate variables. We assume that $\epsilon_{t+1}$ and $\epsilon_{t+1}^i$ are independently normally distributed with zero mean and variance $\sigma_\epsilon^2$ and $\sigma_{\epsilon^i}^2$, respectively. The log PSDF of consumer $i$ is then

$$m_{t+1}^i = E_t m_{t+1}^i + \eta_{m\epsilon} \left( \epsilon_{t+1} + \epsilon_{t+1}^i \right). \tag{36}$$

A quick calculation using (11) suggests that the Sharpe ratio becomes

$$\mathrm{SR}_t^{\max} = -\eta_{m\epsilon} \sqrt{\sigma_\epsilon^2 + \sigma_{\epsilon^i}^2}. \tag{37}$$

For $\eta_{m\epsilon} = -5$ and $\sigma_\epsilon = 0.56\%$, the standard deviation of the idiosyncratic component $\sigma_{\epsilon^i}$ must be 5.4% to create a Sharpe ratio of 0.27. Hence, the idiosyncratic consumption must be about 10 times as large as the aggregate consumption risk.

However, a closer look shows that this calculation is misleading. The reason is that the PSDF now includes a random variable, which (by definition) is not correlated with aggregate variables, such as asset return. So, although the PSDF becomes more variable, its correlation with asset returns decreases. In the case of time-separable CRRA preferences, these effects exactly cancel, leaving HJ bounds and the Sharpe ratio unchanged. To simplify the argument, we use the approximation $\rho(e^x, e^y) \approx \rho(x, y)$ for normal $x$ and $y$ as well as small variances and covariance.[15] The Sharpe ratio in a model with idiosyncratic consumption risk is

$$\mathrm{SR} = -\eta_{m\epsilon} \sqrt{\sigma_\epsilon^2 + \sigma_{\epsilon^i}^2} \left[ \max_{\text{all assets}} -\rho_t(m_{t+1}, r_{t+1}) \right] \tag{38}$$

$$= -\eta_{m\epsilon} \sqrt{\sigma_\epsilon^2 + \sigma_{\epsilon^i}^2} \left[ \max_{\text{all assets}} \frac{\sigma_t(\epsilon_{t+1}, r_{t+1})}{\sigma_r \sqrt{\sigma_\epsilon^2 + \sigma_{\epsilon^i}^2}} \right] \tag{39}$$

$$= -\eta_{m\epsilon}\sigma_\epsilon \left[ \max_{\text{all assets}} \rho_t(\epsilon_{t+1}, r_{t+1}) \right] \tag{40}$$

$$= -\eta_{m\epsilon}\sigma_\epsilon. \tag{41}$$

For time-separable CRRA preferences, the Sharpe ratio is not at all affected by the presence of idiosyncratic shocks as in (35). This argument shows that simply adding some uninsurable income risk to the standard complete markets model will not substantially increase risk premia. Den Haan (1996) and Heaton and Lucas (1996) have studied models in which subsets of agents are subject to uninsurable shocks. Heaton and Lucas (1996) report somewhat higher risk premia in their economy with two groups of agents, each of which is subject to an idiosyncratic shock. The reason for higher-risk premia in such a model is that half of the population is subject to a shock. If many agents are affected by a shock, their reaction will have an effect on the entire market. This causes a correlation between a shock and aggregate dividends, which in turn increases the Sharpe ratio. Note that this argument depends on the market power of those agents who are subject to a common shock. Den Haan (1996) shows that it is much harder to increase risk premia in models in which each individual agent is subject to a shock. Since one agent cannot affect the aggregate market, there is no correlation with an idiosyncratic shock; hence, the Sharpe ratio will remain low.

Constantinides and Duffie (1996) present a more elaborate model in which the conditional variance of the cross-sectional distribution enters the pricing kernel; this variance is varying over time and can be correlated with aggregate dividends. This correlation can increase the Sharpe ratio. To allow for this channel to work through second moments in our setup, we would have to relax the assumption of constant variances of the shocks. What the argument in this section shows, however, is that simple models including idiosyncratic shocks will not substantially increase the risk-return trade-off. Only more complicated models working through second moments have the potential to increase the Sharpe ratio in an economy.

