THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Portfolio efficiency and discount factor bounds with conditional information

Citation for published version:

Abhyankar, A, Basu, D & Stremem, A 2007, 'Portfolio efficiency and discount factor bounds with conditional information: An Empirical Study', *Journal of Banking and Finance*, vol. 31, pp. 419-437.
<https://doi.org/10.2139/ssrn.423962>

Digital Object Identifier (DOI):

[10.2139/ssrn.423962](https://doi.org/10.2139/ssrn.423962)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Journal of Banking and Finance

Publisher Rights Statement:

Abhyankar, A., Basu, D., & Stremem, A. (2007). Portfolio efficiency and discount factor bounds with conditional information: An Empirical Study. *Journal of Banking & Finance*, 31, 419-437

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

PORTFOLIO EFFICIENCY AND DISCOUNT FACTOR BOUNDS
WITH CONDITIONING INFORMATION: AN EMPIRICAL STUDY¹

ABHAY ABHYANKAR

University of Edinburgh

DEVRAJ BASU

ALEXANDER STREMME

Warwick Business School, University of Warwick

NOVEMBER 2005

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 Wayne Ferson and for many constructive suggestions that have helped shape this paper. We also thank participants and the discussant at the 2003 EFA meetings in Glasgow for useful comments, and the anonymous referees for very helpful suggestions. All remaining errors are ours.

PORTFOLIO EFFICIENCY AND DISCOUNT FACTOR BOUNDS WITH CONDITIONING INFORMATION: AN EMPIRICAL STUDY

NOVEMBER 2005

Abstract

Stochastic discount factor bounds provide a useful diagnostic tool for testing asset pricing models by specifying a lower bound on the variance of any admissible discount factor. In this paper, we provide a unified derivation of such bounds in the presence of conditioning information, which allows us to compare their theoretical and empirical properties. We find that, while the location of the ‘unconditionally efficient (UE)’ bounds of Ferson and Siegel (2003) is statistically indistinguishable from the (theoretically) optimal bounds of Gallant, Hansen, and Tauchen (1990) (GHT), the former exhibit better sampling properties. We demonstrate that the difference in sampling variability of the UE and GHT bounds is due to the different behavior of the efficient return weights underlying their construction.

JEL CLASSIFICATION: G11, G12

1 Introduction

Stochastic discount factor (SDF) bounds define the feasible region in the mean-variance plane by providing a lower bound on the variance of an admissible SDF, as a function of its mean. Such bounds have found wide applications in several areas of asset pricing. The optimal use of conditioning information to refine these bounds has been the focus of several recent studies. This procedure incorporates time-variation in the conditional mean and variance of returns and leads to more stringent tests of asset pricing models.

The contribution of this paper is two-fold: first, we develop a unified framework for the construction of stochastic discount factor (SDF) bounds in the presence of conditioning information, and establish a one-to-one correspondence between these bounds and the unconditionally efficient frontier of dynamically managed portfolios. We thus provide two distinct methods of implementing such bounds, either directly from the conditional moments of asset returns, or as the unconditional variance of a dynamic portfolio. While both methods are theoretically equivalent, they can have different empirical properties, with the latter being particularly useful for the implementation of out-of-sample tests. Our second contribution is a comprehensive comparative analysis of the statistical properties of different specifications and implementations of discount factor bounds, using both theoretical arguments as well as an extensive empirical analysis.

Hansen and Jagannathan (1991) show that there is a one-to-one correspondence between variance bounds for pricing kernels and the efficient frontier, when there is no conditioning information. However, extending this correspondence to the case with conditioning information is not straight-forward. Gallant, Hansen, and Tauchen (1990) (GHT) and, in a slightly more restrictive setting, also Ferson and Siegel (2003), derive variance bounds for pricing kernels in the presence of conditioning information. On the other hand, Hansen and Richard (1987) and later Ferson and Siegel (2001) study unconditional mean-variance efficiency in this setting. Our unified approach allows us to extend the Hansen and Jagannathan (1991) correspondence between discount factor bounds and mean-variance efficiency to the case

with conditioning information in the most general setting.

We provide a unified derivation of the different sets of discount factor bounds, which allows us to compare their sampling properties, both from the theoretical and empirical viewpoint. Our expressions enable us to characterize the dynamically managed portfolios that attain the sharpest possible discount factor bounds for a given set of assets and conditioning variables. Moreover, our formulation of the weights of these portfolios facilitates the analysis of their behavior in response to changes in conditioning information.

We find that, while the *location* of Ferson and Siegel’s (2003) ‘unconditionally efficient (UE)’ bounds is statistically indistinguishable from that of the (theoretically optimal) bounds of Gallant, Hansen, and Tauchen (1990), the former exhibit lower sampling variability. Thus, tests based on the UE bounds are likely to have more power than those based on the GHT bounds. Our unified derivation allows us to demonstrate that the difference in sampling variability between the two sets of bounds is due to the different behavior of the portfolio weights underlying their construction.

