Portfolio Efficiency and Discount Factor Bounds With Conditioning Information: A Unified Approach[1]

Abhay Abhyankar
Devraj Basu
Alexander Stremme

Warwick Business School, University of Warwick

August 2002

address for correspondence:

Alexander Stremme
Warwick Business School
University of Warwick
Coventry, CV4 7AL
United Kingdom

e-mail: `a.stremme@wbs.warwick.ac.uk`

[1] We thank Süleyman Basak, Darrell Duffie, Wayne Ferson, and Harald Uhlig for many helpful comments and discussions. We would also like to thank seminar participants at the Financial Markets Group, LSE, the Said Business School, University of Oxford, Imperial College, London, and at the University of York and Exeter University for stimulating discussions. All remaining errors are ours.

Portfolio Efficiency and Discount Factor Bounds With Conditioning Information: A Unified Approach

August 2002

Abstract

In this paper, we develop a unified framework for the study of mean-variance efficiency and discount factor bounds in the presence of conditioning information. We extend the Hilbert space framework of Hansen and Richard (1987) to obtain new characterizations of the efficient portfolio frontier and variance bounds on discount factors, as functions of the conditioning information. We introduce a covariance-orthogonal representation of the asset return space, which allows us to derive several new results, and provide a portfolio-based interpretation of existing results. Our analysis is inspired by, and extends the recent work of Ferson and Siegel (2001,2002), and Bekaert and Liu (2001). Our results have several important applications in empirical asset pricing, such as the construction of portfolio-based tests of asset pricing models, conditional measures of portfolio performance, and tests of return predictability.

JEL Classification: G11, G12

1 Introduction

In this paper, we develop a unified framework for the study of mean-variance efficiency and discount factor bounds in the presence of conditioning information. Stochastic discount factor (SDF) bounds are central in testing asset pricing models. Such bounds define the feasible region in the mean-variance plane by providing a lower bound on the variance of admissible SDFs. Recent studies have focused on the use of conditioning information to refine these bounds. Since by duality, discount factor bounds are directly related to the mean-variance efficient frontier, studying the optimal use of conditioning information in the construction of managed portfolios is hence of central importance. The optimal use of conditioning information is likely to enlarge the opportunity set available to an investor, in contrast to the ad hoc multiplicative use prevalent in the literature. The study of portfolio efficiency with conditioning information, and thus the construction of managed portfolios that utilize such information optimally, is hence of independent interest. Our results extend and complement the existing literature in many important ways, and have several theoretical implications and empirical applications, including the construction of performance measures and tests of asset pricing models.

The main contribution of this paper is two-fold; First, we develop a new portfolio-based framework for the implementation of discount factor bounds, with and without conditioning information. We do this by constructing a new, covariance-orthogonal parameterization of the space of returns on managed portfolios that permits us to derive a general expression for such bounds. Our results connect various different approaches to the construction of such bounds, and allow a direct comparison of their respective properties. In particular, we provide a direct proof of the Gallant, Hansen, and Tauchen (1990) bounds, and an explicit expression for the 'unconditionally efficient' bounds of Ferson and Siegel (2002). Our intuitive decomposition of the bounds, which is new even in the fixed-weight case without conditioning information, yields an interesting interpretation of the original Hansen and Jagannathan (1991) bounds.

Second, to operationalize our theoretical results, we explicitly construct the weights of efficiently managed portfolios, as functions of the conditioning information. While for a specific class of portfolios, these weights have also been reported by Ferson and Siegel (2001), our solutions are more general. Our expressions enable us to characterize the optimal portfolios that maximize unconditional Sharpe ratios and thus attain the sharpest possible discount factor bounds. Moreover, our formulation of the weights facilitates the analysis of their behavior in response to changes in conditioning information. This may have important implications for the statistical properties of tests derived from such portfolios, and for the empirical properties of performance measures.

Mean-variance efficiency, together with the stochastic discount factor approach, are at the heart of modern empirical asset pricing, (see Ferson (2002) for a discussion). Mean-variance theory has found numerous applications, for example in portfolio analysis and asset allocation, empirical tests of asset pricing models, measurement of portfolio performance, and many other fields. The Hilbert space approach to mean-variance theory, pioneered by Chamberlain and Rothschild (1983), provides an elegant and powerful alternative to traditional mean-variance optimization. Hansen and Richard (1987) extend this framework to study the optimal use of conditioning information, which is of increasing importance, given the evidence for asset return predictability, (see Cochrane 1999).

Our work is related to Gallant, Hansen, and Tauchen (1990) (GHT), and Bekaert and Liu

(2001). GHT were the first to use conditioning information to improve the variance bounds for asset pricing models by projecting the SDF unconditionally onto the infinite-dimensional space of 'managed' pay-offs, and calculating the variance of this projection. Bekaert and Liu (2001) provide an alternative implementation of the GHT bounds by finding an optimal transformation of the conditioning instruments which maximizes the implied hypothetical Sharpe ratio. Our methodology allows us to characterize the efficient frontier in their setting, thus recovering the expression for their bounds, and identifying the portfolio which attains them.

Our work is also related to Ferson and Siegel (2001), who characterize the unconditional frontier of 'conditional' returns, stating the weights of efficient portfolios for multiple assets. Our relative contribution is to provide a constructive derivation of these weights, and a theoretical investigation of their behavior. The numerical results reported by Ferson and Siegel (2001) indicate that in their setting, the weights in the case with risk-free asset display a conservative response to extreme values of the conditioning instruments. Our analysis provides a theoretical explanation for this, even in the case without risk-free asset.

Ferson and Siegel (2002) use their characterization of the efficient frontier to construct portfolio-based bounds for discount factors, which they refer to as 'unconditionally efficient (UE)' bounds. Our contribution is to provide an explicit expression for these bounds in terms of their efficient set constants, as a simple application of our general result. In contrast, their construction is purely numerical, based on parameterizing the frontier in terms of the global minimum variance portfolio (GMV) and another, arbitrarily chosen portfolio. Our analysis provides a theoretical basis for these choices.

The remainder of this paper is organized as follows; In Section 2, we provide an overview of the relevant asset pricing theory. In Section 3 we develop our main theoretical results. In the subsequent two Sections, we operationalize these by explicitly characterizing the weights of efficient portfolios. Section 4 focuses on the case without risk-free asset, and provides an analysis of the implied discount factor bounds, while Section 5 covers the case with risk-free asset. Section 6 concludes.

2 Theoretical Background

In this section, we provide a brief outline of the underlying asset pricing theory, and establish our notation. Our exposition is based largely on Hansen and Richard (1987), and Cochrane (2001). We first construct the space of state-contingent pay-offs, and within it the space of traded pay-offs, augmented by the use of conditioning information. Next, we define two different notions of 'return' within this space.

2.1 Set-Up and Notation

We begin by constructing the space of state-contingent pay-offs. Fix a probability space $(\Omega, \mathcal{F}, P)$, endowed with a discrete-time filtration $(\mathcal{F}_t)_t$. We fix $t > 0$, and consider the period beginning at time $t-1$ and ending at t. Denote by L_t^2 the space of all $\mathcal{F}_t$-measurable random variables that are (unconditionally) square-integrable with respect to P. Obviously, L_t^2 is a Hilbert space with respect to the inner product

$$\langle x_t, y_t \rangle := E\big(x_t \cdot y_t \big), \qquad \text{for } x_t, y_t \in L_t^2. \tag{1}$$

To incorporate conditioning information, let $\mathcal{G}_{t-1} \subseteq \mathcal{F}_t$ be a sub-σ-field. We think of $\mathcal{G}_{t-1}$ as summarizing all information on which investors base their portfolio decisions at time $t-1$. In particular, asset prices at time $t-1$ will typically depend on $\mathcal{G}_{t-1}$. In most practical applications, $\mathcal{G}_{t-1}$ will be chosen as the σ-field generated by a set of conditioning *instruments*, variables observable at time $t-1$ that contain information about the distribution of asset returns. Examples considered in the literature include dividend yield (Fama and French 1988), interest rate spreads (Campbell 1987), or consumption-wealth ratio (Lettau and Ludvigson 2001).

Hansen and Richard (1987) show that L_t^2 can be made into a *conditional* Hilbert space, when the inner product is replaced by its conditional version. $E\big(\, x_t \cdot y_t \mid \mathcal{G}_{t-1} \,\big)$. Intuitively, we may think of a conditional Hilbert space as a family of regular, finite-dimensional Hilbert spaces, indexed by the outcomes of the conditioning information. Let $X_t \subseteq L_t^2$ be a (conditionally) complete, linear subspace of L_t^2. In particular, X_t is closed with respect to linear combinations of pay-offs with 'weights' that are $\mathcal{G}_{t-1}$-measurable functions. We interpret X_t as the space of all pay-offs that are attainable by forming 'managed' portfolios from traded assets. As a closed linear subspace of a Hilbert space, X_t is itself also a Hilbert space.