## 5. CONCLUSION

We propose a parametric framework, based on lognormal distributions, that allows us to express the Sharpe ratio as a function of a single preference parameter and moments of the aggregate consumption process. The key preference parameter is the elasticity of the stochastic discount factor associated with the preferences with respect to innovations in consumption. This parameter, together with the volatility of consumption innovations, suffices to calculate the Sharpe ratio. In general, this PSDF elasticity differs from more standard preference parameters such as risk aversion or the elasticity of intertemporal substitution. We compute the PSDF elasticity for a variety of popular preference specifications: time-separable CRRA, various types of habit formation, and recursive preferences of the Epstein–Zin–Weil type. For each case, we show how the PSDF elasticity relates to relative risk aversion and

the elasticity of intertemporal substitution. Computing the PSDF elasticity corresponds to plotting the bounds in the original HJ diagram. Our parametric approach complements the ability of the nonparametric HJ bounds to evaluate preferences regarding their ability to explain risk premia. Finally, we extend the model to allow for idiosyncratic consumption shock. Given our lognormal framework, we show that the HJ bounds are unaffected by idiosyncratic shocks in the case of CRRA preferences.

*NOTES*

1. We focus on the Sharpe ratio restriction. Additional restrictions, such as gain–loss bound implications [Bernardo and Ledoit (1996)], entropy bounds [Stutzer (1995)], growth-optimal bounds [Bansal and Lehmann (1997)], and market frictions [Luttmer (1996)] are outside the scope of the paper.

2. It is easiest to compare the mean–standard deviation frontier to HJ bounds in the case of excess returns. In this case, the HJ bound is just the area above a straight line instead of the area inside a parabola. Consequently, we have this case in mind when we refer to HJ bounds in the remainder of the paper.

3. It is easily shown that for any *time-separable* preferences, $-\eta_{m\epsilon} = \text{RRA} = 1/\text{EIS}$.

4. For normal $x$, $E(e^x) = e^{E(x)+0.5\text{Var}(x)}$.

5. The squared terms in (8) are Jensen terms that disappear when log returns are converted back into simple returns.

6. The Sharpe ratio in the data is at least 0.27 (see Table 1). For small numbers like this, the approximation is very precise: $[e^{(0.265^2)} - 1]^{0.5} = 0.27$.

7. The calculations are based on unconditional moments, although our framework is written in terms of conditional moments. Allowing for conditioning information would raise interesting questions, for example, how the Sharpe ratio might vary over time or how much returns are predictable. However, an extensive econometric analysis of these issues is beyond the scope of this paper.

8. Note that using the S&P as a market portfolio proxy yields a lower bound of the Sharpe ratio.

9. Cochrane (1997) presents a similar calculation.

10. In the remainder of the paper, we use quarterly data; however, the results are not sensitive to the frequency of the data. For example, in annual data the ratio of the mean S&P return to its standard deviation is 0.5, close to two times the corresponding quarterly number. The standard deviation of annual consumption growth is about 0.6%.

These calculations assume that innovations in consumption $\epsilon$ and returns on the market portfolio are perfectly correlated. In the data, this is far from true. Assuming a random walk for consumption, the correlation between consumption innovations and S&P500 returns is only 0.14. If we take this value as given, the true Sharpe ratio must be at least $0.27/0.14 = 1.92$, which implies that $-\eta_{m\epsilon} \geq 344$. For time-separable CRRA, this implies a relative risk aversion coefficient of 344.

11. Again, we refer to the HJ bounds using excess returns.

12. This separation refers to a partial equilibrium setting, that is, given $\sigma_\epsilon$. Of course, in general equilibrium, asset volatilities and correlations depend on preferences as well. See the end of Section 3 for a discussion on how a general equilibrium setup would affect our framework.

