A Note on the Hansen-Jagannathan Volatility Bounds

David Michael Rey
University of Basel/University of St. Gallen

and

Daniel Seiler
Investment Consulting Group/University of St. Gallen

March/April 2001

Authors:
David Michael Rey
University of Basel
WWZ/Department of Finance
Holbeinstrasse 12
4051 Basel
david.rey@unibas.ch
Phone: +41 (0)61 267 3307

Daniel Seiler
Investment Consulting Group
Winkelriedstrasse 19
9000 St. Gallen

daniel.seiler@investconsult.ch
Phone: +41 (0)71 227 6800

WWZ/Department of Finance, Working Paper No. 1/01

A Note on the Hansen-Jagannathan Volatility Bounds

DAVID MICHAEL REY and DANIEL SEILER*

Abstract

HANSEN/JAGANNATHAN (1991) show how to use security market data to restrict the admissible region for means and standard deviations of intertemporal marginal rates of substitution of consumers. They also characterise the duality between the mean-standard deviation frontier for valid stochastic discount factors and the traditional efficient frontier for asset returns. The intend of this paper, however, is to critically review this often proposed duality. On the one hand, emphasising the important role of the maximum-correlation portfolio, we argue that the duality is only valid in complete markets. Yet, in complete markets, the mean-variance frontier for valid discount factors reduces to a single point and the Hansen-Jagannathan volatility bounds are no longer properly defined. On the other hand, since the maximum correlation between a stochastic discount factor and a benchmark portfolio can not be inferred from only a small array of securities data, it is not possible to determine the 'degree of incompleteness' of the market. Therefore, the only opportunity is to draw 'complete markets volatility bounds', i.e. assuming perfect correlation.

I. INTRODUCTION

A major part of the research effort in finance is directed toward understanding why we observe a variety of financial assets with substantial and economically important average return differentials. A variety of asset pricing models have been proposed to explain this phenomenon. These describe how prices (and thus expected rates of return) of claims to future payoffs are determined in securities markets. Differences among the various asset pricing models arise from differences in their assumptions that restrict investors' preferences, endowments, production, information sets and the type of frictions allowed in the markets for real and financial assets. Still, most asset pricing models can be represented in a *stochastic discount factor* form. If assets earn different expected returns, then this stochastic discount factor (SDF) can not be a constant. Cross-sectional differences in expected asset returns thus carry implications for the variance of any valid SDF.

However, the SDFs corresponding to these models are particular parametric functions of the data observed by the econometrician. While empirical studies based on these parametric approaches have led to interesting insights, this parametric approach is usually based on strong assumptions about the underlying economic environment.

* David Michael Rey, lic. oec. HSG; Swiss Institute of Banking and Finance (s/bf), University of St. Gallen; Rosenbergstrasse 52, 9000 St. Gallen, Switzerland; Tel.: +41 (0)71 224 70 26, Email: david.rey@unisg.ch. Daniel Seiler, lic. oec. HSG, dipl. natw. ETH; Investment Consulting Group (ICGroup) AG; Winkelriedstrasse 19, 9000 St. Gallen, Switzerland; Tel.: +41 (0)71 227 68 13, Email: daniel.seiler@investconsult.com. Preliminary version. March/April 2001.

Alternatively, while completely *non-parametric* and assuming as little structure as possible, HANSEN/JAGANNATHAN (1991) show how the behaviour of a given set of asset prices can be used to restrict the admissible region for means and standard deviations of intertemporal marginal rates of substitution (IMRSs) of consumers. In order to calculate a *lower bound on the volatility of the SDF*, they only assume that financial markets obey the *law of one price* and that there are *no arbitrage opportunities*. These assumptions are sufficient to imply that there exists a SDF – which is (almost surely) positive. In other words, HANSEN/JAGANNATHAN (1991) show how to use the volatility bound as a *general diagnostic device*.

Additionally, they also characterise the duality between the mean-standard deviation frontier for valid SDFs and the traditional efficient frontier for asset returns. The intend of this paper, however, is to critically review this often proposed duality.

On the one hand, emphasising the important role of the *maximum-correlation portfolio*, we argue that the duality is *only* valid in complete markets. Yet, in complete markets, the mean-variance frontier for valid discount factors reduces to a single point and the Hansen-Jagannathan volatility bounds are no longer properly defined. On the other hand, since the maximum correlation between a SDF and a benchmark portfolio can not be inferred from only a small array of securities data, it is not possible to determine the '*degree of incompleteness*' of the market. Therefore, the only opportunity is to draw '*complete markets volatility bounds*', i.e. assuming perfect correlation.

Overall, we should therefore not rely on the duality between the standard mean-variance analysis and the mean-standard deviation frontier for valid SDFs. Instead, the Hansen-Jagannathan volatility bounds should better be computed *directly* according to the unique linear least squares approximation of every admissible SDF in the space of available asset payoffs.

The paper is organised as follows. Section 2 gives a brief overview of asset pricing, stochastic discount factors and commonalities of asset pricing models. Section 3 does not only derive the Hansen-Jagannathan bounds, but also stresses the important role of maximum-correlation portfolios and discusses their implication for the volatility bounds both in complete and incomplete markets. The often proposed duality is critically examined in Section 4. Section 5 gives a very brief review of further developments and applications of the Hansen-Jagannathan bounds. Section 6 concludes.