We denote by $\Pi_{t-1} : X_t \longrightarrow \mathcal{G}_{t-1}$ the conditional *pricing rule*, mapping pay-offs $x_t \in X_t$ into their conditional price, $\Pi_{t-1}\big(\, x_t \,\big)$. Note that $\Pi_{t-1}\big(\, x_t \,\big)$ is not constant, but itself a $\mathcal{G}_{t-1}$-measurable function. We assume that the pricing rule satisfies the *'Law of One Price'*, i.e. that Π_{t-1} is a continuous and (conditionally) linear mapping. From the conditional version of the Riesz representation theorem it then follows that there exists a unique element $x_t^* \in X_t$ such that

$$\Pi_{t-1}\big(\, x_t \,\big) = E\big(\, x_t^* \cdot x_t \mid \mathcal{G}_{t-1} \,\big) \qquad \text{for all } x_t \in X_t. \tag{2}$$

The element x_t^* is referred to as the stochastic discount factor (SDF) induced by Π_{t-1}.

2.2 Two Notions of 'Return'

We now introduce two different notions of 'return', implicit in the analysis of Hansen and Richard (1987). While these two concepts have been used separately in distinct strands of the literature, our analysis unifies and clarifies the relation between them.

First, we define the concept of 'conditional' return. These are pay-offs that have unit price conditionally, i.e. for every possible realization of the conditioning information. Specifically,

$$R_t = \Pi_{t-1}^{-1}\{\, 1 \,\} = \{\, r_t \in X_t \mid \Pi_{t-1}(\, r_t \,) \equiv 1 \,\}. \tag{3}$$

This is the space in which Ferson and Siegel (2001) study the efficient portfolio frontier. We are interested in the *ex-ante* efficiency of managed portfolios, rather than the *ex-post* efficiency once conditioning information is known[1]. Therefore, we will focus on unconditional efficiency: an element $r_t^m \in R_t$ with unconditional mean $m \in \mathbb{R}$ is called 'UC' efficient (*unconditionally* efficient with respect to the space of *conditional* returns), if it has minimum unconditional variance among all elements $r_t \in R_t$ with the same mean m.

In contrast to conditional returns, 'unconditional' returns are those which have unit price only in expectation, i.e. on average across all realizations of the conditioning information. We follow the notation of Hansen and Richard (1987, Section 4) and define the function $\hat{\Pi}(\, x_t \,) = E(\, \Pi_{t-1}(\, x_t \,)\,)$. It should be noted, however, that $\hat{\Pi}(\, x_t \,)$ does not in general

[1] As Dybvig and Ross (1985) show, when portfolio managers possess information not known to outside investors, their (conditionally efficient) strategies may seem conditionally inefficient to outsiders. Hence, the notion of unconditional efficiency is a more appropriate measure of performance in this case.

represent a true price of the managed portfolio x_t, but merely its expected cost. Set

$$\hat{R}_t = \hat{\Pi}^{-1}\{\ 1\ \} = \{\ r_t \in X_t \mid \hat{\Pi}\big(\ r_t\ \big) = 1\ \}. \tag{4}$$

The GHT bounds and their implementation by Bekaert and Liu (2001) are implicitly based on the space of unconditional returns. In analogy to the preceding paragraph, an element $\hat{r}_t^m \in \hat{R}_t$ with unconditional mean $m \in \mathbb{R}$ is called 'UU' efficient (*unconditionally* efficient with respect to the space of *unconditional* returns), if it has minimum unconditional variance among all elements $\hat{r}_t \in \hat{R}_t$ with the same mean m.

REMARK: Obviously, any conditional return is also an unconditional return[2], since the portfolio constraint for unconditional returns is weaker than that for conditional returns. As a consequence, extending the return space to include unconditional returns expands the *ex-ante* mean-variance frontier. However, it will in general not be the case that every UU efficient return is automatically UC efficient or vice versa. This is because a given UU efficient return will typically violate the sharper portfolio constraint for conditional returns, while a given UC efficient return will in general not have minimal variance in the larger space of unconditional returns. In fact, we show later that the portfolio weights of UU efficient and UC efficient returns respond very differently to changes in conditioning information.

Despite these differences, all results we derive in this paper hold equally for conditional as well as unconditional returns, and illustrate the power of the (conditional) Hilbert space approach.

[2]In fact, $\hat{R}_t$ is of co-dimension one in the pay-off space X_t, while R_t typically has infinite co-dimension.

2.3 Orthogonal Representation of the Return Spaces

The purpose of this section is to provide an orthogonal representation of the two return spaces, which leads to a natural parameterization of the unconditional efficient frontier. Although the following results hold for both conditional and unconditional returns, we present here, to save space, only the discussion for conditional returns.

We choose $r_t^* = x_t^*/\Pi_{t-1}\big(\, x_t^* \,\big)$ as the *benchmark* return. Hansen and Richard (1987) show (Lemma 3.1) that r_t^* is the unique return with minimal second moment, and is (conditionally) orthogonal to the space of *excess* (zero cost) returns,

$$Z_t = \Pi_{t-1}^{-1}\big\{\, 0 \,\big\} = \big\{\, z_t \in X_t \mid \Pi_{t-1}\big(\, z_t \,\big) \equiv 0 \,\big\}. \tag{5}$$

Since the difference between any two returns is necessarily an excess return, every return $r_t \in R_t$ can be written in the form $r_t = r_t^* + z_t$ for some $z_t \in Z_t$. To decompose z_t further, we define $z_t^* \in Z_t$ to be the Riesz representation of the (conditional) expectation function on Z_t. As a trivial implication, we have $E\big(\, z_t^{*2} \,\big) = E\big(\, z_t^* \,\big)$. Moreover, it is straight-forward to show that z_t^* is the orthogonal projection of the unit pay-off onto the space of excess returns, $z_t^* = \text{proj}(\, 1 \,|\, Z_t \,)$. In particular, if a risk-free asset is traded, $1 \in X_t$, then z_t^* takes the particularly simple form, $z_t^* = 1 - r_t^*/r_f$.

Hansen and Richard (1987) show that any given return $r_t \in R_t$ with mean $m = E\big(\, r_t \,\big)$ admits a representation of the form,

$$r_t = r_t^* + w \cdot z_t^* + \eta, \qquad \text{with} \quad w = \frac{m - E\big(\, r_t^* \,\big)}{E\big(\, z_t^* \,\big)} \in \mathbb{R}, \tag{6}$$

where $\eta \in Z_t$ is orthogonal to z_t^* with $E\big(\, \eta \,\big) = 0$. Moreover, r_t is unconditionally efficient if and only if $\eta \equiv 0$. Intuitively, r_t^* is the (efficient) benchmark return, z_t^* adds mean (efficiently), while η adds variance and thus measures inefficiency. As a direct consequence of the orthogonality we obtain,

Corollary 2.1 *For given mean $m \in \mathbb{R}$, the efficient variance is given by,*

$$\sigma^2\left(r_t^m\right) = E\left(r_t^{*2}\right) + \frac{\left(m - E\left(r_t^*\right)\right)^2}{E\left(z_t^*\right)} - m^2. \tag{7}$$

In particular, if a risk-free asset is traded, this expression simplifies to,

$$\sigma^2\left(r_t^m\right) = -\frac{E\left(r_t^*\right)}{E\left(r_t^*\right) - r_f} \cdot \left(m - r_f\right)^2 =: \frac{\left(m - r_f\right)^2}{\lambda_*^2} \tag{8}$$

From (8) it is clear that in the presence of a risk-free asset, the efficient standard deviation is a linear function of the mean, so that the frontier has the familiar 'wedge' shape. In particular, the maximum Sharpe ratio in this case is given by λ_* in (8). This result also follows from Equation (16) in Jagannathan (1996). Note that, by construction, r_t^* is in fact located on the lower half of the efficient frontier. Hence, $E\left(r_t^*\right) < r_f$, so that λ_* is indeed well-defined.

3 Discount Factor Bounds

In this section, we develop a generic approach to the construction of discount factor bounds. We extend the Hansen and Richard (1987) framework to obtain a new orthogonal parameterization of the unconditionally efficient portfolio frontier in the absence of a risk-free asset. As a consequence, we obtain a simple, generic methodology for the computation of discount factor bounds. The power of this result can be seen from the fact that explicit implementations of both an the UE bounds of Ferson and Siegel (2002), as well as the GHT bounds, can be obtained as a simple application of Theorem 3.3.

