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

Abhay Abhyankar

University of Durham

Devraj Basu

Alexander Stremme

Warwick Business School, University of Warwick

First Draft: May 2002; This Version: December 2004

Currently under Review; Please do not Quote or Distribute

address for correspondence:

Alexander Stremme

Warwick Business School

University of Warwick

Coventry, CV4 7AL

United Kingdom

e-mail: `alex.stremme@wbs.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 London School of Economics, the Said Business School, University of Oxford, Imperial College London, 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

December 2004

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 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 (2004). 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

Keywords: Asset Pricing, Portfolio Efficiency, Stochastic Discount Factors

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. In particular in light of the mounting evidence for asset return predictability, 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 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 conditional performance measures, the study of the economic value of asset return predictability, and tests of conditional 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 orthogonal parameterization of the space of returns on actively managed portfolios, which enables 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 (2003). Our unified approach shows that both sets of bounds can be constructed in very much the same manner, and thus facilitates a direct comparison of their respective theoretical and statistical properties.

Moreover, we show that discount factor bounds can be obtained in two different ways; either directly from the moments of the traded asset returns, or as the variance of a particular efficient actively managed portfolio. The latter is important as it provides a lower bound for the variance of the discount factor even if moments are mis-specified. This is of particular importance when analyzing the out-of-sample performance of asset pricing models, as it provides a non-parametric test for the pricing kernel. Finally, we derive a new decomposition, allowing us to separate the effect of location and shape of the efficient frontier on the level of the bounds. This is important because empirical evidence shows that the location of the frontier (as determined by the moments of the global minimum variance portfolio) can be estimated much more accurately than its slope (Haugen 1997).

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 portfolio that attains the discount factor bounds and thus provide an alternative implementation of the bounds that constitutes a valid test even when the model is incorrectly specified. Moreover, our formulation of the weights of this portfolio facilitates the analysis of their behavior in response to changes in conditioning information. This is important because it enables us to shed light on what drives the different sampling properties of the various sets of discount factor bounds. Our results have many other empirical applications, including the analysis of the optimal use of asset return predictability, tests of conditional asset pricing models, and the study of conditional measures of portfolio performance.

Mean-variance efficiency, together with the stochastic discount factor approach, are at the heart of modern empirical asset pricing, (see Ferson (2003) 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 the traditional constrained optimization approach. 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, (Cochrane 1999). Understanding how to use conditioning information *optimally* is necessary to construct actively managed portfolios that exploit return predictability, improve the power of statistical tests of asset pricing models, and refine measures of portfolio performance.

Our work is related to Gallant, Hansen, and Tauchen (1990) (GHT), and Bekaert and Liu (2004). 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 (2004) 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. Moreover, we explicitly construct the managed portfolio that attains these bounds. This is important as the variance of this portfolio is a valid lower bound for the variance of the discount factor even when the moments of asset returns are mis-specified or incorrectly estimated.

Our work is also related to Ferson and Siegel (2001), who characterize the unconditionally efficient frontier of 'conditional' returns. 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. Our explicit construction of the portfolio weights also allows us to compare conditionally and unconditionally efficient strategies. We show that the optimal unconditionally efficient portfolio corresponds to a conditionally optimal strategy only if the investor's risk aversion is time-varying. This has important implications for the analysis of portfolio performance when returns are predictable.

Ferson and Siegel (2003) 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 the moments of asset returns, 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. Moreover, as for the GHT bounds, we construct the actively managed portfolio that attains the UE bounds, thus providing a robust, non-parametric test that is valid even if the model is mis-specified. Our analysis of the behavior of the efficient weights provides an explanation for the fact that the UE bounds, while theoretically inferior, have better sampling properties than the GHT bounds (**?**).

The remainder of this paper is organized as follows; In Section 2, we provide an overview of the relevant asset pricing theory, and derive a generic expression for the various classes of discount factor bounds. In Section 3 we develop our main theoretical results, providing a portfolio-based characterizations and an intuitive decomposition of the bounds. In the subsequent section, we operationalize our theoretical results by explicitly characterizing the weights of efficient portfolios, and deriving expressions of the bounds in terms of the moments of the base asset returns. Section 5 briefly reviews the analogous results in the case where a risk-free asset is traded. Section 6 concludes.

2 Asset Pricing with Conditioning Information

In this section, we provide a brief outline of the underlying asset pricing theory, and establish our notation. We first construct the space of state-contingent pay-offs, and within it the space of traded pay-offs, attainable by actively managed portfolio strategies whose weights are functions of the conditioning information.

2.1 Set-Up and Notation

We 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 time t. Denote by L^2_t the space of all $\mathcal{F}_t$-measurable random variables that are square-integrable with respect to P. We interpret Ω as the set of 'states of nature', and L^2_t as the space of all (not necessarily attainable) state-contingent pay-offs at time t.

Conditioning Information:

To incorporate conditioning information, we take as given a sub-σ-field $\mathcal{G}_{t-1} \subseteq \mathcal{F}_{t-1}$. 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*[1], variables that are observable at time $t-1$. To simplify notation, we denote by $E_{t-1}(\,\cdot\,)$ the conditional expectation operator with respect to $\mathcal{G}_{t-1}$.

Traded Assets and Managed Pay-Offs:

There are n tradeable risky assets, indexed $k = 1 \ldots n$. We denote the gross return (per dollar invested) of the k-th asset by $r^k_t \in L^2_t$, and by $\tilde{R}_t := (\, r^1_t \,\ldots\, r^n_t \,)'$ the n-vector of risky asset returns. Unless stated otherwise, we assume that no risk-free asset is traded. We define X_t as the space of all elements $x_t \in L^2_t$ that can be written in the form, $x_t = \tilde{R}'_t\, \theta_{t-1}$, with $\theta_{t-1} = (\, \theta^1_{t-1} \,\ldots\, \theta^n_{t-1} \,)'$, where $\theta^k_{t-1} \in L^2_t$ are $\mathcal{G}_{t-1}$-measurable functions. We interpret X_t as the space of 'managed' pay-offs, obtained by forming combinations of the base assets with weights θ^k_{t-1} that are functions of the conditioning information[2].

---

[1]Examples of such variables considered in the literature include dividend yield (Fama and French 1988), interest rate spreads (Campbell 1987), or consumption-wealth ratio (Lettau and Ludvigson 2001).

[2]Note that, in contrast to the fixed-weight case without conditioning information, the space of *managed* pay-offs is infinite-dimensional even when there is only a finite number of base assets.