II. ASSET PRICING AND STOCHASTIC DISCOUNT FACTORS

Although there are differences among asset pricing models, they may not only be viewed as models of SDFs, but also to have important commonalities. All asset pricing models are based on one or more of three central concepts: The *law of one price*, the *no-arbitrage principle* and *financial market equilibrium*.

2.1. Commonalities of Asset Pricing Models

The first important commonality is the *law of one price*, according to which the prices of any two claims which promise the same future payoff must be the same.

The law of one price arises as an implication of the second concept, the *no-arbitrage principle*. The no-arbitrage principle states that market forces tend to align the prices of financial assets to eliminate arbitrage opportunities. Arbitrage opportunities arise when assets can be combined, by buying and selling, to form portfolios that have zero net cost, no chance of producing a loss, and a positive probability of gain. The law of one price follows from the no-arbitrage principle, when it is possible to buy or sell two claims to the same future payoff. If the two claims do not have the same price, and if transaction costs are smaller than the difference between their prices, then an arbitrage opportunity is created.

The third central concept behind asset pricing models is *financial market equilibrium*. Investors' desired holdings of financial assets are derived from an optimisation problem. A necessary condition for financial market equilibrium in a market with no frictions is that the first-order conditions of the investors' optimisation problem be satisfied. This requires that investors be indifferent at the margin to small changes in their asset holdings. Equilibrium asset pricing models follow from the first-order conditions for the investors' portfolio choice problem and from a market-clearing condition. The market-clearing condition states that the aggregate of investors' desired asset holdings must equal the aggregate '*market portfolio*' of securities in supply.

2.2. Stochastic Discount Factor and Euler Equation

It was already stated that virtually all financial asset pricing models may be viewed as models of SDFs. The proposition that there exists a SDF such that the expected product of any asset return with the SDF equals one holds very generally in models that rule out arbitrage opportunities. Equilibrium models with optimising investors imply tight links between the SDF and the marginal utilities of investors' consumption. Thus, by studying the SDF one can relate asset prices to the underlying preferences of investors.

Consumption-based asset pricing models aggregate investors into a single *representative agent*, who is assumed to derive utility from the aggregate consumption of the economy. In these models the SDF is the IMRS for the representative agent. The *Euler equations* – the first-order conditions for optimal consumption and portfolio choices of the representative agent – can be used to link asset returns and consumption.

Consider the *intertemporal choice problem* of an investor who can trade freely in asset *i* and who maximises the expectation of a *time-separable utility function*:

$$\text{Max } E_t\left[\sum_{j=0}^{\infty} \delta^j U\left(C_{t+j}\right)\right], \tag{2.1}$$

where δ is the discount factor, C_{t+j} is the investor's consumption in period $t+j$, and $U(C_{t+j})$ is the period utility of consumption at $t+j$. One of the first-order conditions or Euler equations describing the investor's optimal consumption and portfolio plan is

$$U'(C_t) = \delta E_t\left[(1+R_{i,t+1})U'(C_{t+1})\right]. \tag{2.2}$$

If we divide both the left- and right-hand sides of equation (2.2) by $U'(C_t)$, we get

$$1 = E_t\left[(1+R_{i,t+1})\delta\frac{U'(C_{t+1})}{U'(C_t)}\right] \equiv E_t\left[(1+R_{i,t+1})M_{t+1}\right], \tag{2.3}$$

with M_{t+1} the SDF or *pricing kernel*, equivalent to the IMRS in this model. Note that the IMRS, and hence the SDF, are always positive since marginal utilities are positive.

Expectations in equation (2.3) are taken conditional on information available at time t; however, by taking unconditional expectations and lagging one period, we obtain the following unconditional version,

$$1 = E\left[(1+R_{it})M_t\right]. \tag{2.4}$$

These relationships can be rearranged so that they explicitly determine (unconditional) expected asset returns as

$$E\left[(1+R_{it})\right] = \frac{1}{E[M_t]}\left(1-\mathrm{Cov}[R_{it}], M_t\right). \tag{2.5}$$

Note that if there is an asset whose unconditional covariance with the SDF is zero, then equation (2.5) implies that this asset's expected gross return is the reverse of the unconditional first moment of the SDF,

$$E\left[(1+R_{0t})\right] = \frac{1}{E[M_t]}. \tag{2.6}$$

Overall, this shows that an asset's expected return is greater, the smaller its covariance with the SDF. The intuition behind this result is that an asset whose covariance with M_{t+1} is small tends to have low returns when the investor's marginal utility of consumption is high – that is, when consumption itself is low. Such an asset is risky in that it fails to deliver 'wealth' precisely when 'wealth' is most valuable to the investor. The investor therefore demands a large risk premium to hold it.

Alternatively, equation (2.4) can also be derived merely from the *absence of arbitrage*, without assuming that investors maximise well-behaved utility functions. This can be shown in the familiar discrete setting with states s = 1,..., S, assets i = 1,..., N, $\mathbf{q}$ as the (Nx1) vector of asset prices, and $\mathbf{X}$ as an (SxN) matrix giving the payoffs of each asset in each state. Provided that all asset prices are non-zero, we can further define $\mathbf{G}$ as an (SxN) matrix giving the gross return on each asset in each state.