3.1 Review of Discount Factor Bounds

We begin this section with a brief review of the theory of discount factor bounds. In general, we do not assume the market spanned by X_t to be complete. As a consequence, while x_t^* is unique in the pay-off space X_t, it need not be unique in the larger space L_t^2. Therefore, we call any $m_t \in L_t^2$ an admissible *stochastic discount factor (SDF)*, if it prices all traded assets (conditionally) correctly, that is if,

$$\Pi_{t-1}\big(\, x_t \,\big) = E\big(\, x_t \cdot m_t \mid \mathcal{G}_{t-1} \,\big) \qquad \text{for all } x_t \in X_t. \tag{9}$$

We denote by M_t the space of all such admissible SDFs. If a risk-free asset is traded, it is obvious that all discount factors must have the same mean, $E\big(\, m_t \,\big) = 1/r_f$ for all $m_t \in M_t$. By a simple orthogonality argument, one can show that x_t^* is the unique mean-variance efficient discount factor in M_t. In other words, the variance of x_t^* is a lower bound on the variance of all potential discount factors.

However, more interesting is the case without risk-free asset. In this case, the expectation $E\big(\, m_t \,\big)$ is not uniquely determined. Instead, every choice of $\nu = E\big(\, m_t \,\big)$ corresponds to a 'hypothetical' risk-free return, $1/\nu$. It is the key insight of Hansen and Jagannathan (1991) that the arguments outlined above can be used to construct a lower bound on the variance of a discount factor as a function of its mean. To formalize this idea, for given mean $\nu \in \mathbb{R}$, we define M_t^ν as the set of all $m_t \in M_t$ with $E\big(\, m_t \,\big) = \nu$.

Theorem 3.1 *For fixed $\nu = E(m_t)$, the Hansen-Jagannathan bound can be expressed as,*

$$\inf_{m_t \in M_t^\nu} \frac{\sigma(m_t)}{\nu} = \sup_{r_t \in \hat{R}_t} \frac{E(r_t) - 1/\nu}{\sigma(r_t)} =: \hat{\lambda}_*(\nu) \tag{10}$$

In other words, the HJ bounds can be found by calculating the maximum 'hypothetical' Sharpe ratio, corresponding to the shadow risk-free return $1/\nu$. We use this fact to construct discount factor bounds in the following sections. The proof of this theorem is an extension of the original derivation in Hansen and Jagannathan (1991). A similar argument can also be found in Cochrane (2001).

Remark: Note that the supremum on the right-hand side of (10) is taken over the space of all *unconditional* returns, $\hat{R}_t$. Since this space has co-dimension one in the augmented pay-off space X_t, the resulting bounds are the sharpest possible for given choice of conditioning instruments. Gallant, Hansen, and Tauchen (1990) (GHT) construct these bounds by projecting the discount factor unconditionally onto the pay-off space. An alternative approach to the construction of the GHT bounds, implicitly based on our Theorem 3.1, is developed in Bekaert and Liu (2001). To construct their 'optimally scaled bounds (OSB)', they determine an optimal scaling vector as a function of conditioning information, such that the subspace obtained by scaling the base assets by scalar multiples of this vector attains the maximum hypothetical Sharpe ratio in (10).

Ferson and Siegel (2002) use a similar approach in the construction of their 'unconditionally efficient (UE)' bounds in that they compute the maximum hypothetical Sharpe ratio implied by a shadow risk-free rate $1/\nu$. However, the maximum is taken over the smaller space R_t of *conditional* returns. Hence, for the UE bounds the right-hand side of (10) is replaced by,

$$\lambda_*(\nu) := \sup_{r_t \in R_t} \frac{E(r_t) - 1/\nu}{\sigma(r_t)} < \hat{\lambda}_*(\nu) \tag{11}$$

As a consequence, Ferson and Siegel's UE bounds will in general plot below the GHT bounds, and (10) in this case becomes an inequality. In other words, while the UE bounds are sharper than the HJ bounds without conditioning information, they will reject fewer asset pricing models than the (sharpest) GHT bounds.

3.2 Portfolio-Based Derivation of Discount Factor Bounds

Since the HJ bound is trivial in the case when a risk-free asset is traded, we consider only the case without risk-free asset. To construct the bound for a given mean ν of the discount factor, we need to find the portfolio that maximizes the hypothetical Sharpe ratio in (10). We provide a generic construction of this portfolio, and an expression for the implied hypothetical Sharpe ratio, and hence the discount factor bound. Ferson and Siegel's UE bounds, as well as Bekaert and Liu's implementation of the GHT bounds, can now be obtained by applying our result to the space of conditional and unconditional returns, respectively.

It is clear that the portfolio which attains the maximum hypothetical Sharpe ratio must be unconditionally mean-variance efficient (in the space of risky asset returns). By (6), the entire unconditionally efficient frontier can be parameterized by a single scalar weight, so that the problem of constructing the bounds reduces to a univariate maximization problem. However, the representation of the frontier in terms of r_t^* and z_t^* is motivated by orthogonality arguments with respect to the second moment inner product. While this inner product is the natural choice in the space L_t^2, the focus on mean-variance efficiency suggests instead the use of *covariance* as an inner product. While in general, the covariance function is only positive semi-definite, it will be positive definite when restricted to a space that does not contain a constant element. Thus, in the case without risk-free asset, the pay-off space X_t is a Hilbert space with respect to the covariance inner product. We use this fact to construct a new covariance-orthogonal representation of the efficient frontier that greatly simplifies the computation of discount factor bounds.

Recall that r_t^* is defined as the unique efficient return that is orthogonal to the space Z_t of excess returns, with respect to the second moment inner product. Extending this argument, we consider instead the unique efficient return r_t^0 that is orthogonal with respect to the covariance inner product, i.e. $\mathrm{cov}\big(\, r_t^0\, ,\, z_t\, \big) = 0$ for all $z_t \in Z_t$. To obtain the representation of r_t^0 in terms of the Hansen and Richard (1987) parameterization, we write $r_t^0 = r_t^* + w(\, r_t^0\,) \cdot z_t^*$. Orthogonality then translates into,

$$
\begin{aligned}
0 &= \mathrm{cov}\big(\, r_t^* + w(\, r_t^0\,) \cdot z_t^*\, ,\, z_t\, \big) \\
&= E\big(\, z_t\, \big)\, \big[\, w(\, r_t^0\,)\big(\, 1 - E\big(\, z_t^*\, \big)\, \big) - E\big(\, r_t^*\, \big)\, \big]
\end{aligned}
\tag{12}
$$

Here, the last equality follows from the fact that z_t^* is the Riesz representation of the unconditional expectation on Z_t, so that $E\big(\, z_t^* z_t\, \big) = E\big(\, z_t\, \big)$. Interestingly, the orthogonal efficient return r_t^0 is in fact nothing other than the global minimum variance return. To see this, note that the orthogonality condition (12), evaluated for $z_t = z_t^*$, is equivalent to the first-order condition of the unconstrained variance minimization problem, since

$$
\frac{\partial}{\partial w}\, \sigma^2\big(\, r_t^* + w z_t^*\, \big) = 2 \cdot \mathrm{cov}\big(\, r_t^* + w z_t^*\, ,\, z_t^*\, \big).
$$

We summarize these findings in the following,

Lemma 3.2 *In the case without risk-free asset, the global minimum variance return* r_t^0

admits a canonical representation of the form[3],

$$r_t^0 = r_t^* + w(\, r_t^0\,) \cdot z_t^*, \qquad \textit{with} \quad w(\, r_t^0\,) = \frac{E\big(\, r_t^*\,\big)}{1 - E\big(\, z_t^*\,\big)}. \tag{13}$$

Moreover, r_t^0 *is the unique efficient return that is uncorrelated with all* $z_t \in Z_t$*.*

To span the efficient frontier, we also need to choose an appropriate excess return $z_t^0 \in Z_t$. Recall that z_t^* is chosen as the Riesz-representation of the unconditional expectation on the space of excess returns. Extending this line of argument to the covariance inner product, we seek to find $z_t^0 = w(\, z_t^0\,) \cdot z_t^*$ such that, for all $z_t \in Z_t$,

$$\begin{aligned} E\big(\, z_t\,\big) \;&=\; \operatorname{cov}\big(\, w(\, z_t^0\,) \cdot z_t^*\,,\, z_t\,\big) \\ &=\; w(\, z_t^0\,) E\big(\, z_t\,\big) \left[\, 1 - E\big(\, z_t^*\,\big)\,\right]. \end{aligned} \tag{14}$$

As before, the last equality follows from the fact the z_t^* is the Riesz representation of the expectation functional. As a direct consequence of (14), we obtain, $\sigma^2\big(\, z_t^0\,\big) = E\big(\, z_t^0\,\big)$. Note that, by construction, r_t^0 and z_t^0 span the mean-variance efficient frontier, since for $\kappa \in \mathbb{R}$,

$$r_t^0 + \kappa \cdot z_t^0 = r_t^* + \left[\, w(\, r_t^0\,) + \kappa \cdot w(\, z_t^0\,)\,\right] z_t^*$$

Hence, we can use the representation of the efficient frontier in terms of r_t^0 and z_t^0 to calculate the maximum hypothetical Sharpe ratio and thus the discount factor bound.