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 that attains the discount factor bound. We show that this construction can be linked to the efficient frontier generated by a set of ‘generalized’ returns. These are pay-offs whose price is normalized to one on *average*. Although these pay-offs cannot be attained by forming portfolios of the traded assets, they must be priced correctly by any admissible discount factor. In this sense, they may be regarded as returns on ‘pseudo’ portfolios. We explicitly characterize the efficient frontier in the space of such generalized returns, thus providing an alternative derivation for the GHT bounds as well as the optimally scaled bounds of Bekaert and Liu (2004). We find that the weights of these efficient generalized returns are very similar to those from standard mean-variance analysis. Our approach facilitates the direct comparison of the GHT bounds with the ‘unconditionally efficient (UE)’ bounds of Ferson and Siegel (2003).

Our work is also related to Ferson and Siegel (2001), who study the properties of unconditionally mean-variance efficient portfolios in the presence of conditioning information. They demonstrate in the case of a single risky and risk-free asset that these portfolio weights are not monotonic in the realization of conditioning information, but exhibit a ‘conservative’ response to extreme signals. We provide a theoretical explanation for this behavior even for multiple risky assets. We also show that such a phenomenon occurs even when there is no risk-free asset, and demonstrate how this behavior leads to the lower sampling variability of the UE bounds based on these weights. In contrast, the portfolio weights on which the GHT bounds are based require extreme long and short positions for large values of the conditioning instrument, which accounts for their greater sampling variability.

Ferson and Siegel (2003) propose a bias-correction for bounds with conditioning information. We implement this correction for both sets of bounds and find that it improves the location of the bounds and also reduces sampling variability. In addition, we conduct an out-of-sample analysis of the two sets of bounds and also study the effect of conditional heteroskedasticity and measurement error on the bounds.

The remainder of this paper is organized as follows; In Section 2, we establish our notation and give a brief overview of discount factor bounds. In Section 3, we provide a generic, portfolio-based characterization of these bounds, and in the following Section 4 we derive explicit formulas for their econometric implementation. The results of our empirical analysis are reported in Section 5, and Section 6 concludes. All mathematical proofs are given in the appendix.

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, augmented by the use of conditioning information.

2.1 Traded Assets and Managed Portfolios

Trading takes place in discrete time. For any given period beginning at time $t-1$ and ending at time t , denote by $\mathcal{G}_{t-1}$ the information set available to the investor at the beginning of the period. For notational convenience, we write $E_t(\cdot)$ for the conditional expectation with respect to $\mathcal{G}_{t-1}$.

There are n risky assets, indexed $k = 1 \dots n$. We denote the gross return (per dollar invested) of the k -th asset by r_t^k , and by $\tilde{R}_t := (r_t^1 \dots r_t^n)'$ 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 pay-offs x_t that can be written in the form, $x_t = \tilde{R}_t' \theta_{t-1}$, with $\theta_{t-1} = (\theta_{t-1}^1 \dots \theta_{t-1}^n)'$, where θ_{t-1}^k 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_{t-1}^k that are functions of the conditioning information¹. By construction, the *price* of such a pay-off is given by $e' \theta_{t-1}$, where $e = (1 \dots 1)'$ is an n -vector of ‘ones’.

2.2 Stochastic Discount Factors and Bounds

Stochastic discount factors (SDFs) are a convenient way of describing an asset pricing model. They are characterized in terms of a fundamental valuation equation.

Definition 2.1 *An admissible **stochastic discount factor** is an element m_t such that*

$$E_{t-1}(m_t \tilde{R}_t) = e. \quad (1)$$

In other words, an SDF assigns unit price to the traded asset returns. Note that (1) implies that m_t also prices all managed pay-offs (conditionally) correctly, that is $E_{t-1}(m_t x_t) = e' \theta_{t-1}$

---

¹Note that, in contrast to the fixed-weight case, the space of *managed* pay-offs is infinite-dimensional even when there is only a finite number of base assets.

for all $x_t \in X_t$. Taking unconditional expectations,

$$E(m_t x_t) = E(e' \theta_{t-1}) =: \Pi(x_t) \quad (2)$$

In other words, any SDF that prices the base assets (conditionally) correctly must necessarily be consistent with the ‘*generalized*’ pricing function $\Pi(x_t) = E(e' \theta_{t-1})$. For different choices of θ_{t-1} (and hence different $x_t \in X_t$), we thus obtain a family of testable ‘moment conditions’ that the SDF must satisfy.

A GENERIC EXPRESSION FOR DISCOUNT FACTOR BOUNDS:

While (2) 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 conditions². 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 in their most general form can be formulated as,

Lemma 2.2 *Necessary for a candidate m_t with $E(m_t) = \nu$ to be an admissible SDF is,*

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

where $R_t \subset X_t$ is any arbitrary subspace of X_t such that $\Pi(r_t) = 1$ for all $r_t \in R_t$.

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 of the above inequality can hence be interpreted as the maximum *generalized Sharpe ratio* on R_t , relative to the shadow risk-free rate $1/\nu$.

---

²This is because the space X_t on which the SDF must be tested is infinite-dimensional.