Pricing Function:

Since the base assets are characterized by their *returns*, we set $\Pi_{t-1}(\, r_t^k \,) \equiv 1$, and extend Π_{t-1} to all of X_t by conditional linearity, $\Pi_{t-1}\,(\, x_t \,) = e'\,\theta_{t-1}$ for $x_t = \tilde{R}_t'\,\theta_{t-1} \in X_t$, where $e = (\, 1 \, \ldots \, 1 \,)'$ is an n-vector of 'ones'. Throughout the remainder of this paper, we will refer to $(\, X_t, \, \Pi_{t-1} \,)$ as the *conditional market model* generated by the base assets $\tilde{R}_t$ and the chosen conditioning set $\mathcal{G}_{t-1}$. Finally, we set $\Pi_0(\, x_t \,) = E(\, \Pi_{t-1}(\, x_t \,)\,)$. By construction, both Π_{t-1} and Π_0 are (conditionally) linear and thus satisfy the 'law of one price', a weak from of no-arbitrage condition.

2.2 Stochastic Discount Factors

Stochastic discount factors (SDFs) are a convenient way of describing an asset pricing model. They can be characterized in terms of the following fundamental pricing equation;

Definition 2.1 *By an admissible* ***stochastic discount factor*** *(SDF) for the conditional market model* $(\, X_t, \, \Pi_{t-1} \,)$, *we mean an element* $m_t \in L_t^2$ *such that*

$$E_{t-1}\big(\, m_t \tilde{R}_t \,\big) = e, \quad \textit{where} \quad e = (\, 1 \, \ldots \, 1 \,)' \quad \textit{is an } n\textit{-vector of 'ones'}. \tag{1}$$

Note that (1) implies that m_t also prices all managed pay-offs (conditionally) correctly, that is $E_{t-1}(\, m_t x_t \,) = \Pi_{t-1}(\, x_t \,)$ for all $x_t = \tilde{R}_t'\,\theta_{t-1} \in X_t$. Taking unconditional expectations,

$$E(\, m_t x_t \,) = E(\, \Pi_{t-1}(\, x_t \,)\,) = \Pi_0(\, x_t \,) \tag{2}$$

In other words, any SDF that prices the base assets (conditionally) correctly must necessarily be consistent with the *'generalized'* pricing function $\Pi_0(\, \cdot \,)$. Thus, any subspace $R_t \subset X_t$ for which $\Pi_0 \equiv 1$ on R_t can be considered as a space of *returns*, attainable by a corresponding set of *managed portfolios* in a generalized sense.

Discount Factor Bounds:

For different choices of θ_{t-1} (and hence different $x_t \in X_t$), we thus obtain from (2) a family of testable 'moment conditions' of the form $E(\, m_t\, \tilde{R}_t'\, \theta_{t-1}\,) = E(\, e'\, \theta_{t-1}\,)$ that any candidate SDF must satisfy. While these can be used in many different ways (e.g. GMM) to estimate or test asset pricing models, most of these tests yield necessary but not sufficient conditions[3]. Discount factor *bounds*, first introduced by Hansen and Jagannathan (1991), are one class of such necessary conditions. They are lower bounds on the variance of an SDF, as a function of its mean. Such bounds are a useful diagnostic in that if a candidate does not satisfy the bounds, then it cannot be an admissible SDF. In the extended case with conditioning information, the bounds can be formulated in their most general form as,

Lemma 2.2 *Let* $m_t \in L_t^2$ *with* $\nu = E(\, m_t\,)$*, and* $R_t \subset X_t$ *any arbitrary subspace of* X_t *with* $\Pi_0 \equiv 1$ *on* R_t*. Then, necessary (not sufficient) for* m_t *to be an admissible SDF is,*

$$\frac{\sigma(\, m_t\,)}{\nu} \geq \sup_{r_t \in R_t} \frac{E(\, r_t\,) - 1/\nu}{\sigma(\, r_t\,)} =: \lambda_*(\, \nu\,;\, R_t\,), \tag{3}$$

PROOF: Since (2) is a necessary condition for m_t to be an admissible SDF, we can write $1 = E(\, m_t r_t\,) = \rho \cdot \sigma(\, m_t\,)\, \sigma(\, r_t\,) + \nu \cdot E(\, r_t\,)$ for any $r_t \in R_t$, where ρ is the correlation between m_t and r_t. The inequality then follows trivially from $-1 \leq \rho \leq 1$. □

Note that, if an (unconditionally) risk-free asset was traded with gross return r_f, then any admissible SDF would have to satisfy $r_f = 1/\nu$. Therefore, we refer to $1/\nu$ as the 'shadow' risk-free rate implied by the mean $\nu = E(\, m_t\,)$ of the candidate SDF m_t. The right-hand side $\lambda_*(\, \nu\,;\, R_t\,)$ in (3) can hence be interpreted as the maximum *generalized Sharpe ratio* on R_t, relative to the shadow risk-free rate $1/\nu$. As a consequence, any return $r_t \in R_t$ that attains the supremum in (3) must necessarily be unconditionally mean-variance efficient (i.e. have minimal *unconditional* variance for given unconditional mean) in R_t. We use this fact in the following section to derive a portfolio-based characterization of the bounds.

CLASSIFICATION OF DISCOUNT FACTOR BOUNDS:

---

[3]This is because the space X_t on which the SDF must be tested is infinite-dimensional.

While Lemma 2.2 provides a generic characterization, the different classes of SDF bounds considered in the literature are obtained by choosing different return spaces $R_t \subset X_t$ in (3):

(i) **HJ Bounds:** The Hansen and Jagannathan (1991) (HJ) bounds without conditioning information are obtained from (3) by choosing R_t as the space of *fixed-weight returns*,

$$R_t^0 = \left\{ x_t = \tilde{R}_t' \theta, \text{ where } \theta \in \mathbb{R}^n \text{ with } e' \theta = 1 \right\}$$

(ii) **UE Bounds:** The 'Unconditionally Efficient' (UE) bounds of Ferson and Siegel (2003) are obtained from (3) by choosing R_t as the space of *'conditional returns'*,

$$R_t^C = \left\{ x_t = \tilde{R}_t' \theta_{t-1}, \text{ where } \theta_{t-1} \text{ is } \mathcal{G}_{t-1}\text{-measurable with } e' \theta_{t-1} \equiv 1 \right\}$$

(iii) **GHT Bounds:** The Gallant, Hansen, and Tauchen (1990) (GHT) bounds, and hence also their implementation as the 'optimally scaled' bounds by Bekaert and Liu (2004) are obtained from (3) by choosing R_t as the space of *'generalized returns'*,

$$R_t^G = \left\{ x_t = \tilde{R}_t' \theta_{t-1}, \text{ where } \theta_{t-1} \text{ is } \mathcal{G}_{t-1}\text{-measurable with } E(e' \theta_{t-1}) = 1 \right\}$$

The term 'conditional returns' in (ii) is used to reflect the fact that the portfolio constraint $\Pi_{t-1}(x_t) = e' \theta_{t-1} \equiv 1$ is required to hold conditionally, i.e. for all realizations of the conditioning information. Conversely, the term 'generalized returns' in (iii) indicates the fact that $\Pi_0(x_t) = E(e' \theta_{t-1})$ does not reflect the true price for the pay-off x_t but rather its *expected cost*. Note however that, by (2), any admissible SDF that prices the base assets (conditionally) correctly, must also necessarily price all generalized returns correctly to one. Finally note that, since $R_t^G \subset X_t$ is the largest possible subspace on which $\Pi_0 \equiv 1$, the GHT bounds are by construction the sharpest possible bounds for given set $\mathcal{G}_{t-1}$ of conditioning information. In other words, we expect

$$\lambda_*(\nu ; R_t^0) < \lambda_*(\nu ; R_t^C) < \lambda_*(\nu ; R_t^G).$$

Of course, from a theoretical point of view it is therefore optimal to work with the GHT bounds as these provide the most powerful test for an asset pricing model. However, empirical studies (?) show that, while the UE bounds are statistically indistinguishable from the GHT bounds, the former have better sampling properties.