[3]While the derivation of the global minimum variance portfolio in the framework of Hansen and Richard (1987) has been reported previously (e.g. Cochrane (2001)), the novel feature here is that this portfolio arises naturally in an orthogonal representation of the efficient frontier.

Theorem 3.3 *For given $\nu = E\big(m_t\big)$, the maximum hypothetical Sharpe ratio $\lambda_*(\nu)$ that attains the discount factor bound in (10), admits a decomposition of the form,*

$$\lambda_*^2(\nu) = \lambda_0^2(\nu) + E\big(z_t^0\big), \qquad \text{with} \quad \lambda_0(\nu) = \frac{E\big(r_t^0\big) - 1/\nu}{\sigma\big(r_t^0\big)}. \tag{15}$$

Moreover, the maximum hypothetical Sharpe ratio is attained by the return

$$r_t^0 + \kappa_\nu^* \cdot z_t^0, \qquad \text{with} \quad \kappa_\nu^* = \frac{\sigma^2\big(r_t^0\big)}{E\big(r_t^0\big) - 1/\nu}. \tag{16}$$

Note that, by construction, the portfolio constructed in (16) is the unique efficient portfolio with *zero-beta* rate $1/\nu$.

Proof of Theorem 3.3: For arbitrary $\kappa \in \mathbb{R}$, consider the efficient return $r_t = r_t^0 + \kappa \cdot z_t^0$. The objective is to find κ such as to maximize the implied hypothetical (squared) Sharpe ratio,

$$\frac{\big[E\big(r_t^0 + \kappa \cdot z_t^0\big) - 1/\nu\big]^2}{\sigma^2\big(r_t^0 + \kappa \cdot z_t^0\big)} = \frac{\big[E\big(r_t^0\big) + \kappa \cdot E\big(z_t^0\big) - 1/\nu\big]^2}{\sigma^2\big(r_t^0\big) + \kappa^2 \cdot E\big(z_t^0\big)}$$

The first-order condition for this maximization problem can be written as,

$$\kappa \cdot E\big(z_t^0\big)\big[E\big(r_t^0\big) + \kappa \cdot E\big(z_t^0\big) - 1/\nu\big] = E\big(z_t^0\big)\big[\sigma^2\big(r_t^0\big) + \kappa^2 \cdot E\big(z_t^0\big)\big]$$

The quadratic terms in this expression cancel, due to our choice of z_t^0. Hence, the first-order condition can be easily solved to obtain (16). To prove the decomposition (15) of the maximum hypothetical Sharpe ratio, we re-write the first-order condition as,

$$\lambda_*^2(\nu) = \frac{\big[E\big(r_t^0\big) - 1/\nu\big]^2}{\sigma^2\big(r_t^0\big)} + E\big(z_t^0\big) = \lambda_0^2(\nu) + E\big(z_t^0\big)$$

This completes the proof of Theorem 3.3. □

This result, to our knowledge, is new. It provides not only a very simple way of constructing discount factor bounds, but also a portfolio-based interpretation of these bounds. Also, we would like to emphasize that the above approach to the construction of discount factor bounds is valid even in the fixed-weight case, when there is no conditioning information. If a risk-free asset is traded, Jagannathan (1996) shows that the maximum Sharpe ratio is given by $E(\, z_t^* \,)/(1 - E(\, z_t^* \,))$. On the other hand, our decomposition (15) can be re-written as,

$$\lambda_*^2(\, \nu \,) = \lambda_0^2(\, \nu \,) + \frac{E(\, z_t^* \,)}{1 - E(\, z_t^* \,)}. \tag{17}$$

Our result thus generalizes Equation (16) of Jagannathan (1996), which gives the HJ bound in the case with risk-free asset. We extend this result to the case without risk-free asset, providing an new decomposition of the HJ bounds. Using Theorem 3.1, we now obtain the explicit expression for the discount factor bounds,

Corollary 3.4 *For fixed $\nu = E(\, m_t \,)$, the Hansen-Jagannathan bound can be written as,*

$$\inf_{m_t \in M_t^\nu} \sigma^2(\, m_t \,) \geq \frac{(\, \gamma_1^2 + \gamma_2\gamma_3 \,) \cdot \nu^2 - 2\gamma_1 \cdot \nu + 1}{\gamma_2}, \tag{18}$$

where γ_1, γ_2 are the unconditional mean and variance of r_t^0, respectively, and $\gamma_3 = E(\, z_t^0 \,)$.

In other words, the lower bound on the variance of an SDF is simply a quadratic function of its mean. Our characterization of this function in terms of the moments of the global minimum variance (GMV) portfolio may be useful in implementing discount factor bounds.

Remark: While the results developed in this section are formulated for conditional returns, all arguments are equally valid for unconditional returns, and even in the fixed-weight case without conditioning information. In particular, for unconditional returns, one constructs corresponding orthogonal elements $\hat{r}_t^0$ and $\hat{z}_t^0$, to obtain the respective version of Theorem

3.3. However, it should be noted that in the case of unconditional returns, inequality (18) becomes an equality, since the minimum on the left-hand side is attained by the GHT projection, which can be normalized to become an unconditional return. However, since this is not possible in the case of conditional returns, the implied bounds in this case are truly portfolio-based bounds.

4 Case Without Risk-Free Asset

We now operationalize the results of the preceding section by constructing the weights of r^* and z^*, for both conditional and unconditional returns. We thus recover the weights stated in Ferson and Siegel (2001), and provide a characterization of the efficient frontier implicit in Bekaert and Liu (2001). Using our results from Section 3, we obtain explicit expressions for the UE bounds of Ferson and Siegel (2002) in terms of their 'efficient set' constants, and Bekaert and Liu's (2001) implementation of the GHT bounds. Our unified framework may help clarify the relationship between these two sets of bounds.

We show that for conditional returns, the efficient weights respond 'conservatively' to extreme outcomes of the conditioning instruments, similar to the behavior reported in Ferson and Siegel (2001) in the case with risk-free asset. In contrast, the efficient weights for unconditional returns display a much more 'aggressive' response. This may have important implications for the robustness of the bounds derived from these portfolios.

4.1 Construction of Managed Portfolios

Suppose there are n risky assets, indexed $k = 1 \ldots n$. We denote the (gross) return of the k-th asset by $r_t^k \in L_t^2$. Let $\tilde{R}_t := \big(r_t^1 \, \ldots \, r_t^n \big)'$ denote the vector of risky asset returns. We define the augmented pay-off space X_t as the space of all $x_t = \tilde{R}_t' \theta_{t-1}$, where θ_{t-1} are $\mathcal{G}_{t-1}$-measurable, $\mathbb{R}^n$-valued functions such that $x_t \in L_t^2$. We interpret the elements of this space as 'managed' pay-offs, since the weights θ_{t-1} are functions of the conditioning information. On this space, we define the conditional pricing function $\Pi_{t-1}\big(x_t \big) = e'\theta_{t-1}$, where e is an n-vector of ones. Following Section 2, the space R_t of *conditional* returns in this framework is given by the set of all $r_t = \tilde{R}_t' \theta_{t-1}$ with $e'\theta_{t-1} \equiv 1$. Conversely, the space $\hat{R}_t$ of *unconditional* returns is defined by the (weaker) constraint $E\big(e'\theta_{t-1} \big) = 1$. Denoting the conditional moments of the return vector as,

$$\mu_{t-1} = E\big(\tilde{R}_t \mid \mathcal{G}_{t-1} \big), \qquad \text{and} \qquad \Lambda_{t-1} = E\big(\tilde{R}_t \cdot \tilde{R}_t' \mid \mathcal{G}_{t-1} \big), \tag{19}$$

returns can be written in the form $\tilde{R}_t = \mu_{t-1} + \varepsilon_t$, where μ_{t-1} is the conditional expectation of returns given conditioning information, and ε_t is the residual disturbance with variance-covariance matrix $\Lambda_{t-1} - \mu_{t-1}\mu_{t-1}'$. This is the formulation of the model with conditioning information used in Ferson and Siegel (2001).[4]

[4] Note however that our notation differs slightly from that used in Ferson and Siegel (2001), who define Λ_{t-1} to be the *inverse* of the conditional second-moment matrix.