Now define an (Sx1) vector $\mathbf{p}$ to be a state price vector if it satisfies $\mathbf{X'p} = \mathbf{q}$. Thus, we can represent each asset price as the sum of its state-contingent payoffs times the appropriate states prices. Equivalently, if we divide through by **q**, we get **G'p** = **1**, where **1** is an (*S*x1) vector of ones.

An important result is that there exists a (strictly) positive state price vector if and only if there are no arbitrage opportunities (see e.g. DUFFIE (1996, p. 4). Furthermore, if there exists a positive state price vector, then equation (2.4) is satisfied for some positive random variable M. To see this define $M_s = p_s/\pi_s$, where is the probability of state *s*. For any asset *i* the relationship **G'p** = **1** implies

$$1=\sum_{s=1}^{S} p_s(1+R_{si})=\sum_{s=1}^{S} \pi_s M_s(1+R_{si})=E[(1+R_i)M]. \tag{2.7}$$

In the discrete-state model, asset markets are complete if for each state *s*, one can combine available assets to get a non-zero payoff in *s* and zero payoffs in all other states. A further important result is that the state price vector is unique if and only if asset markets are complete. In this case the SDF is unique, but with incomplete markets there may exist many SDFs satisfying equation (2.4). This result can be understood by considering an economy with several utility-maximising investors. The first-order condition (equation (2.2)) holds for each investor, so each investor's marginal utilities can be used to construct a SDF that prices the assets in the economy. With complete markets, the investors' marginal utilities are perfectly correlated so they all yield the same, unique SDF – with incomplete markets there may be idiosyncratic variation in marginal utilities and hence multiple SDFs that satisfy equation (2.2).

Note that the existence of a SDF is equivalent to the law of one price and that the restriction that the SDF is a strictly positive random variable is equivalent to a no-arbitrage condition. Certainly, as already mentioned above, the no-arbitrage principle does not uniquely identify the SDF unless markets are complete.

III. DERIVATION OF THE HANSEN-JAGANNATHAN BOUNDS

If price data were available from a complete set of security markets, the IMRSs of all consumers could be inferred from *Arrow-Debreu prices*. However, economic agents may not trade in a complete set of contingent-claims markets. Furthermore, it may be practical for an econometrician to use data on only a small array of securities. Because of these limitations, asset market data alone are typically not sufficient to identify IMRSs.

While completely *non-parametric* and assuming as little structure as possible, HANSEN/JAGANNATHAN (1991) show how the behaviour of a given set of asset prices can be used to restrict the admissible region for means and standard deviations of SDFs.

3.1. THE HANSEN-JAGANNATHAN BOUNDS

We thus now ask what asset return data may be able to tell us about the behaviour of the SDF. In order to calculate a *lower bound on the volatility of the SDF*, in what follows we only assume that financial markets obey the *law of one price*. This assumption is sufficient to imply that there *exists* a SDF, but we can not say anything about *positivity* and *uniqueness* of it (which of course is not surprising, since the interesting small array of securities data is very likely to imply an incomplete market setting). In addition, to simplify the exposition, we focus on an *unconditional version of the volatility bound* using only unconditional expectations. We also posit a hypothetical, unconditionally risk-free asset. We take the value of the risk-free asset, or equivalently, the reverse of the unconditional mean of the SDF (see equation (2.6)), as a parameter to be varied as we trace out the bound. We will see, however, that varying the value of the risk-free asset is not very appropriate in complete markets, since the sum of all contingent claim prices completely and uniquely determines the value of this risk-free asset. Strictly speaking, thus, there is no room for a variation of the value of the risk-free asset in a complete market setting.

HANSEN/JAGANNATHAN (1991) begin their analysis with the unconditional version of equation (2.3) and rewrite it in vector form as[1]

$$\mathbf{1} = \mathrm{E}\left[(\mathbf{1}+\mathbf{R}_t)\mathrm{M}_t\right], \tag{3.1}$$

where $\mathbf{R}_t$ is the *N*-vector of time-*t* asset returns and, by assumption, has a non-singular variance-covariance matrix $\boldsymbol{\Omega}$ (i.e., no asset or combination of assets is unconditionally riskless). There may still exist an unconditional zero-beta asset with gross mean return equal to the reciprocal of the unconditional mean of the SDF, but they assume that if there is such an asset its identity is not known a priori. Hence they treat the unconditional mean of the SDF as an unknown parameter,

$$\overline{\mathrm{M}} = \mathrm{E}\left[\mathrm{M}_t^*(\overline{\mathrm{M}})\right] = \mathrm{E}\left[\mathrm{M}_t(\overline{\mathrm{M}})\right]. \tag{3.2}$$

For each possible $\overline{\mathrm{M}}$, they form a candidate SDF $\overline{\mathrm{M}}_t^*(\overline{\mathrm{M}})$ as linear combination of asset returns

$$\mathrm{M}_t^*(\overline{\mathrm{M}}) = \overline{\mathrm{M}} + (\mathbf{R}_t - \mathrm{E}[\mathbf{R}_t])'\beta_{\overline{\mathrm{M}}}, \tag{3.3}$$

and show that the variance of $\overline{\mathrm{M}}_t^*(\overline{\mathrm{M}})$ places a lower bound on the variance of any SDF that has mean $\overline{\mathrm{M}}$ and satisfies equation (3.1).