A CLASSIFICATION OF DIFFERENT SPECIFICATIONS OF BOUNDS:

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 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*,

$$R_t^0 = \{ x_t = \tilde{R}_t' \theta, \text{ where } \theta \in \mathbb{R}^n \text{ with } e' \theta = 1 \} \quad (4)$$

- (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*’,

$$R_t^C = \{ x_t = \tilde{R}_t' \theta_{t-1}, \text{ where } \theta_{t-1} \text{ is } \mathcal{G}_{t-1}\text{-measurable with } e' \theta_{t-1} \equiv 1 \} \quad (5)$$

- (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*’,

$$R_t^G = \{ x_t = \tilde{R}_t' \theta_{t-1}, \text{ where } \theta_{t-1} \text{ is } \mathcal{G}_{t-1}\text{-measurable with } E(e' \theta_{t-1}) = 1 \} \quad (6)$$

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

3 Stochastic Discount Factor Bounds

To construct the bound for a given mean $E(m_t)$ of the discount factor, we need to find the portfolio that maximizes the hypothetical 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 $\Pi \equiv 1$, including in particular the three spaces R_t^0 , R_t^C , or R_t^G defined in the preceding section.

3.1 Discount Factor Bounds and Efficient Portfolios

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 orthogonal³ to the space of *excess* (i.e. zero cost) returns $Z_t = \Pi^{-1}\{0\} \subset R_t$, and $z_t^* \in Z_t$ is a canonically chosen excess return. In other words, the unconditionally efficient frontier in R_t is spanned by r_t^* and z_t^* .

Extending this construction, we consider instead the unique return r_t^0 that is orthogonal to Z_t with respect to the *covariance* inner product⁴, 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) return⁵. In analogy with the Hansen and Richard (1987) construction, we choose $z_t^0 \in Z_t$ so that $E(z_t) = \text{cov}(z_t^0, z_t)$ for all $z_t \in Z_t$. It is easy to show that r_t^0 and z_t^0 also span the unconditionally efficient frontier. In this parametrization, the GMV r_t^0 may be regarded as a measure of *location*, while z_t^0 determines the *shape* of the frontier. While both these parameterizations are theoretically equivalent, the robustness of either with respect to estimation error is quite different, see also Section 5.1.2. We are now ready to state our main result, on which most of the empirical

---

³One can also define r_t^* as the return with minimum unconditional second moment.

⁴In the absence of a risk-free asset, the covariance functional is indeed a well-defined inner product.

⁵This follows directly from the first-order condition of the unconstrained variance minimization problem.

analysis in later sections is based:

Theorem 3.1 *The maximum $\lambda_*(\nu)$ in (3) admits a decomposition of the form,*

$$\lambda_*^2(\nu) = \lambda_0^2(\nu) + \gamma_3 \quad \text{with} \quad \lambda_0(\nu) = \frac{\gamma_1 - 1/\nu}{\gamma_2}, \quad (7)$$

Moreover, necessary for any candidate m_t with $\nu = E(m_t)$ to be an admissible SDF is,

$$\sigma^2(m_t) \geq \frac{(\gamma_1^2 + \gamma_2\gamma_3) \cdot \nu^2 - 2\gamma_1 \cdot \nu + 1}{\gamma_2}. \quad (8)$$

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

PROOF OF THEOREM 3.1: Equation (7) follows from the first-order condition of the maximization problem for $\lambda_*^2(\nu)$ (details are available from the authors upon request). Inequality (8) then follows trivially by Lemma 2.2. $\square$

Lemma 3.2 *The maximum $\lambda_*(\nu)$ in (3) is attained by the return*

$$r_t^\nu = r_t^0 + \kappa^*(\nu) \cdot z_t^0, \quad \text{with} \quad \kappa^*(\nu) = \frac{\gamma_2}{\gamma_1 - 1/\nu}. \quad (9)$$

Moreover, necessary for any candidate m_t with $E(m_t) = \nu$ to be an admissible SDF is,

$$\sigma^2(m_t) \geq \sigma^2\left(\frac{\nu}{\kappa^*(\nu)} \cdot r_t^\nu\right), \quad (10)$$

Here, γ_1 , γ_2 and γ_3 are the moments of r_t^0 and z_t^0 as defined in Theorem 3.1.

Bekaert and Liu (2004) provide an alternative derivation of the GHT bound when the first and second conditional moments are estimated correctly. The bounds are obtained as the squared Sharpe ratio of an ‘optimally scaled’ payoff, given in Equation (22) of their paper. Their derivation is closely related to ours. Specifically, Lemma 3.2 shows that the optimally scaled payoff that attains the discount factor bound is given by $(\nu/\kappa^*(\nu)) \cdot r_t^\nu$. In the case when m_t is indeed an admissible SDF, the optimally scaled payoff can in fact be identified as the unconditional projection of m_t onto the space of managed payoffs X_t , as

$$\frac{\nu}{\kappa^*(\nu)} \cdot r_t^\nu = \frac{\nu\gamma_1 - 1}{\gamma_2} \cdot r_t^0 + \nu \cdot z_t^0 = -\text{proj}(m_t | X_t)$$

When moments are correctly specified, the GHT bounds are obtained as the variance of this payoff, as in (10). 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, a property that carries over to our setting. Our analysis shows that the same holds for the UE bounds; when conditional moments are misspecified, the variance of the conditional return in (9) can be used to provide a valid lower bound on the variance of SDFs. This fact is also particularly useful in out-of-sample estimations of the bounds (the results of our out-of-sample analysis are discussed in Section 5.1.4).