13. A notable exception is the model of Campbell and Cochrane (2000), which has a constant risk-free rate.

14. The elasticity $\eta_{s\epsilon}$ is related to $\eta_{x\epsilon}$ according to $\eta_{s\epsilon} = \bar{x}/(1 - \bar{x})(1 - \eta_{x\epsilon})$.

15. The exact expression is

$$\rho(e^x, e^y) = \frac{e^{\sigma_{xy}} - 1}{\left(e^{\sigma_x^2} - 1\right)^{1/2} \left(e^{\sigma_y^2} - 1\right)^{1/2}}.$$

*REFERENCES*

Abel, A. (1990) Asset prices under habit formation and catching-up-with-the-Joneses. *American Economic Review* 80, 38–42.

Balduzzi, P. & H. Kallal (1997) Risk premia and variance bounds. *Journal of Finance* 52, 1913–1950.

Bansal, R. & B.N. Lehmann (1997) Growth-optimal portfolio restrictions on asset pricing models. *Macroeconomic Dynamics* 1, 333–354.

Barsky, R.B., M.S. Kimball, F.T. Juster, & M.D. Shapiro (1995) Preference Parameter and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Survey. NBER working paper 5213.

Bernardo, A. & O. Ledoit (1996) Gain, Loss and Asset Pricing. Working paper, University of California at Los Angeles.

Boldrin, M., L. Christiano, & J. Fisher (1995) Macroeconomic Lessons for Asset Pricing. NBER working paper 5262.

Boldrin, M., L. Christiano, & J. Fisher (1997) Habit persistence and asset returns in an exchange economy. *Macroeconomic Dynamics* 1, 312–332.

Campbell, J. (1986) Bond and stock returns in a simple exchange model. *Quarterly Journal of Economics* 101, 785–803.

Campbell, J.Y. (1993) Intertemporal asset pricing without consumption data. *American Economic Review* 83, 487–512.

Campbell, J.Y. & J. Cochrane (2000) By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107, 207–251.

Cecchetti, S.G., P.-S. Lam, & N.C. Mark (1994). Testing volatility restrictions on intertemporal marginal rates of substitution implied by Euler equations and asset returns. *Journal of Finance* 49, 123–152.

Cochrane, J. (1997) Where is the market going? Uncertain facts and novel theories. *Economic Perspectives,* Federal Reserve Bank of Chicago, 11, 3–38.

Cochrane, J.H. & L.P. Hansen (1992) Asset pricing explorations for macroeconomics. In J. Blanchard & S. Fisher (eds.), *NBER Macroeconomics Annual,* pp. 115–165. Cambridge, MA: MIT Press.

Constantinides, G.M. (1990) Habit formation: A resolution to the equity premium. *Journal of Political Economy* 98, 519–543.

Constantinides, G.M. & D.Duffie (1996) Asset pricing with heterogeneous consumers. *Journal of Political Economy* 104, 219–240.

Den Haan, W. (1996) Understanding Equilibrium Models with a Small and Large Number of Agents. Working paper, University of California at San Diego.

Dunn, K. & K. Singleton (1986) Modeling the term structure of interest rates under habit formation and durability of goods. *Journal of Financial Economics* 17, 27–55.

Epstein, L.G. & S.E. Zin (1989) Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica* 57, 937–969.

Epstein, L. & S. Zin (1991) Substitution, risk aversion, and the temporal behavior of consumption growth and asset returns II: An empirical analysis. *Journal of Political Economy* 99, 263–286.

Ferson, W. (1983) Expectations of real interest rates and aggregate consumption: Empirical tests. *Journal of Financial and Quantitative Analysis* 18, 477–497.

Ferson, W. & G. Constantinides (1991) Habit formation and durability in aggregate consumption. *Journal of Financial Economics* 29, 199–240.

Giovannini, A. & P. Weil (1989) Risk Aversion and Intertemporal Substitution in the Capital Asset Pricing Model. NBER working paper 2824.

Hansen, L.P. & R. Jagannathan (1991) Implications of security market data for models of dynamic economies. *Journal of Political Economy* 99, 225–262.