If $\overline{\mathrm{M}}_t^*(\overline{\mathrm{M}})$ is to be a SDF, it must satisfy equation (3.1), too,

$$\mathbf{1} = \mathrm{E}\left[(\mathbf{1}+\mathbf{R}_t)\mathrm{M}_t^*(\overline{\mathrm{M}})\right]. \tag{3.4}$$

Expanding the expectation of the above product, we have

$$\begin{aligned}
\mathbf{1} &= \overline{\mathrm{M}}\mathrm{E}[(\mathbf{1}+\mathbf{R}_t)] + \mathrm{Cov}\left[\mathbf{R}_t, \mathrm{M}_t^*(\overline{\mathrm{M}})\right] \\
\mathbf{1} &= \overline{\mathrm{M}}\mathrm{E}[(\mathbf{1}+\mathbf{R}_t)] + \mathrm{E}\left[(\mathbf{R}_t - \mathrm{E}[\mathbf{R}_t])(\mathrm{M}_t^*(\overline{\mathrm{M}}) - \overline{\mathrm{M}})\right] \\
\mathbf{1} &= \overline{\mathrm{M}}\mathrm{E}[(\mathbf{1}+\mathbf{R}_t)] + \mathrm{E}\left[(\mathbf{R}_t - \mathrm{E}[\mathbf{R}_t])(\mathbf{R}_t - \mathrm{E}[\mathbf{R}_t])'\beta_{\overline{\mathrm{M}}}\right]
\end{aligned} \tag{3.5}$$

[1] The following notation is based on CAMPBELL/LO/MACKINLAY (1997, Chapter 8).

$$\mathbf{1} = \overline{M}E[(\mathbf{1}+\mathbf{R}_t)] + \mathbf{\Omega}\beta_{\overline{M}},$$

with $\mathbf{\Omega}$ defined as above. It follows then that

$$\beta_{\overline{M}} = \mathbf{\Omega}^{-1}(\mathbf{1} - \overline{M}E[\mathbf{1}+\mathbf{R}_t]) \tag{3.6}$$

and the variance of the implied SDF is

$$Var[M_t^*(\overline{M})] = \beta'_{\overline{M}}\mathbf{\Omega}\beta_{\overline{M}} = (\mathbf{1} - \overline{M}E[\mathbf{1}+\mathbf{R}_t])'\mathbf{\Omega}(\mathbf{1} - \overline{M}E[\mathbf{1}+\mathbf{R}_t]). \tag{3.7}$$

The right-hand side of equation (3.7) is a *lower bound on the volatility of any stochastic SDF* with mean $\overline{M}$. To see this, note that any other $\overline{M}_t(\overline{M})$ satisfying (3.1), and defined as

$$M_t(\overline{M}) = M_t^*(\overline{M}) + \varepsilon_t \quad \text{with} \quad E[\varepsilon_t] = 0 \tag{3.8}$$

must have the property

$$E[(\mathbf{1}+\mathbf{R}_t)(M_t(\overline{M}) - M_t^*(\overline{M}))] = Cov[\mathbf{R}_t, M_t(\overline{M}) - M_t^*(\overline{M})] = 0. \tag{3.9}$$

Since $\overline{M}_t^*(\overline{M})$ is just a linear combination of asset returns, it follows that

$$Cov[M_t(\overline{M}) - M_t^*(\overline{M}), \varepsilon_t] = 0. \tag{3.10}$$

Thus,

$$\begin{aligned} Var[M_t(\overline{M})] &= Var[M_t^*(\overline{M})] + Var[\varepsilon_t] + 2Cov[M_t^*(\overline{M}), \varepsilon_t] \\ Var[M_t(\overline{M})] &= Var[M_t^*(\overline{M})] + Var[\varepsilon_t] \\ Var[M_t(\overline{M})] &\geq Var[M_t^*(\overline{M})]. \end{aligned} \tag{3.11}$$

In fact, we can go beyond the inequality in equation (3.11) to show that

$$Var[M_t(\overline{M})] = \frac{Var[M_t^*(\overline{M})]}{(Corr[M_t^*(\overline{M}), M_t(\overline{M})])^2}, \tag{3.12}$$

so a SDF can only have a variance close to the lower bound if it is highly correlated with the combination of asset returns $\overline{M}_t^*(\overline{M})$.