3 Stochastic Discount Factor Bounds

To construct the bound for a given mean $\nu = E(m_t)$ of the candidate discount factor, we need to find the portfolio that maximizes the generalized Sharpe ratio in (3). In this section, we provide a generic construction of this portfolio and hence the bounds, which is valid for any space of returns. For what follows, we denote by $R_t \subset X_t$ any subspace on which $\Pi_0 \equiv 1$, including in particular the three spaces R_t^0, R_t^C and R_t^G or fixed-weight, conditional and generalized returns, respectively, as defined in the preceding section.

3.1 Generic Discount Factor Bounds

Any return $r_t \in R_t$ that attains the supremum in (3) must necessarily be unconditionally mean-variance efficient in R_t. Hence, we need to characterize the efficient frontier in R_t. It follows from Hansen and Richard (1987) that every unconditionally efficient $r_t \in R_t$ can be written in the form $r_t = r_t^* + w \cdot z_t^*$ for some $w \in \mathbb{R}$, where $r_t^* \in R_t$ is the unique return orthogonal[4] to $Z_t = \Pi_0^{-1}\{0\} \subset X_t$, and $z_t^* \in Z_t$ is a canonically chosen *excess* (i.e. zero cost) return. In other words, the efficient frontier in R_t is spanned by r_t^* and z_t^*.

We modify this construction and consider instead the unique return r_t^0 that is orthogonal to Z_t with respect to the *covariance* inner product[5], i.e. $\text{cov}(r_t^0, z_t) = 0$ for all $z_t \in Z_t$. Note that r_t^0 is nothing other than the *global minimum variance* (GMV) return[6]. In analogy with

[4]One can also define r_t^* as the return with minimum unconditional second moment.

[5]In the absence of a risk-free asset, the covariance functional is indeed a well-defined inner product.

[6]This follows directly from the first-order condition of the unconstrained variance minimization problem.

the Hansen and Richard (1987) construction, we choose $z_t^0 \in Z_t$ so that $E(z_t) = \mathrm{cov}(z_t^0, z_t)$ for all $z_t \in Z_t$, i.e. z_t^0 is the Riesz representation of the expectation functional on Z_t. Given the defining properties of r_t^0 and z_t^0, it is now easy to show that[7],

$$r_t^0 = r_t^* + \frac{E(r_t^*)}{1 - E(z_t^*)} \cdot z_t^* \qquad \text{and} \qquad z_t^0 = \frac{1}{1 - E(z_t^*)} \cdot z_t^* \tag{4}$$

Note that by construction, r_t^0 and z_t^0 are linear combinations of r_t^* and z_t^*, and hence also span the mean-variance efficient frontier. In our parametrization, the GMV r_t^0 may be regarded as a measure of *location*, while z_t^0 determines the *shape* of the frontier. We discuss the differences between the r_t^* and the r_t^0 parameterizations in Section 3.4. We are now ready to state our first main result,

Theorem 3.1 *Necessary for $m_t \in L_t^2$ with $\nu = E(m_t)$ to be an admissible SDF is,*

$$\sigma^2(m_t) \geq \frac{(\gamma_1^2 + \gamma_2\gamma_3) \cdot \nu^2 - 2\gamma_1 \cdot \nu + 1}{\gamma_2}, \tag{5}$$

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, with coefficients that are functions of the unconditional moments of the GMV return r_t^0 and the canonical excess return z_t^0.

PROOF OF THEOREM 3.1: Follows directly from Proposition 3.2 below. □

3.2 Portfolio-Based Characterization of the Bounds

Because the bounds in (5) are expressed in terms of the moments of r_t^0 and z_t^0, they constitute a valid test only if these moments are correctly specified. In particular, we show below that

[7] 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.

the unconditional moments of r_t^0 and z_t^0 are functions of the *conditional* moments of the base asset returns, which may be difficult to estimate and, more importantly, are subject to model specification risk. To eliminate this weakness, we show below (Lemma 3.3) that the bounds can also be obtained as the variance of a managed portfolio. The advantage of this characterization is that the (in-sample) variance of this portfolio *always* constitutes a lower bound for the variance of an SDF, even if moments are mis-specified. We begin by proving the fundamental proposition from which our main results follow.

Proposition 3.2 *For given $\nu = E(m_t)$, the maximum generalized Sharpe ratio $\lambda_*(\nu)$ that attains the discount factor bound in (3), admits a decomposition of the form,*

$$\lambda_*^2(\nu) = \lambda_0^2(\nu) + E(z_t^0), \qquad \text{with} \quad \lambda_0(\nu) = \frac{E(r_t^0) - 1/\nu}{\sigma(r_t^0)}. \tag{6}$$

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

$$r_t^\nu = r_t^0 + \kappa^*(\nu) \cdot z_t^0, \qquad \text{with} \quad \kappa^*(\nu) = \frac{\sigma^2(r_t^0)}{E(r_t^0) - 1/\nu}. \tag{7}$$

PROOF: Appendix A.1. □

To our knowledge, this result 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 our approach 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, using (4) our decomposition (6) can be re-written as,

$$\lambda_*^2(\nu) = \lambda_0^2(\nu) + \frac{E(z_t^*)}{1 - E(z_t^*)}. \tag{8}$$

We thus generalize Equation (16) of Jagannathan (1996), to the case without risk-free asset.

REMARK: Equation (6) means that the maximum generalized Sharpe ratio that attains the discount factor bounds is driven by two distinct components; the generalized Sharpe ratio $\lambda_0^2(\nu)$ of the GMV (which measures the location of the efficient frontier), and the term

$E(z_t^0)$ which captures the shape of the frontier (the higher this value, the 'wider' is the frontier). In empirical applications, our decomposition enables us to separate the effect that these two factors have on both location as well as the statistical properties of discount factor bounds (see also Section 3.4 below).

Finally, note also that the return r_t^ν defined in (7) is the unique efficient return with *zero-beta* rate $1/\nu$, i.e. the 'tangency' portfolio relative to the 'shadow' risk-free rate $1/\nu$.