4 Implementing Discount Factor Bounds

In the preceding section, we derived generic expressions for discount factor bounds in the presence of conditioning information. For these expression to be of any practical use, we need to derive explicit formulae for the returns that attain the bounds, and compute their conditional moments. We define,

$$\mu_{t-1} = E_{t-1}(\tilde{R}_t), \quad \text{and} \quad \Lambda_{t-1} = E_{t-1}(\tilde{R}_t \cdot \tilde{R}_t'). \quad (11)$$

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)⁶. Finally, we set

$$A_{t-1} = e'\Lambda_{t-1}^{-1}e, \quad B_{t-1} = \mu_{t-1}'\Lambda_{t-1}^{-1}e, \quad D_{t-1} = \mu_{t-1}'\Lambda_{t-1}^{-1}\mu_{t-1} \quad (12)$$

These are the conditional versions of the ‘efficient set’ constants a , b and d 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.

---

⁶Note 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.

4.1 Ferson and Siegel's UE Bounds:

The ‘unconditionally efficient’ (UE) bounds of Ferson and Siegel (2003) are obtained from the generic formulation in Theorem 3.1 by using the space R_t^C of conditional returns in (3). To facilitate comparison, we follow Ferson and Siegel (2003) and formulate the bounds in terms of their ‘efficient set’ constants $\alpha_1 = E(1/A_{t-1})$, $\alpha_2 = E(B_{t-1}/A_{t-1})$ and $\alpha_3 = E(D_{t-1} - B_{t-1}^2/A_{t-1})$.

Proposition 4.1 *The UE bounds for a candidate m_t with $E(m_t) = \nu$ can be written as,*

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

Moreover, the conditional return $r_t^\nu \in R_t^C$ from (9) that attains the maximum generalized Sharpe ratio in (3) and hence the UE bounds can be written as $r_t^\nu = \tilde{R}_t' \theta_{t-1}$, where

$$\theta_{t-1} = \Lambda_{t-1}^{-1} \left(\frac{1 - w(\nu)B_{t-1}}{A_{t-1}} e + w(\nu) \mu_{t-1} \right) \quad \text{and} \quad w(\nu) = \frac{\alpha_1\nu - \alpha_2}{\alpha_2\nu - (1 - \alpha_3)} \quad (14)$$

PROOF OF PROPOSITION 4.1: We show in Appendix A.1 that α_1 and α_2 are the second and first moments of $r_t^* \in R_t^C$, and $\alpha_3 = E(z_t^*)$. Expression (13) then follows from Theorem 3.1 and the fact that $\gamma_1 = \alpha_2/(1 - \alpha_3)$, $\gamma_2 = \alpha_1 - \alpha_2^2/(1 - \alpha_3)$ and $\gamma_3 = \alpha_3/(1 - \alpha_3)$. The proof of the second assertion is given in Appendix A.1. $\square$

We can identify (14) as the weights of the efficient conditional return with unconditional mean $\alpha_2 + w(\nu)\alpha_3$ (see also Ferson and Siegel 2001). This portfolio has zero-beta rate $1/\nu$. The behavior of these weights as functions of the return moments and the conditioning information determines the sampling properties of the bounds.

4.2 GHT Bounds

The GHT bounds of Gallant, Hansen, and Tauchen (1990) are obtained from the generic formulation in Theorem 3.1 by using the space R_t of generalized returns in (3). Following the

notation of Bekaert and Liu (2004), we denote by a , b , and d the unconditional expectations of the efficient set constants A_{t-1} , B_{t-1} , and D_{t-1} introduced above. In analogy with the preceding section, we furthermore define $\hat{\alpha}_1 = 1/a$, $\hat{\alpha}_2 = b/a$ and $\hat{\alpha}_3 = d - b^2/a$.

Proposition 4.2 *The GHT bounds for a candidate m_t with $E(m_t) = \nu$ can be written as,*

$$\sigma^2(m_t) \geq \frac{(\hat{\alpha}_1 \hat{\alpha}_3 + \hat{\alpha}_2^2) \cdot \nu^2 - 2\hat{\alpha}_2 \cdot \nu + (1 - \hat{\alpha}_3)}{\hat{\alpha}_1(1 - \hat{\alpha}_3) - \hat{\alpha}_2^2}, \quad (15)$$

Moreover, the generalized return $r_t^\nu \in R_t^G$ from (9) that attains the maximum Sharpe ratio in (3) and hence the GHT bounds can be written as $r_t^\nu = \tilde{R}_t' \theta_{t-1}$ with

$$\theta_{t-1} = \Lambda_{t-1}^{-1} \left(\frac{1 - w(\nu)b}{a} e + w(\nu) \mu_{t-1} \right), \quad \text{where } w(\nu) = \frac{\hat{\alpha}_1 \nu - \hat{\alpha}_2}{\hat{\alpha}_2 \nu - (1 - \hat{\alpha}_3)} \quad (16)$$

PROOF OF PROPOSITION 4.2: We show in Appendix A.2 that $\hat{\alpha}_1$ and $\hat{\alpha}_2$ are the second and first moments of $r_t^* \in R_t^G$, and $\hat{\alpha}_3 = E(z_t^*)$. Expression (15) then follows from Theorem 3.1 in the same way as in the proof of Proposition 4.1. The proof of the second assertion is given in Appendix A.2. $\square$

Note that, using the relationship between a , b , d and the $\hat{\alpha}_i$, it is easy to show that (15) can be re-arranged to give Equation (25) in Bekaert and Liu (2004). Moreover, one can show that the ‘optimally scaled’ pay-off defined in Equation (22) of their paper can be normalized to give the efficient return defined in (16) in the above proposition.