3.2. BENCHMARK AND MAXIMUM-CORRELATION PORTFOLIOS

We can restate these results in a more familiar way by introducing the idea of a *benchmark portfolio*, which we may define as

$$1 + R_{bt}(\overline{M}) \equiv \frac{M_t^*(\overline{M})}{E[M_t^*(\overline{M})^2]}. \tag{3.13}$$

It is straightforward to check that this return satisfies equation (3.1) on returns and has – among others – the following two properties:

- It is *mean-variance efficient*. That is, no other portfolio has smaller variance and the same mean.
- Any stochastic discount factor $\overline{M}_t(\overline{M})$ has a greater correlation with R_{bt} than with any other

portfolio. For this reason, R_{bt} is sometimes referred to as a *maximum-correlation portfolio.*

Very importantly, note that according to DUFFIE (1996, p. 12), *if markets are complete*, then the benchmark return R_{bt} is *perfectly correlated* with the state-price deflator. In this complete market case, we shall call the benchmark portfolio '*market portfolio*' with return R_{mt}. Otherwise, in incomplete markets, the benchmark return R_{bt} is *less* than perfectly correlated with the state-price deflator. In summary, we thus have

$$\left| \operatorname{Corr}\left[R_{bt}, M_t\left(\overline{M}\right)\right] \right| < 1 \tag{3.14}$$

for any *benchmark portfolio*, and

$$\left| \operatorname{Corr}\left[R_{mt}, M_t\left(\overline{M}\right)\right] \right| = 1 \tag{3.15}$$

for the *market portfolio*.

Equivalently, since $\overline{M}_t^*\left(\overline{M}\right)$ is simply a combination of asset returns, we may write

$$\left| \operatorname{Corr}\left[M_t^*\left(\overline{M}\right), M_t\left(\overline{M}\right)\right] \right| < 1 \tag{3.14'}$$

for any *benchmark portfolio*, and

$$\left| \operatorname{Corr}\left[M_t^*\left(\overline{M}\right), M_t\left(\overline{M}\right)\right] \right| = 1 \tag{3.15'}$$

for the *market portfolio* instead of equation (3.14) and (3.15), respectively. In the latter case of the market portfolio, we thus have

$$\operatorname{Var}\left[M_t\left(\overline{M}\right)\right] = \frac{\operatorname{Var}\left[M_t^*\left(\overline{M}\right)\right]}{\left(\operatorname{Corr}\left[M_t^*\left(\overline{M}\right), M_t\left(\overline{M}\right)\right]\right)^2} = \operatorname{Var}\left[M_t^*\left(\overline{M}\right)\right], \tag{3.12'}$$

indicating the *lower bound on the volatility of any stochastic SDF* with mean $\overline{M}$ (see equation (3.7)). Consequently, this lower bound is a '*complete market volatility bound*'.

A second look at equation (3.11) gives an alternative way to see this. Only in the case that

$$\operatorname{Var}\left[\varepsilon_t\right] = 0 \tag{3.16}$$

we have the equality

$$\operatorname{Var}\left[M_t\left(\overline{M}\right)\right] = \operatorname{Var}\left[M_t^*\left(\overline{M}\right)\right], \tag{3.17}$$

again indicating the lower bound on the volatility of any stochastic SDF with mean $\overline{M}$. In other words, if the variance of the random variable ε is zero, we are again in *a complete market setting*:

$$\left| \operatorname{Corr}\left[M_t^*\left(\overline{M}\right), M_t^*\left(\overline{M}\right) + \varepsilon_t\right] \right| = 1 \text{ if and only if } \operatorname{Var}\left[\varepsilon_t\right] = 0. \tag{3.18}$$

In contrast, in the case that we have

$$\left| \operatorname{Corr}\left[M_t^*\left(\overline{M}\right), M_t^*\left(\overline{M}\right) + \varepsilon_t\right] \right| < 1 \text{ if and only if } \operatorname{Var}\left[\varepsilon_t\right] > 0, \tag{3.19}$$

markets are *incomplete*.

It may be interesting to note that the random variable ε has *no pricing effect*. Since both $\overline{M}_t^*\left(\overline{M}\right)$ and $\overline{M}_t\left(\overline{M}\right)$ are a valid SDF, we can write

$$1 = E\left[(1+R_{it})M_t^*(\overline{M})\right] = E\left[(1+R_{it})M_t(\overline{M})\right], \tag{2.4'}$$

despite that

$$Var\left[M_t(\overline{M})\right] > Var\left[M_t^*(\overline{M})\right] \text{ for } Var\left[\varepsilon_t\right] > 0 \text{ and } E\left[(R_{it}\varepsilon_t)\right] = 0 .$$

This is therefore surprising as the standard deviation of the SDF is often interpreted as the *market price of risk*. Thus, the conclusion might be that if one increases the variance of the random variable ε (of course, this is only possible in incomplete markets), one does also increases the market price of risk. As a consequence, this would also increase expected returns or, equivalently, lower prices). Unless in complete markets, however, the standard deviation of the SDF should not generally be interpreted as the market price of risk. Instead, one should interpret the *Sharpe ratio* such as given in equation (4.1) as the market price of risk. In doing so, it is directly evident from equation (4.4) that if one increases the variance of the random variable ε, one does certainly also increase the standard deviation of the SDF, but, at the same time, decrease the correlation coefficient between the returns of the benchmark portfolio and the SDF. Thus, both the Sharpe ratio in equation (4.1) and the so-defined market price of risk are not affected. This is why the random variable ε has *no pricing effect*.

3.3. IMPLICATIONS FOR THE HANSEN-JAGANNATHAN BOUNDS

Hence, the overall conclusion in Section 3.2 was that the Hansen-Jagannathan volatility bounds, as they are usually drawn (i.e. according to equation (3.7)), are '*complete market volatility bounds*'. But does a 'complete market volatility bound' actually make sense? The answer is straightforward: No, it does not. In complete markets, any valid SDF is unique and thus has a single mean and standard deviation, which is completely known. Consequently, in complete markets, the mean-variance frontier for valid stochastic discount factors reduces to a single point and the Hansen-Jagannathan volatility bounds are no longer properly defined.