3.3 Relationship to 'Optimally Scaled' Returns

From 3.1 we know that the discount factor bounds can be obtained as a function of the (unconditional) moments of r_t^0 and z_t^0. We will show below (Section 4) that these moments in turn are functions of the *conditional* moments of the base asset returns. As the latter are notoriously difficult to estimate non-parametrically, any implementation of the bounds based directly on the moments may be subject to considerable measurement and/or model specification error.

In the case of the GHT bounds, Bekaert and Liu (2004) provide an alternative derivation that obtains the bounds as the variance of an 'optimally scaled' payoff, given in Equation (22) of their paper. Their derivation is closely related to ours, as the following lemma shows;

Lemma 3.3 *Necessary for $m_t \in L_t^2$ with $\nu = E(m_t)$ to be an admissible SDF is,*

$$\sigma(m_t) \geq \frac{\nu}{\kappa^*(\nu)} \cdot \sigma(r_t^0 + \kappa^*(\nu) \cdot z_t^0), \tag{9}$$

where $\kappa^*(\nu) \in \mathbb{R}$ is defined as in Proposition 3.2 above.

Proof: The proof of this lemma follows directly from Proposition 3.2.

It is now easy to show that in the case of generalized returns, the right-hand side of (9) is equivalent to the 'optimal scaling' transformation used in Bekaert and Liu (2004) to implement the GHT bound. However, our unified approach implies that the same approach is valid for all sets of discount factor bounds, including the 'UE' bounds of Ferson and Siegel (2003) as well as the HJ bounds in the fixed-weight case. Lemma 3.3 is important because it shows that discount factor bounds can be obtained by estimating the variance of a specific managed return. In particular, this facilitates the analysis of the *out-of-sample* performance of a candidate asset pricing model.

Note moreover that, when m_t is indeed an admissible SDF, the optimally scaled payoff in (9) can in fact be identified as the unconditional projection of m_t onto the space of managed payoffs X_t, since

$$\frac{\nu}{\kappa^*(\nu)} \cdot r_t^\nu = \frac{E(m_t)E(r_t^0) - 1}{\sigma^2(r_t^0)} \cdot r_t^0 + E(m_t) \cdot z_t^0 = -\operatorname{proj}\big(m_t \,|\, X_t\big)$$

When moments are correctly specified, the bounds are obtained as the variance of this projection, as in (9). This is in fact the original definition of the GHT bounds used in Gallant, Hansen, and Tauchen (1990). Moreover, even when the conditional moments are incorrectly estimated, the variance of the optimally scaled return still provides a valid lower bound to the variance of pricing kernels. Our analysis shows that this property not only holds for the GHT bounds, but indeed for all the different classes of bounds considered in the literature.

Our unified approach shows that the UE and GHT bounds (and also the HJ bounds in the case without conditioning information) can be constructed in very much the same manner. It is also easy to see that theoretically, the UE bounds will plot below the GHT bounds as they are obtained from the maximum Sharpe ratio in the space R_t^C of *conditional* returns, which is contained in the space R_t^G of generalized returns. However, in their empirical analysis **?**) show that, while the difference between the UE and GHT bounds is statistically insignificant, the UE bounds possess better sampling properties. This is because the portfolio weights of the efficient conditional return (7) display a more 'conservative' response to extreme changes in the conditioning instruments than those of the respective generalized return (see also Section 4 below).

3.4 Comparison of the two Parameterizations

As mentioned before, the efficient frontier with conditioning information can also be parameterized using the minimum second moment return r_t^* and the corresponding excess return z_t^*. This parametrization forms the basis of the Hansen and Richard (1987) analysis. Our parametrization is based instead on the GMV portfolio r_t^0 and the appropriate covariance-orthogonal excess return z_t^0. Geometrically, the GMV captures the *location* of the frontier, while z_t^0 determines its *shape*. Apart from dramatically simplifying the derivation of the bounds, our representation has other interesting properties.

?) investigate the sampling properties of the building blocks of both these parameterizations. While both r_t^* and z_t^* are similarly sensitive to sampling variability and measurement error, their findings for the GMV-based parametrization are very different. The estimates of the moments of the GMV are very robust and almost unaffected by measurement error; nearly all of the variability that shows up in the estimates of the bounds is due to z_t^0. In other words, the *location* of the efficient frontier seems very stable, while its shape (curvature) is much less robust. These results resemble those in the fixed-weight case (Haugen 1997), where the GMV is the portfolio whose moments can be measured most accurately.

4 Explicit Construction of Discount Factor Bounds

We now operationalize the results of the preceding section by explicitly constructing the portfolio weights of the efficient return r_t^ν that attains the SDF bounds, both for conditional and generalized returns. In the process, we recover the efficient portfolio weights stated in Ferson and Siegel (2001), and provide a characterization of the efficient frontier implicit in Bekaert and Liu (2004). Using our results from Section 3, we obtain explicit expressions for the UE bounds of Ferson and Siegel (2003), and Bekaert and Liu's (2001) implementation of the GHT bounds. Our unified framework helps clarify the relationship between these two sets of bounds.

4.1 Efficient 'Conditional' Returns and UE Bounds

We first study the properties of the efficient frontier for conditional returns. Using our results from Section 3, we then characterize the return that attains the maximum generalized Sharpe ratio in (3), from which we obtain an explicit expression for the 'unconditionally efficient (UE)' bounds. To begin with, we define,

$$\mu_{t-1} = E_{t-1}\big(\tilde{R}_t\big), \qquad \text{and} \qquad \Lambda_{t-1} = E_{t-1}\big(\tilde{R}_t \cdot \tilde{R}_t'\big). \tag{10}$$

In other words, returns can be written as $\tilde{R}_t = \mu_{t-1} + \varepsilon_t$, where μ_{t-1} is the conditional expectation of returns given conditioning information, and ε_t is the residual disturbance with variance-covariance matrix $\Sigma_{t-1} = \Lambda_{t-1} - \mu_{t-1}\mu_{t-1}'$. This is the formulation of the model with conditioning information used in Ferson and Siegel (2001)[8]. Finally, we set

$$A_{t-1} = e'\Lambda_{t-1}^{-1}e, \qquad B_{t-1} = \mu_{t-1}'\Lambda_{t-1}^{-1}e, \qquad D_{t-1} = \mu_{t-1}'\Lambda_{t-1}^{-1}\mu_{t-1} \tag{11}$$

These are the conditional versions of the 'efficient set' constants from classic mean-variance theory. We choose this notation in order to highlight the structural similarities between the UE and GHT bounds, and to facilitate a direct comparison.