Our approach thus demonstrates that both sets (UE and GHT) bounds can be obtained in very much the same fashion. Moreover, our results show that both sets of bounds admit two different characterizations; either in terms of the efficient set constants α_i and $\hat{\alpha}_i$ respectively, or directly as the variance of the optimally managed pay-off r_t^ν . While both approaches yield the same result in population, they may have rather different properties in finite samples. Moreover, the portfolio-based implementation is particularly useful to assess the out-of-sample performance of the bounds. The difference in behavior (see also the following section) of the efficient weights in (14) and (16) is largely responsible for the different sampling properties of the UE and GHT bounds, respectively.

4.3 Properties of Efficient Portfolio Weights

In this section, we examine the behavior of the weights of the (generalized) portfolios that attain the two sets of bounds. In particular, we are interested in the response of the weights to extreme values of the conditioning instruments. While both sets of weights can be shown to converge to finite limits (see below), the *speed* of convergence is quite different. The conditional return weights (14) that generate the UE bounds converge much faster, due to the conditional normalization constant. Conversely, the weights (16) of the generalized return that attain the GHT bounds exhibit an almost linear response to values of the instrument within a reasonable range. This difference in behavior is largely responsible for the different sampling properties of the two sets of bounds. An empirical analysis of this phenomenon is provided in Section 5.1.2.

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 of the base asset return innovations does not depend on y_{t-1} (i.e. a linear regression setting). To investigate the asymptotic properties of these weights for large values of the conditioning instrument, we use the Sherman-Morrison formula (see Appendix A.3). Using this identity and the definition of the efficient set constants, it is easy to see that $\Lambda_{t-1}^{-1}\mu_{t-1}$ and hence also B_{t-1} tend to zero as $y_{t-1} \rightarrow \pm \infty$, while both $\Lambda_{t-1}^{-1}e$ and A_{t-1} converge to finite limits. Hence, for extreme values of the instrument, the weights (14) of the *conditional* return that attains the UE bounds converge to

$$\theta_{t-1} \rightarrow \frac{(\beta'\Sigma^{-1}\beta)\Sigma^{-1}e - (\beta'\Sigma^{-1}e)\Sigma^{-1}\beta}{(e'\Sigma^{-1}e)(\beta'\Sigma^{-1}\beta) - (\beta'\Sigma^{-1}e)^2} \quad \text{as } y_{t-1} \rightarrow \pm \infty. \quad (17)$$

These are in fact the asymptotic weights of the minimum second moment return r_t^* as it can be shown that $z_t^* \rightarrow 0$ as $y_{t-1} \rightarrow \pm \infty$ in the Hansen and Richard (1987) decomposition of the efficient frontier. Moreover, it is easy to see that the conditional mean of the unconditionally efficient return defined by (14) converges to $w(\nu)$ as $y_{t-1} \rightarrow \pm \infty$, 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 (16) of the *generalized* return that attains the GHT bounds converges to

$$\theta_{t-1} \longrightarrow \frac{1 - w(m)b}{a} \left[\Sigma^{-1}e - \frac{\beta' \Sigma^{-1}e}{\beta' \Sigma^{-1}\beta} \Sigma^{-1}\beta \right] \quad \text{as } y_{t-1} \longrightarrow \pm \infty. \quad (18)$$

From this analysis we see that the major difference in the speed of convergence is determined by the presence of the conditional normalization constant B_{t-1} in (14), which is missing from the corresponding generalized return weights.

5 Empirical Analysis

In this section, we report the results of our empirical analysis. We estimate two different models, a simple linear predictive regression, and a model displaying conditional heteroscedasticity following Ferson and Siegel (2003). We use monthly returns on the five Fama-French industry portfolios as base assets, and consumption-wealth ratio (CAY) as predictive variable. To analyze the sampling properties of the discount factor bounds, we conduct an extensive simulation analysis based on the estimated models.

5.1 Constant Volatility Model

We specialize the set-up of Section 2 to the case of a single instrument with a predictive regression as in Equation (1) of Ferson and Siegel (2001). Specifically, let y_{t-1} denote a (univariate) conditioning instrument, and $\mathcal{G}_{t-1} = \sigma(y_{t-1})$. For notational convenience, we set $y_{t-1}^0 = y_{t-1} - E(y_{t-1})$. Throughout this first section, we assume that the instrument only affects the conditional mean of the base asset returns;

Assumption 5.1 *Throughout this section, we assume that returns can be described as,*

$$\tilde{R}_t = \mu_0 + \beta \cdot y_{t-1}^0 + \varepsilon_t. \quad (19)$$

where the vector of residuals ε_t is independent of y_{t-1} , has zero conditional mean and constant variance-covariance matrix Σ .

In the notation of Section 4, this means $\mu_{t-1} = \mu_0 + \beta \cdot y_{t-1}^0$. In the following section, we estimate this regression and calculate the implied unconditionally efficient portfolio frontier and corresponding discount factor bounds for both conditional (UE bounds) and generalized (GHT bounds) returns.