4.2 Constructing the Efficient Frontier

In this section, we construct the efficient frontier in the absence of a risk-free asset. To simplify notation, we will drop the time index for the remainder of this section. We define (conditional) 'efficient set' constants,

$$A = e'\Lambda^{-1}e, \qquad B = \mu'\Lambda^{-1}e, \qquad D = \mu'\Lambda^{-1}\mu \tag{20}$$

Finally, we denote by the lowercase letters a, b, and d the unconditional expectations of the quantities A, B, and D, respectively.[5]

Theorem 4.1 (Constructing r^*) *In the case without risk-free asset,*

(i) for ***conditional*** *returns,* $r^* = \tilde{R}'\theta$, *with* $\theta = \frac{1}{A}\Lambda^{-1}e$

In particular, $E(r^*) = E(\frac{B}{A})$ *and* $\sigma^2(r^*) = E(\frac{1}{A}) - E(\frac{B}{A})^2$

(ii) for ***unconditional*** *returns,* $\hat{r}^* = \tilde{R}'\theta$, *with* $\theta = \frac{1}{a}\Lambda^{-1}e$

In particular, $E(\hat{r}^*) = \frac{b}{a}$ *and* $\sigma^2(\hat{r}^*) = \frac{1}{a} - (\frac{b}{a})^2$

PROOF: Appendix A.1. □

Note in particular that the expressions for the weights in both cases are identical up to the normalization constants, A and a, respectively. These constants ensure that the portfolio

[5] The constants a, b, and d are identical to those defined in Bekaert and Liu (2001) if the base asset prices are normalized to one in their analysis.

constraints are satisfied. Their difference reflects the fact that for conditional returns, the constraint is required to hold for every realization of the conditioning information, while for unconditional returns it is required to hold only in expectation.

Theorem 4.2 (Constructing z^*) *In the case without risk-free asset,*

*(i) for **conditional** returns,* $z^* = \tilde{R}'\theta$, *with* $\theta = \Lambda^{-1}\big(\mu - \frac{B}{A} e \big)$

In particular, $E\big(z^* \big) = E\big(D - \frac{B^2}{A} \big)$

*(ii) for **unconditional** returns,* $\hat{r}^* = \tilde{R}'\theta$, *with* $\Lambda^{-1}\big(\mu - \frac{b}{a} e \big)$

In particular, $E\big(\hat{z}^* \big) = d - \frac{b^2}{a}$

Proof: Appendix A.2. □

4.3 Conditional Returns and Discount Factor Bounds

We now study the properties of the efficient frontier for conditional returns, using our results from the preceding section. Using our results from Section 3.2, we then characterize the implied discount factor bounds for conditional returns. From this we obtain an explicit expression for the 'unconditionally efficient (UE)' bounds considered in Ferson and Siegel (2002), in terms of their efficient set constants.

(a) Efficient Portfolio Weights

Ferson and Siegel (2001) describe the efficient frontier in terms of 'efficient set' constants. Following their notation, we can identify the moments in Theorem 4.1 (i) as $E\big(r^* \big) = \alpha_2$

and $\sigma^2\big(r^* \big) = \alpha_1 - \alpha_2^2$. In other words, α_2 and α_1 are the first and second moments of r^*, respectively. Similarly, we find $\alpha_3 = E\big(z^* \big)$.

Lemma 4.3 *The weights of the UC efficient return r^m for given mean $m \in \mathbb{R}$ are,*

$$\theta_m^* = \Lambda^{-1}\big(\frac{1 - w_m^* B}{A} e + w_m^* \mu \big), \qquad \textit{where} \quad w_m^* = \frac{m - \alpha_2}{\alpha_3} \tag{21}$$

The above expression has an interesting economic interpretation, which enables us to analyze the behavior of the weights in response to changes in conditioning information. Similar to the fixed-weight case, the efficient portfolio weights in (21) consist of two components, an equally weighted 'market' portfolio whose weights are proportional to the unit vector e, and a 'managed' component whose weights are proportional to the vector of conditional expected base asset returns, μ. The first component responds uniformly across all assets to changes in the conditioning instruments, while the weights of the second respond to information about the return of individual assets. Motivated by this observation, we interpret the former as a 'market-timing' component, while the latter captures effects of asset 'selectivity'. The market-timing component clearly reflects the more conservative part of the efficient portfolio, while the selectivity component reflects a more aggressive response to conditioning information.

The overall behavior of the efficient portfolio is determined by the allocation of weights across these two components. From the definition of the conditional constant A, it follows that the sum of weights on the market-timing component is simply $1 - w_m^* B$. It is easy to see that for extreme values of the conditioning instruments, B converges to zero. In other words, the efficient portfolio responds to extreme information by shifting all weight into the conservative market-timing component.[6] In Section 4.4, we will see that the efficient weights for unconditional returns, from which Bekaert and Liu's (2001) implementation of the GHT bounds is constructed, display a more aggressive behavior, shifting all weight into the selectivity component. This may have important implications for the robustness of bounds constructed from these weights. In particular, the appearance of the vector μ in the selectivity component will amplify any measurement errors of the conditional moments for extreme values of the conditioning information.

(b) Discount Factor Bounds

Following Section 3, discount factor bounds can be obtained by maximizing the hypothetical Sharpe ratio implied by a given $\nu = E\big(m\big)$. Applying Corollary 3.4 to the space of conditional returns, we can now give an explicit expression for Ferson and Siegel's (2002) UE bounds in terms of their 'efficient set' constants,

Corollary 4.4 (UE Bound) *Necessary for a candidate m to be an admissible SDF is,*

$$\sigma^2(\,m\,) \geq \frac{(\alpha_1\alpha_3 + \alpha_2^2)\cdot\nu^2 - 2\alpha_2\cdot\nu + (1-\alpha_3)}{\alpha_1(1-\alpha_3) - \alpha_2^2}, \quad \textit{where} \quad \nu = E\big(\,m\,\big) \tag{22}$$

In other words, the UE bound takes the form of a second-order polynomial in ν, where the coefficients are related to the efficient set constants. This formulation of the bounds should be useful in practical implementations.

Proof: In terms of the 'efficient set' constants defined in Ferson and Siegel (2001), the mean and variance of the global minimum variance return can be written as, $\gamma_1 = \alpha_2/(1-\alpha_3)$ and

[6] Ferson and Siegel (2001) report a similar result only in the case with risk-free asset. In Section 5.3, we provide a theoretical explanation for these findings.

$\gamma_2 = \alpha_1 - \alpha_2^2/(1-\alpha_3)$, respectively. Similarly, we obtain $\gamma_3 = \alpha_3/(1-\alpha_3)$. Substituting this into Equation (18) and re-arranging, we obtain the desired result. □

Note that, since the space of conditional returns is smaller than the space of unconditional returns, the UE bounds will not in general attain the highest possible discount factor bound, and (22) will in general be a strict inequality. However, Ferson and Siegel (2002) show that the UE bounds possess attractive empirical properties. First, they remain valid lower bounds for the variance of stochastic discount factors even when the conditional moments are misspecified. Second, the conservative behavior of the weights of the portfolio that attains the bounds may translate into robust sampling properties.

Remark: By Theorem 3.3, the portfolio that attains the discount factor bounds can be written as, $r^0 + \kappa_\nu^* \cdot z^0$. Using Theorems 4.1 (i) and 4.2 (i), the weights of this portfolio are,

$$\theta_\nu^* = \Lambda^{-1}\Big(\frac{1-\kappa_\nu^* B}{A} e + \kappa_\nu^* \mu\Big), \qquad \text{with} \quad \kappa_\nu^* = \frac{\alpha_1 \nu - \alpha_2}{\alpha_2 \nu - (1-\alpha_3)} \tag{23}$$

From the preceding section, we can identify this as the weights of the efficient return for mean $\alpha_2 + \alpha_3 \cdot \kappa_\nu^*$. The zero-beta rate associated with this portfolio is, by construction, $1/\nu$.

4.4 Unconditional Returns and Discount Factor Bounds

Gallant, Hansen, and Tauchen (1990) consider discount factor bounds in the space of *unconditional* returns. While their bounds are constructed by projecting the discount factor onto the pay-off space, Bekaert and Liu (2001) use an approach based on maximizing the implied hypothetical Sharpe ratio in the spirit of Theorem 3.1. Using our framework, we highlight the implicit portfolio interpretation of their approach.