(A) Efficient Portfolio Weights

Ferson and Siegel (2001) describe the efficient frontier in terms of a set of constants α_1, α_2 and α_3. In our notation, we can write these as[9]

$$\alpha_1 = E(\,B_{t-1}/A_{t-1}\,), \quad \alpha_2 = E(\,1/A_{t-1}\,), \quad \text{and} \quad \alpha_3 = E(\,D_{t-1} - B_{t-1}^2/A_{t-1}\,).$$

Note that, using Proposition A.1, we can identify these constants as $E\,(\,r_t^*\,) = \alpha_1$ and $\sigma^2\,(\,r_t^*\,) = \alpha_2 - \alpha_1^2$. In other words, α_1 and α_2 are the first and second moments of the

[8]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.

[9]Note that our notation differs slightly from that used in Ferson and Siegel (2001), where the roles of α_1 and α_2 are reversed. Our notation is such that α_1 and α_2 are the *first* and *second* moments of r_t^*, respectively.

minimum second moment return r_t^*. Similarly, Proposition A.2 implies $\alpha_3 = E\,(\, z_t^*\,)$. Using this notation, we can now characterize the efficient frontier in the space R_t^C of conditional returns;

Lemma 4.1 *The unconditionally efficient conditional return $r_t^m \in R_t^C$ for given unconditional mean $m \in \mathbb{R}$ can be written as $r_t^m = \tilde{R}_t'\theta_{t-1}$, where*

$$\theta_{t-1} = \Lambda_{t-1}^{-1}\Big(\frac{1 - wB_{t-1}}{A_{t-1}}\, e + w\, \mu_{t-1} \Big), \quad \textit{where} \quad w = \frac{m - \alpha_1}{\alpha_3}, \tag{12}$$

While this result has been reported in Ferson and Siegel (2001), we include it here for two reasons. First, we wish to highlight the connection between the efficient set constants and the moments of r_t^* and z_t^*. Second, a direct comparison with Theorem 4.3 in the next section allows us to analyze the similarities and differences between the efficient portfolio weights for conditional and generalized returns, respectively.

Proof of Lemma 4.1: Appendix A.2. □

Note that, using (4), we can obtain the weights of the GMV return r_t^0 by setting $m = \alpha_1/(1-\alpha_3)$ in (12). Similarly, the weights of r_t^* can be obtained by setting $w = 0$.

(B) Discount Factor Bounds

The discount factor bounds are now a trivial implication of the results derived so far;

Corollary 4.2 *Necessary for $m_t \in L_t^2$ with $\nu = E(\, m_t\,)$ to be an admissible SDF is,*

$$\sigma^2(\, m_t\,) \geq \frac{(\alpha_2\alpha_3 + \alpha_1^2)\cdot \nu^2 - 2\alpha_1 \cdot \nu + (1-\alpha_3)}{\alpha_1(1-\alpha_3) - \alpha_1^2}. \tag{13}$$

In other words, the UE bound takes the form of a second-order polynomial in the mean ν of the candidate SDF, where the coefficients are functions of the efficient set constants.

Proof of Corollary 4.2: From (4), it follows trivially that $\gamma_1 = \alpha_1/(1-\alpha_3)$ and $\gamma_2 = \alpha_2 - \alpha_1^2/(1-\alpha_3)$, respectively. Similarly, we obtain $\gamma_3 = \alpha_3/(1-\alpha_3)$. Substituting this into Equation (5) yields the desired result. □

(C) Bounds from Portfolios

From Lemma 3.3 we know that the discount factor bound can also be obtained as the variance of the efficient return $r_t^{\nu} = r_t^0 + \kappa^*(\nu) \cdot z_t^0$. Using Lemma 4.1, together with the fact that r_t^{ν} has zero-beta rate $1/\nu$, we can write $r_t^{\nu} = \tilde{R}_t' \theta_{t-1}$, where

$$\theta_{t-1} = \Lambda_{t-1}^{-1}\Big(\frac{1 - wB_{t-1}}{A_{t-1}}\, e + w\, \mu_{t-1} \Big), \qquad \text{with} \quad w = \frac{\alpha_2 \nu - \alpha_1}{\alpha_1 \nu - (1 - \alpha_3)} \tag{14}$$

From the preceding section, we can identify this as the weights of the efficient return that has unconditional mean $m = \alpha_1 + w\alpha_3$. The zero-beta rate associated with this portfolio is, by construction, $1/\nu$.

There are thus two methods of obtaining discount factor bounds; either directly from the conditional moments of the base asset returns as in (13), or via the variance of a specific return using Lemma 3.3. When the conditional moments are correctly specified, these two methods yield the same answer. If, however, the conditional moments are misspecified, then the variance of the return constructed above still yields a valid lower bound for the variance of an SDF. Clearly, this is particularly useful for studying the out-of-sample properties of the bounds. Moreover, empirical studies (?) seem to indicate that the bounds from portfolios have marginally better sampling properties than the bounds obtained directly from the conditional moments.

4.2 Efficient 'Generalized' Returns and GHT Bounds

In analogy with the preceding section, we now study the properties of the efficient frontier for *generalized* returns, which has not been done previously. Using our results from Section 3, we then characterize the return that attains the maximum generalized Sharpe ratio in (3), from which we obtain an explicit expression for the GHT bounds. We also derive a portfolio-based characterization of these bounds, which establishes the link between our approach and the 'optimal scaling' approach of Bekaert and Liu (2004).

We use lowercase letters a, b and d to denote the unconditional expectations of the conditional constants A_{t-1}, B_{t-1}, and D_{t-1} introduced in the preceding section. Note that these are normalized versions of the efficient set constants defined in Bekaert and Liu (2004).

(A) Efficient Portfolio Weights

Although Bekaert and Liu (2004) do not study the efficient frontier for generalized returns, it is implicit in their 'optimal scaling' approach. To highlight the similarities between the expressions for conditional and generalized returns, we define,

$$\hat{\alpha}_1 = b/a, \quad \hat{\alpha}_2 = 1/a, \quad \text{and} \quad \hat{\alpha}_3 = d - b^2/a.$$

As before, using Propositions A.3 and A.4, we can identify $\hat{\alpha}_1$ and $\hat{\alpha}_2$ as the first and second moments of r_t^* in the case of generalized returns, and $\hat{\alpha}_3 = E(\, z_t^*\,)$.

Theorem 4.3 *The unconditionally efficient generalized return $r_t^m \in R_t^G$ for given unconditional mean $m \in \mathrm{I\!R}$ can be written as $r_t^m = \tilde{R}_t'\theta_{t-1}$, where*

$$\theta_{t-1} = \Lambda_{t-1}^{-1}\Big(\,\frac{1 - wb}{a}\,e + w\,\mu_{t-1}\,\Big), \qquad \text{where} \quad w = \frac{m - \hat{\alpha}_1}{\hat{\alpha}_3} \tag{15}$$

Proof: Appendix A.3. □

To our knowledge, this result is new. Comparing the above expression with (12), we note that the functional form of the efficient weights is identical in both cases. The only difference is that the conditional constants A_{t-1} and B_{t-1} that appear in the case of conditional returns are replaced by their unconditional counterparts. While this difference may seem marginal, it is responsible for the difference in response of the weights to extreme changes of the conditioning information.