5.1.1 Estimation Results

As base assets, we use 510 monthly returns on the five Fama-French industry portfolios⁷, observed over the period from 1959:01 to 2001:07. As conditioning instrument, we use the lagged, de-measured, monthly consumption-wealth ratio (CAY) for the same period, as constructed in Lettau and Ludvigson (2001). The summary statistics of the data are reported in Table 3. Note that in that table we also report data on the S&P 500 index which we will use later as factor in the estimation of the conditional heteroscedasticity model. The results of the predictive regression are reported in Table 1. We use the estimated $\mu_{t-1} = \hat{\mu}_0 + \hat{\beta} \cdot y_{t-1}^0$ to form the efficient set constants defined in Section 4, which allows us to construct the efficient frontier for conditional and generalized returns, and the implied discount factor bounds.

In Figure 1, we plot the efficient frontiers in the fixed-weight setting (dashed line) together with the frontiers for conditional and generalized returns. It is clear that the frontiers with conditioning information are wider than the fixed-weight frontier, which indicates that the efficient use of conditioning information indeed expands the opportunity set available to

---

⁷These data are available at <http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/>.

	$\hat{\mu}$	$\hat{\beta}$	$\hat{\Sigma}$ (Conditional VCV Matrix)					R^2
			Manuf	Utils	Shops	Money	Other	
Manuf	1.0100	0.5001	0.0020	0.0010	0.0022	0.0018	0.0018	0.0259
Utils	1.0087	0.3763	0.0010	0.0015	0.0010	0.0014	0.0010	0.0189
Shops	1.0114	0.6274	0.0022	0.0010	0.0031	0.0022	0.0020	0.0260
Money	1.0110	0.5721	0.0018	0.0014	0.0022	0.0025	0.0018	0.0269
Other	1.0089	0.5181	0.0018	0.0010	0.0020	0.0018	0.0020	0.0272
maximum $R^2 =$								0.0321

Table 1: Estimation Results

This table displays the estimated coefficients $\hat{\mu}$ and $\hat{\beta}$ of the predictive regression (19) of monthly gross returns on the five Fama-French industry portfolios on the conditioning variable CAY, as well as the conditional variance-covariance matrix $\hat{\Sigma}$ of the residuals in this regression.

The maximum R^2 is obtained by finding the convex combination of the dependent variables which maximizes the R^2 of the corresponding univariate regression.

the investor. The frontiers for conditional (bold-faced line) and generalized (light-weight line) returns are virtually indistinguishable. Interestingly, the efficient use of conditioning information does not seem to affect the location of the global minimum variance (GMV) portfolio (the standard deviation of the GMV falls from 0.0359 in the fixed-weight case to 0.0358 and 0.0355 for conditional and generalized returns, respectively). However, the maximum monthly Sharpe ratio (assuming a risk-free rate of approximately 4% annually), rises from 0.153 to 0.178 and 0.180, respectively. Note also that the base assets (shown as circles in the figure) plot well inside even the fixed-weight frontier. This property is however not specific to the choice of assets or conditioning instruments. In further experiments using other data sets (not reported here), we found the same pattern of behavior.

Figure 2 compares and contrasts the behavior of the weights (as functions of the predictive

variable) of the GMV r_t^0 and the minimum second moment return r_t^* . While the weights of the GMV are comparatively stable in both cases and converge quickly to their asymptotic values, the weights of r_t^* highlight the difference between the UE and GHT bounds: for small values of the conditioning variables, both sets of weights show very similar behavior. However, for larger values of the instruments, the UE weights converge quickly, displaying the ‘conservative response’ discussed in Section 4, while the GHT weights display an almost linear response to the signal, requiring extreme long and short positions in the corresponding portfolio. From theory we know that the GHT weights would also converge eventually, but this is for extreme values of the instrument, far beyond the range observed in the data. The asymptotic weights of the efficient conditional return, as well as the corresponding fixed-weight return, are reported in Table 2.

Note however that the range of values of CAY shown in the graph is wider than that covered by the actual time series. For the values typically observed in the data, the two sets of weights are almost identical, which explains the fact that the frontiers for conditional and generalized returns in Figure 1 are virtually indistinguishable.

5.1.2 Sampling Properties of the Bounds

We first compare the bounds with and without conditioning information. Figure 3 shows that both the UE and GHT bounds plot above the fixed weight bounds and are statistically different from them for most values of $E(m_t)$. This shows that the optimal use of conditioning information raises the discount factor bounds significantly.

Since the distribution of the discount factor bound estimator is not known explicitly, we use simulation analysis to obtain its empirical distribution. To this end, we fit an AR(1) process to the observed time series of CAY,

$$y_t^0 = \alpha y_{t-1}^0 + \eta_t$$

The parameters we obtain are $\alpha = 0.874$, with $\sigma_\eta = 0.0070$. Using this specification and

	Efficient Portfolio Weights	
	Fixed-Weight	Conditional Return (asymptotic)
Manuf	0.4536	1.9764
Utils	0.7097	1.8982
Shops	-0.0873	-1.1951
Money	-0.3224	-1.6668
Others	0.2464	-0.0374

Table 2: Asymptotic Portfolio Weights

This table reports the weights of a typical efficient portfolio in the fixed-weight case (ignoring conditioning information), and the asymptotic weights (for extreme values of the predictor variable) of the corresponding efficient conditional return.

the estimates from the predictive regression, we simulate 100,000 time series of the conditioning instrument and the base asset returns, each equal in length to the original series (510 observations). Along each series, regression (19) is estimated. For each estimation, we calculate the corresponding efficient set constants A_{t-1} , B_{t-1} and D_{t-1} , from which the discount factor bounds are computed using (13) and (15). This procedure yields 100,000 simulated estimates of the bounds, the empirical distribution of which is used to quantify the sampling variability. As a benchmark we estimate the bounds along a simulated time series of one million observations. In what follows, these latter estimates will be referred to as the ‘true’ bounds.