(a) Efficient Portfolio Weights

Using Theorems 4.1 (ii) and 4.2 (ii), we obtain,

Lemma 4.5 *The weights of the UU efficient return $\hat{r}^m$ for given mean $m \in \mathbb{R}$ are,*

$$\theta_m^* = \Lambda^{-1}\Big(\frac{1 - w_m^* b}{a} e + w_m^* \mu\Big), \qquad \textit{where} \quad w_m^* = \frac{am - b}{ad - b^2} \tag{24}$$

Note that, as in case for conditional returns, the efficient weight can be decomposed into a conservative 'market-timing' part, and a 'selectivity' component. An argument similar to that for conditional returns show that, for extreme values of the conditioning instruments, the weight shifts away from the conservative component into the 'selectivity' part. Thus, the efficient weights in the case of unconditional returns display a more aggressive response. By Corollary 2.1, the efficient variance for given mean m is,

$$\sigma^2(\hat{r}^m) = \frac{d - 2bm + (a - \delta)m^2}{\delta}, \qquad \text{with} \quad \delta = ad - b^2 \tag{25}$$

This characterizes the efficient frontier implicit in the analysis of Bekaert and Liu (2001).

(b) Discount Factor Bounds

In this section, we provide an alternative derivation of the GHT bounds, applying Theorem 3.3 to the case of unconditional returns. While an expression for these bounds was also derived by Bekaert and Liu (2001), we add to their analysis by providing a portfolio-based interpretation.

Corollary 4.6 (GHT Bound) *Necessary for a candidate m to be an admissible SDF is,*

$$\sigma^2(m) \geq \frac{d \cdot \nu^2 - 2b \cdot \nu + (a - \delta)}{1 - d} \quad \textit{where} \quad \nu = E(m) \tag{26}$$

This is Equation (25) in Bekaert and Liu (2001).

PROOF: First, using the 'efficient set' constants of Bekaert and Liu (2001), the mean and variance of the global minimum variance return $\hat{r}^0$ can be written as, $\gamma_1 = b/(a-\delta)$ and $\gamma_2 = (1-d)/(a-\delta)$. Similarly, we obtain $\gamma_3 = \delta/(a-\delta)$. Applying Corollary 3.4 then gives the desired result. □

Note that (26) in this case is sharp, since the right-hand side is attained by the variance of the GHT projection. Similar to the UE bounds, this expression also takes the form of a second-order polynomial in ν. The methodology in Bekaert and Liu (2001) yields the OSB as a quartic over a quadratic polynomial which, if moments are correctly specified, reduces to the above expression.

REMARK: In our analysis, the portfolio that attains the discount factor bound is given from Theorem 3.3 as $\hat{r}^0 + \kappa_\nu^* \cdot \hat{z}^0$. Rewriting this using Theorems 4.1 (ii) and 4.2 (ii), we obtain the weights of this portfolio as,

$$\theta_\nu^* = \Lambda^{-1}\Big(\frac{1-d}{a-\delta-b\nu} \cdot e + \frac{b-\nu}{a-\delta-b\nu} \cdot \mu \Big) \tag{27}$$

It is straight-forward to show that these weights in fact coincide with the optimal scaling vector given in Equations (22) and (23) of Bekaert and Liu (2001), suitably normalized.

5 Case With Risk-Free Asset

In this section, we construct the efficient frontier for conditional and unconditional returns when there is a risk-free asset. This enables us to provide an explanation for the 'conservative response' of the efficient weights for conditional returns, as reported in Ferson and Siegel

(2001). In contrast, we show that the weights for unconditional returns exhibit a much more aggressive behavior in that they may require extreme long and short positions in individual assets. We also obtain explicit expressions for the maximum Sharpe ratios in both cases, which we use in a separate paper to construct a portfolio-based test of whether asset return predictability significantly expands the mean-variance frontier.

5.1 Constructing Managed Portfolios

As in Section 4, we consider the case where the pay-off space is 'spanned' by a finite set of base assets. As before, we denote by $\tilde{R}_t := \left(r_t^1 \ \ldots \ r_t^n \right)'$ the vector of risky asset returns. However, we now consider the case in which a risk-free asset is traded. We denote its (gross) return by r_f. In this case, the augmented space of 'managed' pay-offs now consist of elements of the form $x_t = \theta_{t-1}^0 r_f + (\tilde{R}_t - r_f e)' \theta_{t-1}$. Note that we allow portfolios that have have 'managed' positions in the risk-free asset, which themselves are hence no longer risk-free, since the weight θ_{t-1}^0 may vary with conditioning information.

Note that, in contrast to Section 4, the weights θ_{t-1} on the risky assets are now applied to their *excess* returns. As a consequence, the space R_t of *conditional* returns in this framework is given by those pay-offs for which $\theta_{t-1}^0 \equiv 1$. However, the space $\hat{R}_t$ of *unconditional* returns is defined by the (less strict) constraint $E\left(\theta_{t-1}^0 \right) = 1$. We define

$$\Sigma_{t-1} = \mathrm{Var}\left(\tilde{R}_t \mid \mathcal{G}_{t-1} \right) = \Lambda_{t-1} - \mu_{t-1} \cdot \mu_{t-1}' \tag{28}$$

Note that, in contrast to Ferson and Siegel (2001), we derive the efficient portfolio weights in the case with risk-free asset in terms of the conditional variance-covariance matrix Σ_{t-1} of returns, rather than the matrix of second moments Λ_{t-1}. This will enable us to derive an expression for the Sharpe ratio for unconditional returns, which is similar to Equation (16) in Jagannathan (1996).

5.2 Constructing the Efficient Frontier

In this section, we derive the weights of r^* and z^* as functions of conditioning information, for conditional and unconditional returns. This enables us to construct the efficient portfolio frontier in both cases. To simplify notation, we will drop the time index for the remainder of this section. We define,

$$H^2 = (\mu - r_f e)' \Sigma^{-1} (\mu - r_f e), \tag{29}$$

In analogy to Section 4, we denote by $h^2 = E\big(H^2\big)$ the unconditional expectation of H^2. Note that, in the case of one risky and one risk-free asset, the quantity H^2 is in fact the maximum squared conditional Sharpe ratio.

Theorem 5.1 (Constructing r^*) *In the case without risk-free asset,*

(i) for ***conditional*** *returns,* $r^* = r_f + \big(\tilde{R} - r_f e\big)'\theta$,

with $\theta = -\big(\dfrac{r_f}{1+H^2}\big) \cdot \Sigma^{-1}\big(\mu - r_f e\big)$

(ii) for ***unconditional*** *returns,* $\hat{r}^* = \theta^0 r_f + \big(\tilde{R} - r_f e\big)'\theta$,

with $\theta = -\big(\dfrac{r_f}{1+h^2}\big) \cdot \Sigma^{-1}\big(\mu - r_f e\big)$ *and* $\theta^0 = \big(\dfrac{1+H^2}{1+h^2}\big)$

Proof: Appendix A.3. □

Note that, for unconditional returns the risk-free asset is dynamically managed, in that its weight θ^0 is a function of conditioning information. The weights of z^* can easily be derived from the above expressions, since we know that in the case with risk-free asset, $z^* = 1 - r^*/r_f$.

5.3 Efficient Frontier of Conditional Returns

We now calculate the efficient portfolio weights for conditional returns, thus providing a constructive proof of the results stated in Theorem 2 of Ferson and Siegel (2001). In the presence of a risk-free asset, Ferson and Siegel (2001) formulate their efficient portfolio weights in terms of a quantity ζ (see Equation (11) in their paper), which may be interpreted as a measure of the effect of return predictability on the mean-variance frontier. Using a simple matrix identity, this quantity can be shown to equal,

$$\zeta = E\Big(\frac{H^2}{1+H^2} \Big) \tag{30}$$

As a consequence, we can rewrite the unconditional moments of r^* as $E\big(r^* \big) = r_f (1 - \zeta)$ and $\sigma^2\big(r^* \big) = r_f \zeta (1 - \zeta)$. The explicit weights for r^* and z^*, give us the weights of the efficient return for a given mean m as,

$$\theta^*_m = \frac{w^*_m - r_f}{1+H^2} \cdot \Sigma^{-1}\big(\mu - r_f e \big), \qquad \text{where} \quad w^*_m = \frac{m - r_f(1 - \zeta)}{\zeta} \tag{31}$$

Using the same matrix identity, this expression can be shown to be identical to that stated in Equation (12) of Ferson and Siegel (2001). Our expression (31), while similar to the efficient portfolio weights in the absence of conditioning information, differs from the latter in that the normalization factor $1+H^2$ is in fact time-varying. The presence of this time-varying quantity is an artefact of the *conditional* portfolio constraint. For small values of the conditioning information, the response of the weights is determined mainly by the numerator $\Sigma^{-1}(\mu -$

$r_f e$), as the normalization coefficient is close to one. In contrast, for larger values of the conditioning information, the normalization coefficient increasingly dominates the response. This is illustrated in Figure 1 in Ferson and Siegel (2001), which graphs the portfolio weights as a function of a conditioning variable in the case where the conditional mean μ is linear. In contrast, in the case of unconditional returns, the normalization coefficient is unconditionally constant and therefore does not affect the responsiveness of the weights (see Section 5.4 (a)).