(B) Discount Factor Bounds

The GHT bound can now be obtained in the same fashion as the UE bound;

Corollary 4.4 *Necessary for $m_t \in L_t^2$ with $\nu = E(\, m_t\,)$ to be an admissible SDF is,*

$$\sigma^2(\, m_t\,) \geq \frac{(\hat{\alpha}_2\hat{\alpha}_3 + \hat{\alpha}_1^2)\cdot\nu^2 - 2\hat{\alpha}_1\cdot\nu + (1 - \hat{\alpha}_3)}{\hat{\alpha}_1(1 - \hat{\alpha}_3) - \hat{\alpha}_1^2}. \tag{16}$$

Proof: The proof is identical to that of Corollary 4.2 in the preceding section. □

Note that (16) 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 (2004) yields the GHT bounds as a quartic over a quadratic polynomial which, if moments are correctly specified, reduces to the above expression. Specifically, substituting into (16) the corresponding expressions for the $\hat{\alpha}_i$ in terms of the constants a, b and d, we obtain Equation (25) in Bekaert and Liu (2004).

(C) Bounds from Portfolios

Using Theorem 4.3, together with the fact that r_t^ν has zero-beta rate $1/\nu$, we can write the efficient return that attains the bounds as $r_t^\nu = \tilde{R}_t' \theta_{t-1}$, where

$$\theta_{t-1} = \Lambda_{t-1}^{-1} \Big(\frac{1 - wb}{a}\, e + w\, \mu_{t-1} \Big), \qquad \text{where} \quad w = \frac{\hat{\alpha}_2 \nu - \hat{\alpha}_1}{\hat{\alpha}_1 \nu - (1 - \hat{\alpha}_3)} \tag{17}$$

Again, substituting for the $\hat{\alpha}_i$ in terms of a, b and d, it is straight-forward to show that these weights indeed coincide with the optimal scaling vector given in Equations (22) and (23) of Bekaert and Liu (2004), suitably normalized. The GHT bound is then obtained from the variance of this generalized return via (9).

4.3 Properties of the Efficient Portfolio Weights

Throughout this section, we will assume that the conditional mean is a linear function of a single conditioning instrument, $\mu_{t-1} = \mu(y_{t-1}) = \mu_0 + \beta y_{t-1}$ for some $\mathcal{G}_{t-1}$-measurable y_{t-1}. Moreover, we assume that the conditional variance-covariance matrix Σ_{t-1} of the base asset return innovations does not depend on y_{t-1} (we will hence write simply Σ). In other words, we assume that the base asset returns are given by a linear predictive regression, as in Equation (1) of Ferson and Siegel (2001).

To investigate the asymptotic properties of the efficient weights for large values of the conditioning instrument, we use a well-known matrix identity, stated for completeness in Appendix A.4. Using this identity and the definition of the efficient set constants, it is easy to see that $\Lambda_{t-1}^{-1}\mu_{t-1} \longrightarrow 0$ and $B_{t-1} \longrightarrow 0$ as $y_{t-1} \longrightarrow \pm\infty$, while both $\Lambda_{t-1}^{-1}e$ and A_{t-1} converge to finite limits. Hence, for extreme values of the instrument, the weights of the efficient *conditional* return, as given in (12), converge to

$$\theta_{t-1} \longrightarrow \frac{(\beta'\Sigma^{-1}\beta)\Sigma^{-1}e - (\beta'\Sigma^{-1}e)\Sigma^{-1}\beta}{(e'\Sigma^{-1}e)(\beta'\Sigma^{-1}\beta) - (\beta'\Sigma^{-1}e)^2} \qquad \text{as } y_{t-1} \longrightarrow \pm\infty. \tag{18}$$

These are in fact the asymptotic weights of the minimum second moment return r_t^* as it can be shown that $z_t^* \longrightarrow 0$ as $y_{t-1} \longrightarrow \pm\infty$ in the Hansen and Richard (1987) decomposition of the efficient frontier. Moreover, it is easy to see that the conditional mean of the efficient return defined by (12) converges to w as $y_{t-1} \longrightarrow \pm\infty$, similar to the case with risk-free asset. In contrast, just as in the case with risk-free asset, the conditional mean of the corresponding *conditionally* efficient strategy can be shown to diverge for extreme values of the instrument.

An argument similar to that made above shows that the weights of the efficient *generalized* return, as given in (15), converge to

$$\theta_{t-1} \longrightarrow \frac{1 - wb}{a}\left[\Sigma^{-1}e - \frac{\beta'\Sigma^{-1}e}{e'\Sigma^{-1}e}\,\Sigma^{-1}\beta\right] \qquad \text{as } y_{t-1} \longrightarrow \pm\infty. \tag{19}$$

While both the conditional as well as the generalized efficient returns converge to fixed limits for extreme values of the conditioning instrument, a numerical analysis based on estimated values for Σ and β shows that the weights of the conditional return converge much quicker towards their asymptotic values. In contrast, while the weights of the generalized return will, as we have just shown, eventually converge, they display an almost linear response for any reasonable range of values of the conditioning instrument. This difference in response can be largely attributed to the tighter portfolio constraint for conditional returns, which limits the extent to which the weights can respond to changes in conditioning information. In turn, the different response to extreme signals is largely responsible for the different sampling properties (**?**) of the UE and GHT bounds.

5 Case With Risk-Free Asset

While not the central focus of this paper, we analyze in this section the case when a risk-free asset is traded. Since in this case, discount factor bounds are trivial, we focus instead on the properties of efficient portfolios.

Assume that in addition to the n risky assets, an (unconditionally) risk-free asset is traded, whose (gross) return we denote r_f. In this case, the augmented space X_t of 'managed' pay-offs now consist of elements of the form $x_t = \theta^0_{t-1} r_f + (\tilde{R}_t - r_f e)'\theta_{t-1}$. Note that we allow portfolios that have 'managed' positions in the risk-free asset, which themselves are hence no longer risk-free, since the weight θ^0_{t-1} may vary with conditioning information. In contrast to Section 4, the weights θ_{t-1} on the risky assets are now applied to their *excess* returns, which implies that the pricing functional now takes the particularly simple form $\Pi_{t-1}(x_t) = \theta^0_{t-1}$. As a consequence, the space R^C_t of *conditional* returns in this framework is given by those pay-offs for which $\theta^0_{t-1} \equiv 1$. Conversely, the space R^G_t of *generalized* returns is defined by the (less strict) constraint $E\left(\theta^0_{t-1}\right) = 1$. We define

$$\Sigma_{t-1} = \mathrm{Var}\left(\tilde{R}_t \mid \mathcal{G}_{t-1} \right) = \Lambda_{t-1} - \mu_{t-1} \cdot \mu'_{t-1} \tag{20}$$

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 generalized returns, which is similar to Equation (16) in Jagannathan (1996).