But how would the Hansen-Jagannathan volatility bounds look like in an *incomplete* market setting? We simply do not know. Of course, COCHRANE/HANSEN (1992) refine the Hansen-Jagannathan bounds to consider information about the correlation between a given stochastic discount factor and the vector of asset returns and therefore provide a tighter set of restrictions that the 'complete market volatility bounds'. In doing so, we must rewrite equation (3.3) – which anyway only holds in complete markets –, as

$$M_t(\overline{M}) = \overline{M} + (\mathbf{R}_t - E[\mathbf{R}_t])'\beta_{\overline{M}} + \varepsilon_t , \tag{3.20}$$

which may be interpreted as regressing $\overline{M}_t(\overline{M})$ on the vector of asset returns. COCHRANE/HANSEN (1992) now show that the following equation holds:

$$Var\left[M_t(\overline{M})\right] = \frac{Var\left[M_t^*(\overline{M})\right]}{R^2} , \tag{3.21}$$

where R^2 denotes the *goodness of fit*. This relation reveals a clear tradeoff: If a candidate SDF is less than perfectly correlated with the return vector, it must be more volatile than implied by the 'complete market volatility bounds'. Thus, any valid SDF must either be highly correlated with asset returns, or have a higher standard deviation than indicated by the 'complete market volatility bound' for a given set of assets. But if a SDF that is completely uncorrelated with asset returns will fail any pricing test. Put it slightly different, a candidate SDF may be very volatile and perfectly fall into the 'complete market volatility bound', but still price miserably if it is uncorrelated with the returns on a given set of assets.

Figure 3.1: The 'correlation puzzle'

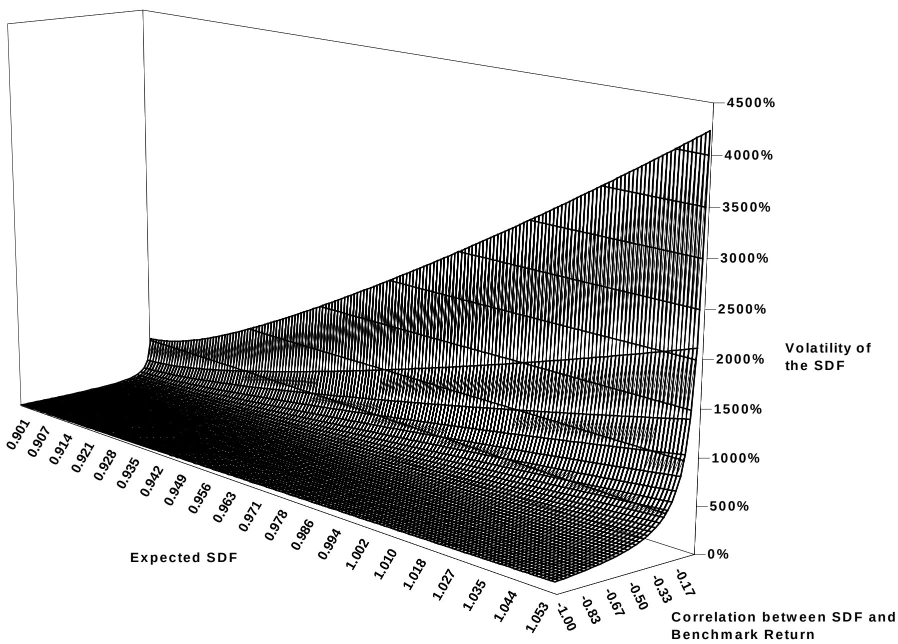

Source: Own calculations.

However, without direct data on the SDF, the regression coefficients in equation (3.20) can not be estimated. In other words, given only a small array of securities data, we can not infer the (maximum) correlation between the benchmark portfolio and any valid SDF. Since the variance of any valid SDF only depends on the second moment matrix of the *N* asset returns (see equation (3.7)), the lower volatility bound depends only on the assets available to the econometrician (and definitely not on any particular asset pricing model may being studied). Given this space of available payoffs, it is not possible to determine the '*degree of incompleteness*' of the market. Therefore, the only opportunity is to draw 'complete markets volatility bounds', i.e. assuming perfect correlation.

IV. Duality between Efficient Frontier and HJ Bounds

It is often stated that the lower bound on the volatility of a stochastic discount factor is closely related to the *standard mean-variance analysis* that has long been used in the financial economics literature.