Hansen, L.P & K. Singleton (1982) Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica* 50, 1269–1288.

Heaton, J. (1995) An empirical investigation of asset pricing with temporally dependent preference specifications. *Econometrica* 63, 681–717.

Heaton, J. & D. Lucas (1996) Evaluating the effects of incomplete markets on risk sharing and asset pricing. *Journal of Political Economy* 104, 443–487.

Kocherlakota, N.R. (1996) The equity premium: It's still a puzzle. *Journal of Economic Literature* 34, 42–71.

Kreps, D.M. & E.L. Porteus (1978) Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46, 185–200.

Lettau, M. (1998) Idiosyncratic Risk and Volatility Bounds, or, Can Models with Idiosyncratic Risk Solve the Equity Premium Puzzle? CEPR discussion paper 1795.

Lettau, M. (2000) Inspecting the Mechanism: The Determination of Asset Prices in the RBC Model. Revised from CEPR discussion paper 1884, March 2000.

Lettau, M. & H. Uhlig (2000) Can habit formation be reconciled with business cycle facts? *Review of Economic Dynamics* 3, 79–99.

Lucas, R. (1978) Asset prices in an exchange economy. *Econometrica* 46, 1429–1445.

Luttmer, E. (1996) Asset pricing in economies with frictions. *Econometrica* 64, 1439–1467.

MacKinlay, C. (1995) Multifactor models: Do they explain deviations from the CAPM? *Journal of Financial Economics* 38, 3–28.

Mehra, R. & E. Prescott (1985) The equity premium puzzle. *Journal of Monetary Economics* 15, 145–161.

Rubinstein, M. (1976) The strong case for the generalized logarithmic utility model as the premier model of financial markets. *Journal of Finance* 31, 551–571.

Shiller, R. (1981) Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71, 421–436.

Shiller, R. (1982) Consumption, asset markets and macroeconomic fluctuations. *Carnegie-Rochester Conference Series on Public Policy* 17, 203–238.

Stutzer, M. (1995) A Bayesian approach to diagnosis of asset pricing models. *Journal of Econometrics* 68, 367–397.

Weil, P. (1991) Nonexpected utility in macroeconomics. *Quarterly Journal of Economics* 105, 29–42.

# APPENDIX

## A.1. THE SHARPE RATIO IN POSTWAR DATA

Table A.1 summarizes some important and well-known facts about asset markets. The data are sampled at quarterly frequency and range from 48Q1 to 95Q4. Returns are real and reported in percent. We consider four different assets: the 30-day T-bill, long-term government bonds, a portfolio of small stocks, and the S&P 500 portfolio of larger companies. Table A.1 reports the mean and the standard deviation of the four assets for the entire sample as well as for three subsamples. The subsamples are chosen somewhat arbitrarily to demonstrate three distinct postwar episodes for asset prices. The pattern of asset returns is well known. Equities have, on average, a much higher return than short-term and long-term bonds. Their standard deviation is also substantially higher. The exception are the 1980's, when returns of large stocks were low (however, small stocks did much better). Small stocks in general have a higher standard deviation, which is partially compensated by a somewhat higher average return.

  We also compute the Sharpe ratio defined as the slope of the capital market line (CML). Since the true market portfolio is not observable, as a proxy we use the portfolio with the