5.1 Efficient Frontier and Sharpe Ratio

We begin by defining the analogue to the efficient set constants introduced in Section 4,

$$H^2_{t-1} = (\mu_{t-1} - r_f e)' \Sigma^{-1}_{t-1} (\mu_{t-1} - r_f e), \tag{21}$$

Similar to Section 4, we denote by $h^2 = E\left(H_{t-1}^2\right)$ the unconditional expectation of H_{t-1}^2. From classic mean-variance theory it is well-known that the quantity H_{t-1}^2 is in fact the maximum squared *conditional* Sharpe ratio (i.e. the maximum Sharpe ratio achievable by portfolios that are fixed-weight efficient relative to conditional mean and variance).

(A) Efficient Conditional Returns

When a risk-free asset is traded, the maximum Sharpe ratio λ_*^2 is in fact attained by the return r_t^* from the Hansen and Richard (1987) representation of the efficient frontier. Modifying the proof of Proposition A.1 to account for the presence of a risk-free asset, one can show[10] that the maximum (squared) Sharpe ratio in the space R_t^C of *conditional* returns is given by the expression $\lambda_*^2 = \zeta/(1-\zeta)$, where

$$\zeta = E\big(\frac{H_{t-1}^2}{1+H_{t-1}^2}\big), \tag{22}$$

as defined in Ferson and Siegel (2001). Following the same arguments as in Section 4.1, it is now straight-forward to show that the unconditionally efficient conditional return $r_t^m \in R_t^C$ for given unconditional mean $m \in \mathbb{R}$ can be written as $r_t^m = r_f + (\,\tilde{R}_t - r_f e\,)'\theta_{t-1}$, where

$$\theta_{t-1} = \frac{w - r_f}{1+H_{t-1}^2} \cdot \Sigma_{t-1}^{-1}\big(\,\mu_{t-1} - r_f e\,\big), \quad \text{with} \quad w = \frac{m - r_f(1-\zeta)}{\zeta}. \tag{23}$$

Using a simple matrix identity (see Appendix A.4), this expression can be shown to be identical to that stated in Equation (12) of Ferson and Siegel (2001). Our expression (23), while similar to the efficient portfolio weights in the absence of conditioning information, differs from the latter in that the normalization factor $1 + H_{t-1}^2$ is in fact time-varying. The presence of this time-varying quantity, an artefact of the *conditional* portfolio constraint, is responsible for the 'conservative response' of the portfolio weights to extreme values of the conditioning information, as reported in Ferson and Siegel (2001).

(B) Efficient Generalized Returns

[10] In the interest of space, we do not include the proof in this paper, details are available from the authors upon request.

Similarly, extending the proof of Proposition A.3 to include a risk-free asset, one can show that the maximum (squared) Sharpe ratio in the space R_t^G of *generalized* returns is simply given by $\lambda_*^2 = E(\, H_{t-1}^2\,) = h^2$. In other words, the (squared) maximum *unconditional* Sharpe ratio is simply the unconditional expectation of the squared *conditional* Sharpe ratio. While our results apply also to multiple risky assets, the analogous result for a single risky asset has been reported in Cochrane (1999).

It now follows that the unconditionally efficient generalized return $r_t^m \in R_t^G$ for given unconditional mean $m \in \mathbb{R}$ can be written as $r_t^m = \theta_{t-1}^0 r_f + (\,\tilde{R}_t - r_f e\,)'\theta_{t-1}$, where

$$\theta_{t-1} = \frac{w - r_f}{1 + h^2} \cdot \Sigma_{t-1}^{-1}\big(\,\mu_{t-1} - r_f e\,\big), \quad \text{with} \quad w = \frac{m - r_f}{h^2}, \tag{24}$$

and θ_{t-1}^0 is a function of H_{t-1}^2, normalized so that $E(\,\theta_{t-1}^0\,) = 1$. Note that the functional form of (24) is identical to the weights obtained from classic mean-variance theory in the case without conditioning information.

5.2 Properties of Efficient Portfolio Weights

In this section we further explore the properties of the weights of unconditionally efficient conditional returns and generalized returns. We first discuss the connections with utility maximization, and show that the weights of the unconditionally efficient conditional return have properties similar to fixed-weight efficient portfolios in the case when estimation risk is taken into account. As in Section 4.3, we assume for the remainder of this section that the conditional mean is a linear function of a single conditioning instrument, $\mu_{t-1} = \mu(\,y_{t-1}\,) = \mu_0 + \beta y_{t-1}$, and that the conditional variance-covariance matrix Σ is constant.

Ferson and Siegel (2001) and Avramov and Chordia (2003) show that the unconditionally efficient weights arise from maximizing the conditional expected value of a quadratic utility function. The resulting unconditionally efficient strategy is conditionally efficient, but not vice-versa. We now explore the differences between conditionally and unconditionally efficient strategies. Consider first an investor who maximizes *conditional* quadratic utility with (conditional) risk aversion coefficient Γ_{t-1}. Standard portfolio theory implies that the optimal portfolio is of the form

$$r_t = r_f + (\tilde{R}_t - r_f e)'\theta_{t-1} \quad \text{where} \qquad \theta_{t-1} = \frac{1}{\Gamma_{t-1}} \Sigma^{-1} (\mu_{t-1} - r_f e).$$

Evidently, the conditional expected excess return of this strategy is H^2_{t-1}/Γ_{t-1}, where H^2_{t-1} is the maximum squared *conditional* Sharpe ratio as defined above. Conversely, suppose now that the investor chooses a managed strategy such as to maximize *unconditional* quadratic utility with risk aversion coefficient γ. The optimal portfolio can be shown[11] to be of the form

$$r_t = r_f + (\tilde{R}_t - r_f e)'\theta_{t-1} \quad \text{with} \qquad \theta_{t-1} = \frac{1}{\gamma} \cdot \frac{1 + h^2}{1 + H^2_{t-1}} \Sigma^{-1}(\mu_{t-1} - r_f e), \qquad (25)$$

where $h^2 = E(H^2_{t-1})$ as defined above. Hence, the unconditionally optimal strategy implied by a risk aversion coefficient γ, when viewed as a conditionally efficient strategy, corresponds to a conditional risk aversion coefficient of

$$\Gamma_{t-1} = \gamma \cdot \frac{1 + H^2_{t-1}}{1 + h^2}.$$

In other words, the unconditionally efficient strategy corresponds to a conditionally efficient strategy that is associated with a *time-varying* coefficient of risk aversion. The portfolio strategy of an unconditional utility maximizer behaves like that of a conditional utility maximizer whose risk aversion changes with the value of the conditioning information. In particular, while h^2 is a constant, it is clear that in the linear specification considered here the conditional Sharpe ratio H^2_{t-1} will tend to infinity when $y_{t-1} \longrightarrow \pm\infty$. Thus, for extreme values of the conditioning information, the unconditionally efficient strategy when viewed as a conditionally efficient one has a much higher coefficient of relative risk-aversion. Conversely, for the same level of risk aversion, the unconditionally efficient strategy becomes increasingly conservative for large values of the conditioning variable, reducing the weights on the risky assets. This is analogous to the behavior of efficient fixed-weight portfolios that incorporate

[11] This follows from (23).

estimation risk. Such strategies reduce the weight in the tangency portfolio relative to strategies that do not account for estimation risk (Brown 1976).