REMARK: Note that in the case with risk-free asset, the maximum Sharpe ratio is in fact attained by r^*. From the results of the preceding sections, we find that the maximum (squared) Sharpe ratio for conditional returns is given by,

$$\lambda_*^2 = E\Big(\frac{H^2}{1+H^2} \Big) / E\Big(\frac{1}{1+H^2} \Big), \tag{32}$$

which, in the notation of Ferson and Siegel (2001), can be written as $\lambda_*^2 = \zeta/(1-\zeta)$.

5.4 Efficient Frontier of Unconditional Returns

In this section, we construct the efficient frontier for unconditional returns. The weights of $\hat{r}^*$ and $\hat{z}^*$ give us the risky asset weights for the efficient return for given mean m as,

$$\theta_m^* = \frac{m - r_f}{h^2} \cdot \Sigma^{-1}\big(\mu - r_f e \big) \tag{33}$$

In the case without conditioning information, μ and Σ are constant, so that the above expression reduces to the familiar efficient weights in classical mean-variance theory. Note that, in contrast to the Ferson-Siegel weights in (31), the normalization factor in this case is in fact constant. As a consequence, these weights will not in general exhibit the 'conservative response' of the Ferson-Siegel weights. In fact, in the linear regression specification considered in Ferson and Siegel (2001), the corresponding efficient weights for unconditional returns will be linear! This implies that extreme portfolio positions are possible,[7] which is due to the less restrictive portfolio constraints. Moreover, measurement errors in the conditional moments will be amplified linearly for extreme values of the conditional instruments, which may have implications for the robustness of these weights.

Remark: As argued above, we can compute the maximum Sharpe ratio in this case directly from the moments of $\hat{r}^*$. Due to the unconditional normalization constants, the maximum (squared) Sharpe ratio for unconditional returns takes the simple form,

$$\lambda_*^2 = \left(\frac{h^2}{1+h^2}\right) / \left(\frac{1}{1+h^2}\right) = h^2 \tag{34}$$

Note that this expression is very similar to that for the Sharpe ratio in the fixed-weight case without conditioning information. In fact, h^2 is simply the unconditional expectation of the maximum squared conditional Sharpe ratio. While our result applies to multiple risky assets, the case of a single risky asset has been analyzed in Cochrane (1999).

6 Conclusion

We provide a unified framework for the study of mean-variance efficiency and discount factor bounds in the presence of conditioning information. First, we develop a new portfolio-based framework for the implementation of discount factor bounds with and without conditioning information. To do this, we construct a new, covariance-orthogonal parameterization of the

[7] The behavior of the weights is thus similar to the sensitivity of traditional fixed weight portfolios, see also Green and Hollifield (1992) for a discussion of this issue.

space of returns on managed portfolios. As a direct implication of our results, we obtain a general, portfolio-based methodology for the implementation of discount factor bounds. Our results connect various different approaches to the construction of such bounds, and allow a direct comparison of their respective properties. Second, we explicitly construct the weights of efficiently managed portfolios as functions of the conditioning information. This enables us to characterize the optimal portfolios that maximize unconditional Sharpe ratios and thus attain the sharpest possible discount factor bounds. Moreover, our formulation of the weights facilitates the analysis of their behavior in response to changes in conditioning information.

Our analysis has several important empirical applications. First, the expression for the maximum Sharpe ratio in the presence of conditioning information can be used to study the effect of return predictability. Second, the techniques developed in this paper can be used to construct portfolio-based tests of conditional asset pricing models. Finally, our results can also be used to construct measures of portfolio performance in the presence of conditioning information, a topic we are currently investigating.

References

Bekaert, G., and J. Liu (2001): "Conditioning Information and Variance Bounds on Pricing Kernels," working paper, Columbia Business School.

Campbell, J. (1987): "Stock Returns and the Term Structure," *Journal of Financial Economics*, 18, 373–399.

Chamberlain, G., and M. Rothschild (1983): "Arbitrage, Factor Structure, and Mean-Variance Analysis on Large Asset Markets," *Econometrica*, 51(5), 1281–1304.

Cochrane, J. (1999): "Portfolio Advice for a Multi-Factor World," *Economic Perspectives, Federal Reserve Bank of Chicago*, 23(3), 59–78.

Cochrane, J. (2001): *Asset Pricing*. Princeton University Press, Princeton, New Jersey.

Dybvig, P., and S. Ross (1985): "Differential Information and Performance Measurement using a Security Market Line," *Journal of Finance*, 60, 383–399.

Fama, E., and K. French (1988): "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3–25.

Ferson, W. (2002): "Test of Multifactor Pricing Models, Volatility Bounds and Portfolio Performance," in *Handbook of the Economics of Finance*, ed. by G. Constantinides et al. forthcoming, Elsevier Science.

Ferson, W., and A. Siegel (2001): "The Efficient Use of Conditioning Information in Portfolios," *Journal of Finance*, 56(3), 967–982.

Ferson, W., and A. Siegel (2002): "Stochastic Discount Factor Bounds with Conditioning Information," forthcoming, Review of Financial Studies.

Gallant, R., L. Hansen, and G. Tauchen (1990): "Using Conditional Moments of Asset Payoffs to Infer the Volatility of Intertemporal Marginal Rates of Substitution," *Journal of Econometrics*, 45(1), 141–179.

Green, R., and B. Hollifield (1992): "When will Mean-Variance Efficient Portfolios be Well Diversified?," *Journal of Finance*, 47(5), 1785–1809.

Hansen, L., and R. Jagannathan (1991): "Implications of Security Markets Data for Models of Dynamic Economies," *Journal of Political Economy*, 99(2), 225–262.

Hansen, L., and S. Richard (1987): "The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models," *Econometrica*, 55(3), 587–613.

Jagannathan, R. (1996): "Relation between the Slopes of the Conditional and Unconditional Mean-Standard Deviation Frontier of Asset Returns," in *Modern Portfolio Theory and its Applications: Inquiries into Asset Valuation Problems*, ed. by S. Saito et al. Center for Academic Societies, Osaka, Japan.

Lettau, M., and S. Ludvigson (2001): "Consumption, Aggregate Wealth and Expected Stock Returns," *Journal of Finance*, 56(3), 815–849.

A Appendix

A.1 Proof of Theorem 4.1:

PART (I): CONDITIONAL RETURNS

By Lemma 3.3 of Hansen and Richard (1987), the second moment minimization problem for conditional returns can be solved conditionally. Thus, for any given realization of the conditioning information, we seek to find θ such as to

$$\text{minimize} \quad \theta' \Lambda \theta \qquad \text{subject to} \quad e'\theta = 1$$

We set up the (conditional) Lagrangean,

$$L(\theta) = \frac{1}{2}\big(\theta' \Lambda \theta\big) - \alpha\big(e'\theta - 1\big)$$

where α is the Lagrangean multiplier for the conditional portfolio constraint. The first-order condition with respect to θ for the minimization problem is,

$$\Lambda \theta = \alpha e \qquad \text{which implies} \quad \theta = \alpha \Lambda^{-1} e$$

To determine the Lagrangean multiplier α, we use the portfolio constraint,

$$1 = e'\theta = \alpha(e'\Lambda^{-1}e) = \alpha A \qquad \text{which implies} \quad \theta = \frac{1}{A}\Lambda^{-1}e$$

To compute the moments of r^*, we use the law of iterated expectations to obtain,

$$E\big(r^*\big) = E\big(\mu'\theta\big) = E\big(\frac{B}{A}\big) \qquad \text{and} \quad E\big(r^{*2}\big) = E\big(\theta'\Lambda\theta\big) = E\big(\frac{A}{A^2}\big)$$

This completes the proof for conditional returns. $\square$

Part (ii): Unconditional Returns

We use calculus of variation. Suppose θ is a solution, and ϕ is an arbitrary vector of (managed) weights. Define,

$$\theta_\varepsilon = (1-\varepsilon)\theta + \varepsilon \frac{\phi}{E\big(e'\phi\big)}$$