**TABLE A.1.** Asset-market facts, U.S. data, 1948–1995[a]

| Asset | Real returns, % per quarter | | | |
| --- | --- | --- | --- | --- |
| | 1948–1995 | 1948–1972 | 1973–1982 | 1983–1995 |
| T-bills | | | | |
|    Mean | 0.24 | 0.13 | −0.01 | 0.64 |
|    Std. dev. | 0.72 | 0.71 | 0.96 | 0.59 |
| Long gov't bond premium | | | | |
|    Mean | 0.21 | −0.19 | −0.47 | 1.40 |
|    Std. dev. | 4.80 | 2.83 | 7.07 | 5.47 |
| S&P500 premium | | | | |
|    Mean | 2.00 | 2.49 | 0.12 | 2.49 |
|    Std. dev. | 7.49 | 6.75 | 9.53 | 6.96 |
| Small stocks premium | | | | |
|    Mean | 2.70 | 2.64 | 3.63 | 2.01 |
|    Std. dev. | 11.13 | 10.26 | 14.69 | 9.93 |
| Sharpe ratio | | | | |
|    Mean | 0.27 | 0.39 | 0.33 | 0.41 |
|    95th percentiles | (0.12, 0.44) | (0.19, 0.64) | (0.04, 0.70) | (0.14, 0.75) |

[a]Returns are measured at quarterly frequency. Risk premia are computed as the difference between the asset return and the T-bill rate. The Sharpe ratio is computed as the maximal ratio of mean excess return to standard deviation of the excess return of portfolios formed by the three risky assets. The 95th percentiles of the Sharpe ratio are calculated using a bootstrap with 10,000 replications. *Source*: Ibbotson Associates.

highest ratio of average return to standard deviation formed by the three risky assets. For the whole sample, this portfolio is basically the S&P500 Index itself; adding long bonds or small stocks does not add much. In the subperiods, however, the Sharpe ratio is substantially higher than that of the deviation of the S&P500 Index. The effect is especially pronounced in the 1970's when large stocks did very poorly. The Sharpe ratio for the entire sample is 0.27 on a quarterly basis though it is higher in each of the three subperiods. For the remainder of this paper, we use 0.27, which is in line with other studies, for example, MacKinlay (1995) who advocates a quarterly Sharpe ratio of 0.3.

To get a rough idea about the precision of this estimate, we perform a bootstrap with 10,000 replications. The 95th percentiles for the 1948–1995 sample are 0.12 and 0.44, indicating that there is quite a bit of uncertainty about the estimate. The quantiles are even wider than for the whole sample, which is not surprising because the subsamples are shorter.

Figure A.1 shows the return data for the sample from 1948–1996 in the standard textbook mean–standard deviation frontier.

## A.2. APPROXIMATION ERROR

In this section, we briefly address how accurate some linear approximations are. First, consider the loglinear approximation of the PSDF in the Catching-up-with-the-Joneses model (23). Figure A.2 shows the nonlinear log of the PSDF and the linearized log PSDF for two values of the steady-state ratio of habit and consumption. Habit and consumption at time $t$ are assumed to be at their steady-state values. Shocks to the ratio of habit and consumption in $t+1$ are on the $y$ axis (e.g., 0.02 means that $X_{t+1}/C_{t+1}$ is 2% above the steady-state
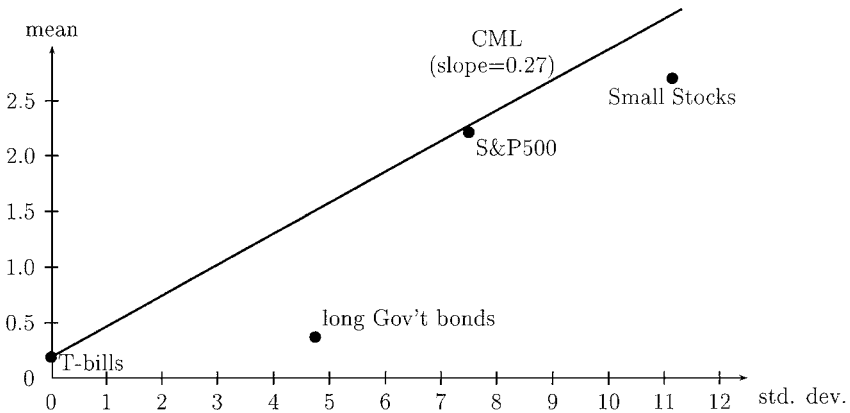
**FIGURE A.1.** Mean–standard deviation frontier: CML with postwar quarterly data from Table A.1.
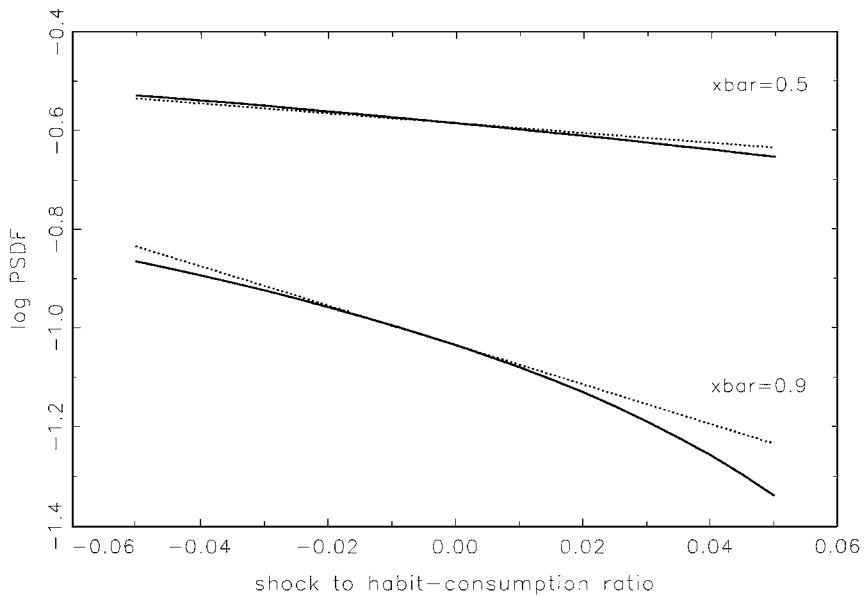