In order to further emphasize the different sampling properties of the bounds, we repeat the above procedure with a hypothetical predictor variable, for which the residual standard

deviation in the AR(1) specification is $\sigma_\eta = 0.0176$, instead of $\sigma_\eta = 0.0070$ as estimated⁸. For the simulation, we adjust the variance of the residual in (19), so that the unconditional variance of the base asset returns is unchanged.

Figure 4 plots the mean of the discount factor bound estimator (solid line), together with the 95% confidence interval (vertical error bars). The dashed line indicates the location of the respective “other” bound to facilitate comparison. The left hand panels (x.1) plot the UE bounds, while the GHT bounds are shown in the right hand panels (x.2). The top row of panels (1.x) correspond to the original instrument (CAY), while the bottom row (2.x) correspond to the hypothetical predictor instrument with higher variance.

Increasing the R^2 of the predictive regression shifts both sets of bounds upwards and increases their sampling variability. While the minima of the two bounds shift by similar amounts, the increase in curvature of the GHT bounds is more pronounced. The mean of a candidate SDF is likely to be near the minima of the bounds (about .997, assuming a risk-free of 4%). Figure 4 clearly indicates that in this region the two bounds are statistically indistinguishable while the UE bounds have lower sampling variability than the GHT bounds. This is further illustrated in Figure 5, which shows the empirical distribution of the estimators at $E(m_t) = .998$ (see also column (a) in Table 5).

It is evident that the difference in sampling error becomes much more pronounced as the variance of the predictor variable increases. This is because of the differing response of the efficient weights for extreme values of the signal (see also Section 4 and Figure 2).

The standard deviation for the UE bounds is consistently lower than that of the GHT bounds. However, the standard deviation of either bound is quite high. Column (a) of Table 5. reports location, sampling variability and 95% confidence intervals for $E(m_t) = 0.998$. In the case of the hypothetical instrument, the mean of the GHT bound is only 13.6% higher

---

⁸This specification leads to a total standard deviation for the predictor variable of 0.0362, which is similar to the standard deviation of the short rate.

while the 95% confidence interval is more than 21% wider for the UE bound.

Our portfolio-based approach helps us understand this differing behavior of the two sets of bounds. The moments of r_t^0 for UE and GHT are robust to sampling and measurement error. The difference in the behavior of the bounds is driven entirely by the difference in the sampling variability of z_t^0 (see Table 5).

The overall conclusion to be drawn from this analysis is that, while the UE bounds are theoretically sub-optimal, they are statistically indistinguishable from the optimal GHT bounds while having lower sampling variability. Our results and analysis extend those in Ferson and Siegel (2003) in that our portfolio-based approach allows us to directly compare the sampling properties of the GHT and UE bounds.

5.1.3 Small Sample Bias Correction

Ferson and Siegel (2003) propose a small sample correction for bounds with conditioning information. They show that incorporating the bias-correction improves both accuracy and sampling variability of the bound, particularly when the number of time series observations is small. Figure 6 plots the mean of the discount factor bound estimator with and without bias corrections for a sample of 60 observations (5 years), together with the 95 % confidence interval (vertical error bars). The dashed line indicates the location of the ‘true’ discount factor bounds. The left hand panels (x.1) plot the UE bounds, while the GHT bounds are shown in the right hand panels (x.2). The top row of panels (1.x) correspond to bounds without bias corrections while the bottom row (2.x) correspond to those with. In both cases there is a clear upward bias in the bounds estimator. In fact for a wide range of discount factor means, the true bounds falls outside the confidence intervals around the un-adjusted estimates. Incorporating the bias correction dramatically improves the accuracy and lowers the sampling variability of the bounds; the size of the confidence interval shrinks by about 12% in both cases. Note that the difference in sampling variability between the UE and GHT bounds persists after the small sample correction. All our subsequent analysis incorporates

the bias correction.

5.1.4 Out-of-Sample Analysis

To study the out-of-sample performance of the bounds estimators, we split each simulated sample of length 510 into an in-sample period of 270 observations and an out-of-sample period of 240 observations. We estimate the predictive regression in-sample and use the in-sample conditional moments to construct the weights of the returns in Propositions 4.1 and 4.2 that attain the bounds. We then estimate the unconditional moments of these returns out-of-sample using (8), to obtain the out-of-sample bounds. Figure 7 plots the in-sample and out-of-sample estimates of the discount factor bound, together with the 95% confidence interval around the in-sample estimates. The dashed lines indicates the out-of-sample estimates. The out-of-sample bounds are consistently lower (by about 10%) than the in-sample bounds, with the GHT bounds performing marginally better. However, this difference is not statistically significant for most values of $E(m_t)$. While the in-sample bounds are always statistically different from zero, we cannot reject that the hypothesis that the out-of-sample bounds are statistically different from zero, for a range of $E(m_t)$. The sampling variability of the out-of-sample bounds is of the same order of magnitude as the in-sample bounds, and is in fact lower around the minimum of the bounds. The confidence bounds are slightly narrower near the minimum and wider at the extremes. The UE bounds continue to have lower sampling variability than the GHT bounds, out-of-sample.