By normalization, θ_ε is an admissible perturbation in the sense that it generates a one-parameter family of unconditional returns. Since θ solves the minimization problem, the following first-order condition must hold,

$$\frac{d}{d\varepsilon}\bigg|_{\varepsilon=0} E\big(\theta_\varepsilon' \Lambda \theta_\varepsilon\big) = 0$$

$$\text{which implies} \quad 0 = E\big(\theta' \Lambda \big[E\big(e'\phi\big)\theta - \phi\big]\big) = E\big(\big[E\big(\theta'\Lambda\theta\big)e' - \theta'\Lambda\big]\phi\big)$$

Since this equation must hold for every ϕ, it implies,

$$\theta = E\big(\theta'\Lambda\theta\big)\Lambda^{-1}e =: \alpha\Lambda^{-1}e$$

To determine the normalization constant α, we use the portfolio constraint,

$$1 = E\big(e'\theta\big) = \alpha E\big(e'\Lambda^{-1}e\big) = \alpha a \qquad \text{which implies} \quad \theta = \frac{1}{a}\Lambda^{-1}e$$

To compute the moments of $\hat{r}^*$, we use the law of iterated expectations to obtain,

$$E\big(\hat{r}^*\big) = E\big(\mu'\theta\big) = \big(\frac{b}{a}\big) \qquad \text{and} \quad E\big(\hat{r}^{*2}\big) = E\big(\theta'\Lambda\theta\big) = \big(\frac{a}{a^2}\big)$$

This completes the proof of Theorem 4.1. □

A.2 Proof of Theorem 4.2:

Part (i): Conditional Returns

To determine z^*, we use the fact that it is the Riesz representation of the conditional expectation on the space of excess returns. Since any excess return can be written as $z = (z + r^*) - r^* =: r - r^*$, this implies

$$E\big((r - r^*)(z^* - 1) \mid \mathcal{G} \big) = 0 \qquad \text{for all} \quad r \in R$$

Write $z^* = \tilde{R}'\theta$ and $r = \tilde{R}'\phi/(e'\phi)$ for some arbitrary vector of weights ϕ. Using the conditional moments and the fact that z^* is conditionally orthogonal to r^*, we obtain,

$$0 = E\big(rz^* - (r - r^*) \mid \mathcal{G} \big) = \frac{\theta'\Lambda\phi}{e'\phi} - \mu'\big(\frac{\phi}{e'\phi} - \frac{1}{A}\Lambda^{-1}e \big)$$

$$\text{which implies} \quad \big[\Lambda\theta - (\mu - \frac{B}{A}e) \big]'\phi = 0$$

Since this equation must hold for any ϕ, it implies,

$$\theta = \Lambda^{-1}\big(\mu - \frac{B}{A}e \big)$$

To compute the expectation of z^*, we calculate,

$$E\big(z^* \big) = E\big(\mu'\theta \big) = E\big(\mu'\Lambda^{-1}(\mu - \frac{B}{A}e) \big) = E\big(D - \frac{B^2}{A} \big)$$

This completes the proof for conditional returns. □

Part (ii): Unconditional Returns

For unconditional returns, $\hat{z}^*$ is the Riesz representation of the unconditional expectation on the space of excess returns. Hence,

$$E\big((r - \hat{r}^*)(\hat{z}^* - 1) \big) = 0 \qquad \text{for all} \quad r \in \hat{R}$$

As before, we write $\hat{z}^* = \tilde{R}'\theta$ and $r = \tilde{R}'\phi / E\big(e'\phi\big)$ for some arbitrary ϕ. Using the law of iterated expectations and the fact that $\hat{z}^*$ is orthogonal to $\hat{r}^*$, we obtain,

$$0 = E\big(r\hat{z}^* - (r - \hat{r}^*)\big) = E\Big(\frac{\theta'\Lambda\phi}{E\big(e'\phi\big)} - \mu'\Big(\frac{\phi}{E\big(e'\phi\big)} - \frac{1}{a}\Lambda^{-1}e\Big)\Big)$$

$$\text{which implies} \quad E\Big(\big[\theta - (\mu - \frac{b}{a}e)\big]'\phi\Big) = 0$$

Since this equation must hold for any ϕ, it implies,

$$\theta = \Lambda^{-1}\big(\mu - \frac{b}{a}e\big)$$

To compute the expectation of $\hat{z}^*$, we calculate,

$$E\big(\hat{z}^*\big) = E\big(\mu'\theta\big) = E\big(\mu'\Lambda^{-1}(\mu - \frac{b}{a}e)\big) = \big(d - \frac{b^2}{a}\big)$$

This completes the proof of Theorem 4.2. □

A.3 Proof of Theorem 5.1:

We apply Theorem 4.1 to the case with risk-free asset. To do this, extend the vector of base asset returns $\tilde{R}^+ = (r_f, \tilde{R})'$. Denote by $\mu_+ = (r_f, \mu)'$ the extended conditional mean vector, and by Λ_+ the corresponding second-moment matrix,

$$\Lambda_+ = \begin{pmatrix} r_f^2 & r_f\mu' \\ r_f\mu & \Lambda \end{pmatrix}$$

Using partitioned matrix inversion, we obtain,

$$\Lambda_+^{-1} = \frac{1}{r_f^2}\begin{pmatrix} (1 + \mu'\Sigma^{-1}\mu) & -r_f \cdot \mu'\Sigma^{-1} \\ -r_f \cdot \Sigma^{-1}\mu & r_f^2 \cdot \Sigma^{-1} \end{pmatrix}$$

Part (i): Conditional Returns

Let $e_+ = (\,1,\, e\,)'$ denote the extended $(n+1)$-vector of 'ones'. Following Theorem 4.1, we define the normalization constant,

$$A_+ = e_+'\Lambda_+^{-1}e_+ = \frac{1}{r_f^2}\big[\,1 + (\mu - r_f e)'\Sigma^{-1}(\mu - r_f e)\,\big] = \frac{1+H^2}{r_f^2}$$

Write $r^* = \theta^0 r_f + \tilde{R}'\theta$ and set $\theta_+ = (\,\theta^0,\, \theta\,)'$. From Theorem 4.1 (i), we obtain,

$$\theta_+ = \frac{1}{A_+}\Lambda_+^{-1}e_+ \qquad \text{which implies} \quad \theta = -\frac{r_f}{1+H^2}\Sigma^{-1}\big(\,\mu - r_f e\,\big)$$

$$\text{and} \quad \theta^0 = \frac{1 + \mu'\Sigma^{-1}(\,\mu - r_f e\,)}{1+H^2} = \frac{1 + H^2 + r_f \cdot e'\Sigma^{-1}(\,\mu - r_f e\,)}{1+H^2} = 1 - e'\theta$$

To compute the moments of r^*, we define the analogue of the efficient set constant B for the extended asset vector,

$$B_+ = \mu_+'\Lambda_+^{-1}e_+ = \frac{1}{r_f}$$

Using the moments from Theorem 4.1 (i), we obtain,

$$E\big(\,r^*\,\big) = E\big(\,\frac{B_+}{A_+}\,\big) = E\big(\,\frac{r_f}{1+H^2}\,\big) \qquad \text{and} \quad E\big(\,r^{*2}\,\big) = E\big(\,\frac{1}{A_+}\,\big) = E\big(\,\frac{r_f^2}{1+H^2}\,\big)$$

This completes the proof for conditional returns. □

Part (ii): Unconditional Returns

For unconditional returns, the normalization constant becomes,

$$a_+ = E\big(\,A_+\,\big) = \frac{1+h^2}{r_f^2}$$

Write $\hat{r}^* = \theta^0 r_f + \tilde{R}'\theta$ and set $\theta_+ = (\,\theta^0,\, \theta\,)'$. From Theorem 4.1 (ii), we obtain,

$$\theta_+ = \frac{1}{a_+}\Lambda_+^{-1}e_+ \qquad \text{which implies} \quad \theta = -\frac{r_f}{1+h^2}\Sigma^{-1}\big(\,\mu - r_f e\,\big)$$

$$\text{and} \quad \theta^0 = \frac{1 + \mu'\Sigma^{-1}(\,\mu - r_f e\,)}{1 + h^2} = \frac{1 + H^2 + r_f \cdot e'\Sigma^{-1}(\,\mu - r_f e\,)}{1 + h^2} = \frac{1 + H^2}{1 + h^2} - e'\theta$$

To compute the moments, note that $b_+ = E\big(\,B_+\,\big) = 1/r_f$. Using the moments from Theorem 4.1 (ii), we obtain,

$$E\big(\,\hat{r}^*\,\big) = \big(\,\frac{b_+}{a_+}\,\big) = \big(\,\frac{r_f}{1 + h^2}\,\big) \qquad \text{and} \quad E\big(\,\hat{r}^{*2}\,\big) = \big(\,\frac{1}{a_+}\,\big) = \big(\,\frac{r_f^2}{1 + h^2}\,\big)$$

This completes the proof of Theorem 5.1. □