The behaviour of the conditional means is also interesting. For the same level of risk aversion, the conditional mean of the conditionally efficient strategy clearly tends to infinity as $y_{t-1} \longrightarrow \pm\infty$. In contrast, the excess mean of the unconditionally efficient strategy tends to $(1 + h^2)/\gamma$. The unconditionally efficient strategy thus resembles the conditionally efficient strategy for small values of the conditioning information, and a conditionally efficient strategy that is constrained to keep the conditional mean fixed for extreme values.

In contrast, the risky asset weights (24) of the unconditionally efficient *generalized* return are identical to those of the corresponding conditionally efficient strategy, given the same level of risk aversion. The difference in this case lies in the weight placed on the risk-free asset, which in the case of generalized returns is time-varying. Thus, the unconditionally efficient generalized return does not display the conservative response to extreme values of the conditioning variable. In fact, the weights on the risky assets in this case depend linearly on the instrument, requiring potentially extreme long and short positions.

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 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

AVRAMOV, D., and T. CHORDIA (2003): "Stock Returns are Predictable: A Firm-Level Analysis," working paper, University of Maryland.

BEKAERT, G., and J. LIU (2004): "Conditioning Information and Variance Bounds on Pricing Kernels," *Review of Financial Studies*, 17(2), 339–378.

BROWN, S. (1976): "Optimal Portfolio Choice under Uncertainty: A Bayesian Approach," PhD thesis, University of Chicago.

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.

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

FERSON, W. (2003): "Test of Multifactor Pricing Models, Volatility Bounds and Portfolio Performance," in *Handbook of the Economics of Finance*, ed. by G. Constantinides et al. 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 (2003): "Stochastic Discount Factor Bounds with Conditioning Information," *Review of Financial Studies*, 16(2), 567–595.

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.

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.

Haugen, R. (1997): *Modern Investment Theory*. Prentice Hall, NJ.

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 Mathematical Appendix

A.1 Proof of Proposition 3.2:

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 squared hypothetical Sharpe ratio,

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

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

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

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 (7). To prove the decomposition (6) of the maximum hypothetical Sharpe ratio, we re-write the first-order condition as,

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

This completes the proof of Proposition 3.2. □

A.2 Proof of Lemma 4.1:

In Propositions A.1 and A.2 below we characterize the portfolio weights for the conditional returns r_t^* and z_t^*. From this follows,

$$E(r_t^*) = E(B_{t-1}/A_{t-1}) = \alpha_1, \quad \text{and} \quad E(z_t^*) = E(D_{t-1} - B_{t-1}^2/A_{t-1}) = \alpha_3.$$

The desired result then follows from the Hansen and Richard (1987) representation of the efficient frontier. □

Proposition A.1 *The conditional return r_t^* with minimum second moment is given by,*

$$r_t^* = \tilde{R}_t' \theta_{t-1} \quad \textit{with} \quad \theta_{t-1} = \frac{1}{A_{t-1}} \Lambda_{t-1}^{-1} e$$

PROOF: Throughout the proof, we will omit the time subscript to simplify notation. By Lemma 3.3 of Hansen and Richard (1987), the second moment minimization problem for conditional returns can be solved conditionally. 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,

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

This completes the proof of Proposition A.1. □

Proposition A.2 *The projection z_t^* of* 1 *onto the space of conditional excess returns is,*

$$z_t^* = \tilde{R}_t'\theta_{t-1} \quad \textit{with} \quad \theta_{t-1} = \Lambda_{t-1}^{-1}\big(\mu_{t-1} - \frac{B_{t-1}}{A_{t-1}}e\big)$$

PROOF: Throughout the proof, we will omit the time subscript. We use the fact that z^* 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

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

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,

$$0 = E_{t-1}\big(rz^* - (r - r^*)\big) = \frac{\theta'\Lambda\phi}{e'\phi} - \mu'\big(\frac{\phi}{e'\phi} - \frac{1}{A}\Lambda^{-1}e\big)$$

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

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

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

This completes the proof of Proposition A.2. □

A.3 Proof of Theorem 4.3:

In Propositions A.3 and A.4 below we characterize the portfolio weights for the generalized returns r_t^* and z_t^*. From this, we obtain,

$$E(\, r_t^* \,) = b/a = \hat{\alpha}_1, \quad \text{and} \quad E(\, z_t^* \,) = d - b^2/a = \hat{\alpha}_3.$$

The desired result then follows from the Hansen and Richard (1987) representation of the efficient frontier. □

Proposition A.3 *The generalized return r_t^* with minimum second moment is given by,*

$$r_t^* = \tilde{R}_t' \theta_{t-1} \quad \textit{with} \quad \theta_{t-1} = \frac{1}{a} \Lambda_{t-1}^{-1} e$$

PROOF: Throughout the proof, we will omit the time subscript. We use calculus of variation. Suppose θ is a solution, and ϕ is an arbitrary vector of (managed) weights. Define,

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

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

$$\left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} E\big(\, \theta_\varepsilon' \Lambda \theta_\varepsilon \,\big) = 0$$

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

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

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

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

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

This completes the proof of Proposition A.3. □

Proposition A.4 *The projection z_t^* of* 1 *onto the space of generalized excess returns is,*

$$r_t^* = \tilde{R}_t' \theta_{t-1} \quad \textit{with} \quad \theta_{t-1} = \Lambda_{t-1}^{-1} \left(\mu_{t-1} - \frac{b}{a} e \right)$$

Proof: Throughout the proof, we will omit the time subscript. For unconditional returns, z^* is the Riesz representation of the unconditional expectation. Hence,

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

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

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

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

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

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

This completes the proof of Proposition A.4. □

A.4 Matrix Identity used in Section 5:

Suppose $\Sigma \in \mathbb{R}^{n \times n}$ is symmetric and $\mu \in \mathbb{R}^n$. If both Σ and $(\Sigma - \mu\mu')$ are invertible, then

$$\left(\Sigma + \mu\mu' \right)^{-1} = \Sigma^{-1} - \frac{\Sigma^{-1} \mu \mu' \Sigma^{-1}}{1 + \mu' \Sigma^{-1} \mu}$$

This relation is trivial to verify, we do not provide a proof here.