To see this, we start again with equation (3.1) and rewrite it for a benchmark portfolio,

$$1 = E[(1 + R_{bt})M_t]. \tag{4.1}$$

With r_{bt} the *excess* return of the benchmark portfolio, we have

$$0 = E[r_{bt}M_t]. \tag{4.2}$$

Using the definition of a covariance, equation (4.2) can be rewritten as

$$0 = \rho\,\sigma[r_{bt}]\sigma[M_t] + E[(r_{bt})]E[M_t], \tag{4.3}$$

with ρ the correlation coefficient between the (excess) return of the benchmark portfolio and the SDF (see equation (3.14)). Some simple manipulations result in the equality

$$\frac{E[r_{bt}]}{(-\rho)\sigma[r_{bt}]} = \frac{E[R_{bt}] - R_{0t}}{(-\rho)\sigma[R_{bt}]} = \frac{\sigma[M_t]}{E[M_t]}. \tag{4.4}$$

To investigate this link between the standard mean-variance analysis and the Hansen-Jagannathan bounds, we first consider the case of *complete markets*.

4.1. Complete Markets

It was already shown in equation (3.15) that the SDF and the return of the benchmark portfolio is *perfectly negative* correlated. We therefore called the benchmark portfolio '*market portfolio*'. We can thus rewrite equation (4.4) as

$$\frac{E[R_{mt}] - R_{0t}}{\sigma[R_{mt}]} = \frac{\sigma[M_t]}{E[M_t]}. \tag{4.5}$$

It was also shown that any benchmark portfolio (thus including the market portfolio) is *mean-variance efficient*.

The left-hand side of this expression is the *Sharpe ratio* for the market portfolio. This is well-known from the field of *portfolio theory*: The Sharpe ratio is the maximal slope of the line drawn from the risk-free rate of return and tangent to the traditional minimum-variance boundary at the mean-variance efficient market portfolio. At least at a first glance, we thus have a simple correspondence between the Hansen-Jagannathan bounds and the traditional minimum variance boundary for the given assets. But does this duality make sense?

Recall that both the tangent portfolio (i.e. the market portfolio) and the Sharpe ratio depend on the value of the risk-free asset. It was already mentioned above, in complete markets, the value of the risk-free asset as well as the SDF are fully and uniquely determined. This has two implications: First, it is not very appropriate to take the value of the risk-free asset, or equivalently, the reverse of the unconditional mean of the SDF, as a parameter to be varied as we trace out the (complete market) volatility bound. Second, in complete markets, the mean-variance frontier for valid discount factors collapses to a single point and the Hansen-Jagannathan volatility bounds are no longer properly defined.

In what follows, we investigate whether the often proposed duality also holds in an *incomplete market setting*.

4.2. INCOMPLETE MARKETS

First of all, note that even in incomplete markets it is possible to calculate and draw a mean-variance efficient frontier. Of course, this is what we usually have: Given only a *small* array of securities data, an incomplete market setting is the most likely case we are exposed to. But can we interpret the 'Sharpe ratio' in the same way?

To see this, recall that in incomplete markets the correlation between the returns of the benchmark portfolio and the SDF is no longer perfect. There is thus no way to simplify equation (4.4)

$$\frac{E[r_{bt}]}{(-\rho)\sigma[r_{bt}]} = \frac{E[R_{bt}] - R_{0t}}{(-\rho)\sigma[R_{bt}]} = \frac{\sigma[M_t]}{E[M_t]}. \tag{4.4}$$

Certainly, the above analysis and Figure 3.1 have already shown that the lower the correlation coefficient, the higher the volatility of the SDF (note that the expected value of the SDF does not change). It is also evident from equation (4.4) that the lower the correlation is, the higher is the 'Sharpe ratio', i.e. the value of the expression on the left-hand side. Recall that a properly defined Sharpe ratio is the maximal slope of the line drawn from the risk-free rate of return through the mean-variance efficient benchmark portfolio. Given our interesting (small) array of securities data, we can only interpret

$$\frac{E[R_{bt}] - R_{0t}}{\sigma[R_{bt}]} \tag{4.6}$$

as a Sharpe ratio. There is thus no way to take the correlation coefficient into account in a traditional mean-variance setting. Of course, it is possible to *assume* that the correlation is perfect and trace out the volatility bound, but then we are of course back in the complete market setting described in the previous section. Thus, despite that in an incomplete market setting it would be appropriate to take the value of the risk-free asset as a parameter to be varied as we trace out the volatility bound, *the often proposed duality breaks down*. We should therefore not rely on the duality between the standard mean-variance analysis and the mean-standard deviation frontier for valid stochastic discount. Instead, the Hansen-Jagannathan volatility bounds should better be computed *directly* according to the unique linear least squares approximation of every admissible stochastic discount factor in the space of available asset payoffs (see equation (3.7)).

V. EXTENSIONS AND APPLICATIONS

Additional extensions, refinements and applications of the Hansen-Jagannathan volatility bounds have become manifold in recent years.

5.1. IMPLICATIONS OF NON-NEGATIVITY

So far, based on equation (3.1) – which is equivalent to the law of one price –, we have ignored the restriction that the SDF must be non-negative. If there are no arbitrage opportunities, it implies that the SDF is a strictly positive random variable. Already HANSEN/JAGANNATHAN (1991) show how to obtain a tighter bound on the standard deviation of the SDF by making use of the restriction that there are no arbitrage opportunities. However, computing the bounds imposing positivity requires a *numerical search procedure*. Their results indicate that the bounds imposing positivity are nearly coincident with the simpler bounds in the portion of the parabola where the standard deviation is low, and differ much from the simpler bounds when the standard deviation is relatively high.