**FIGURE A.2.** Approximation error for steady-state ratio of habit and consumption.

level); $\gamma$ is set to unity. The figure shows that the loglinear approximation works fairly well. The difference between the true nonlinear PSDF and the linearized version is large only if the habit is, on average, very close to consumption and shocks to the ratio are very large.

Next, consider the linearization of the maximum Sharpe ratio in (11). Figure A.3 plots (10) and (11) as functions of $\sigma_m$ and confirms that the linearization is very close to the nonlinear function for reasonable values of $\sigma_m$.
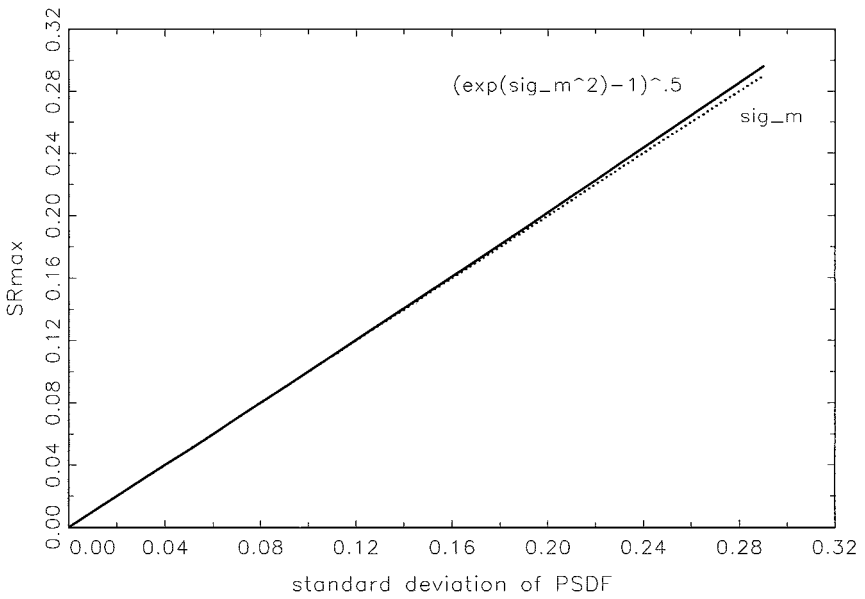
**FIGURE A.3.** Approximation error for Sharpe ratio.