5.2. CONDITIONAL MOMENTS

To simplify the exposition, the above analysis is solely based on an *unconditional* version of the volatility bound using only unconditional expectations – which obviously is leaving out information. HANSEN/JAGANNATHAN (1991) also show how to incorporate conditioning variables into the analysis. Put it simple, principally everything could have been stated for conditional means and variances, too.

5.3. FURTHER DEVELOPMENTS

Applications of the Hansen-Jagannathan volatility bounds now include explicit *market integration tests* (CHEN/KNEZ (1995), *mutual fund performance measurement* (CHEN/KNEZ (1996), FERSON/SCHADT (1996), FARNSWORTH/FERSON/JACKSON/TODD (1998)), the use of the Hansen-Jagannathan regions to perform *mean-variance spanning tests* (DE SANTIS (1995), BEKAERT/URIAS (1996)) and incorporate transaction costs (LUTTMER (1993), HE/MODEST (1992) and COCHRANE/HANSEN (1992)).

Additionally, SNOW (1991) extends the Hansen-Jagannathan volatility bounds to include *higher moments* of the asset returns. COCHRANE/HANSEN (1992), BURNSIDE (1994), and CECCHETTI/LAM/MARK (1994) show how to take *sampling errors* into account when examining whether a particular candidate SDF satisfies the Hansen-Jagannathan bounds. Since the methods we have examined so far are developed, for the most part, under the null hypothesis that the asset pricing model under consideration by the econonometrician assigns the right prices to all assets. An alternative is to assume that the model is wrong and examine how wrong the model is. HANSEN/JAGANNATHAN (1994) discuss a possible way to examine what is missing in a model and assign a scalar measure of the model's specification.

5.4. THE EQUITY PREMIUM PUZZLE

Using long-run annual US data, two stylised facts are commonly emphasised. Firstly, the average excess return on US stocks over short-term debt – notably the equity premium – is as high as about 6% per year. Secondly, aggregate consumption is very smooth, so covariances with consumption growth are small. Putting these facts together, the power utility model can only fit the equity premium if the coefficient of relative risk aversion is very large (i.e. too large to be plausible). This is the *equity premium puzzle* of MEHRA/PRESCOTT (1985). The Hansen-Jagannathan volatility bounds provide an alternative characterisation of this puzzle. Their framework gives a clear answer to why the (basic) consumption-based models fail: Consumption rates are simply not volatile enough and too little correlated with asset returns to explain all restrictions of the standard valuation model.

VII. CONCLUSION

HANSEN/JAGANNATHAN (1991) show how to use security market data to restrict the admissible region for means and standard deviations of intertemporal marginal rates of substitution of consumers. They also characterise the duality between the mean-standard deviation frontier for valid stochastic discount factors and the traditional efficient frontier for asset returns.

However, we critically reviewed this often proposed duality. On the one hand, we emphasised the important role of the maximum-correlation portfolio and argued that the duality is only valid in complete markets. Yet, in complete markets, the mean-variance frontier for valid discount factors reduces to a single point and the Hansen-Jagannathan volatility bounds are no longer properly defined. On the other hand, since the maximum correlation between a SDF and a benchmark portfolio can not be inferred from only a small array of securities data, it is not possible to determine the '*degree of incompleteness*' of the market. Therefore, the only opportunity is to draw '*complete markets volatility bounds*', i.e. assuming perfect correlation.

Overall, we should therefore not rely on the duality between the standard mean-variance analysis and the mean-standard deviation frontier for valid SDFs. Instead, the Hansen-Jagannathan volatility bounds should better be computed *directly* according to the unique linear least squares approximation of every admissible SDF in the space of available asset payoffs.

BIBLIOGRAPHY

CAMPBELL, J.Y. (2000): 'Asset Pricing at the Millennium'; *The Journal of Finance*, Vol. LV, No. 4, 1515-1567.

CAMPBELL, J.Y./LO, A.W./MACKINLAY, A.C. (1997): *The Econometrics of Financial Markets*; Princeton University Press, Princeton, New Jersey.

COCHRANE, J.H./HANSEN, L.P. (1992): 'Asset pricing explorations for macroeconomics'; *NBER Working Paper*, No. 4088, (June).

DROBETZ, W. (2000): *Global Stock Markets*; Gabler Edition Wissenschaft, Wiesbaden.

DUFFIE, D. (1996): *Dynamic Asset Pricing Theory*; 2nd edition, Princeton University Press, Princeton, New Jersey.

FERSON, W.E. (1995): 'Theory and Empirical Testing of Asset Pricing Models'; *Handbooks in OR & MS* (JARROW, R. et al, Eds.), Vol. 9, 145-200.

FERSON, W.E./JAGANNATHAN, R. (1996): 'Econometric Evaluation of Asset Pricing Models'; *Handbooks of Statistics* (MADDALA, G.A. and RAO, C.R., Eds.), Vol. 14.

HANSEN, L.P./JAGANNATHAN, R. (1991): 'Implications of Security Market Data for Models of Dynamic Economies'; *Journal of Political Economy*, Vol. 99, No. 2, 225-262.

MEHRA, R./PRESCOTT, E.C. (1985): 'The Equity Premium: A Puzzle'; *Journal of Monetary Economics*, 15, 145-161.