To further assess the out-of-sample performance of the estimation, we regress the unconditional moments of the returns obtained out-of-sample on the in-sample estimates. We find that the means are very similar, but the sampling variability of the in-sample estimates does not explain very much of the out-of-sample variability. The in-sample and out-of-sample GMVs are almost uncorrelated, while the regression of out-of-sample on in-sample z_t^0 has an R^2 of 13%.

5.1.5 Measurement Error

We now study the effects of measurement error in the predictive regression on the bounds. We work with the original sample size of 510 observations. We run four simulation experiments identical to the one described above, except that we assume that some or all the parameters are estimated without measurement error. This is done by replacing the respective estimated parameter in each simulation by the ‘true’ value that was used for data-generation. We consider four cases: (a) both parameters, μ_0 and β are estimated, (b) μ_0 estimated and β assumed known, (c) β estimated and μ_0 assumed to be known, and (d) both parameter known. Table 5 reports the results of this exercise for the first three cases. Note that removing estimation risk in μ leads to a greater reduction in sampling variability than removing estimation risk in β . This is because μ affects the mean of both the location and shape returns, r_t^0 and z_t^0 , while the GMV is largely unaffected by β . In terms of location, both UE and GHT bounds have the same accuracy when measurement error is removed. Measurement error introduces an upwards bias which is greater for the GHT bounds. The sampling variability of the GHT bounds is higher than that of the UE bounds, even when measurement error is removed. Overall, the UE bounds seem more robust to estimation risk than the GHT bounds.

5.2 The Conditional Heteroskedasticity (CH) Model

We now modify our setup to incorporate conditional heteroscedasticity (CH) following Ferson and Siegel (2003). Overall, conditional heteroscedasticity has little or no effect on the results obtained in the preceding section. The R^2 of the variance regression is less than 1.5%. Both location and variability of the bound estimator change very little, as does the comparison between the UE and GHT bounds. For low $E(m_t)$ the bounds with CH are marginally lower than the bounds obtained from the linear model, whereas for high $E(m_t)$ the opposite is true. The same pattern holds for sampling variability and error bounds. However, none of these differences are statistically significant. In particular, the CH model does not perform

any better out-of sample than the linear model.

6 Conclusion

The main contribution of this paper is to provide a detailed comparison between various stochastic discount factor bounds with conditioning information. We do this by exploiting the explicit link between the stochastic discount factor approach and portfolio efficiency in the presence of conditioning information. We find that the ‘unconditionally efficient (UE)’ bounds of Ferson and Siegel (2003) are statistically indistinguishable from the (theoretically) optimal bounds of Gallant, Hansen, and Tauchen (1990), while having smaller sampling variability. We demonstrate that the difference in sampling variability of the UE and GHT bounds is due to the different behavior of the portfolio weights underlying their construction.

References

- BEKAERT, G., and J. LIU (2004): “Conditioning Information and Variance Bounds on Pricing Kernels,” *Review of Financial Studies*, 17(2), 339–378.
- 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.
- 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 4.1:

In Lemmas A.1 and A.2 below we characterize the weights for the conditional returns r_t^* and z_t^* that span the unconditionally efficient frontier in R_t^C . From this follows,

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

The desired result (14) then follows from Lemma 3.2. $\square$

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

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

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,

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

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

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

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

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

This completes the proof of Lemma A.1. $\square$

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

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

PROOF: Throughout the proof, we will omit the time subscript. We use the fact that z_t^* is the Riesz representation of the conditional expectation on the space of excess returns. Since any excess return can be written as $z = (z + r_t^*) - r_t^* =: r - r_t^*$, this implies

$$E_{t-1} \left((r - r_t^*)(z_t^* - 1) \right) = 0 \quad \text{for all } r \in R_t^C$$

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

$$0 = E_{t-1} \left(r z_t^* - (r - r_t^*) \right) = \frac{\theta' \Lambda \phi}{e' \phi} - \mu' \left(\frac{\phi}{e' \phi} - \frac{1}{A} \Lambda^{-1} e \right)$$

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

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 Lemma A.2. □

A.2 Proof of Proposition 4.2:

In Lemmas A.3 and A.4 below we characterize the weights for the generalized returns r_t^* and z_t^* that span the unconditionally efficient frontier in R_t^G . From this, we obtain,

$$E(r_t^*) = b/a = \hat{\alpha}_2, \quad \text{and} \quad E(z_t^*) = d - b^2/a = \hat{\alpha}_3.$$

The desired result (16) then follows from Lemma 3.2. □

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

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

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,

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

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

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,

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

This completes the proof of Lemma A.3. □

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

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

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

$$E((r - r_t^*)(z_t^* - 1)) = 0 \quad \text{for all } r \in R_t^G$$

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

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

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

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 Lemma A.4. □

A.3 Sherman-Morrison formula used in Section 4.3:

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

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

Variable	Base Assets					Factor	Instr
	Manuf	Utils	Shops	Money	Other	Index	CAY
Mean	1.01001	1.00871	1.01138	1.01096	1.00886	1.00968	0.00000
Std	0.04505	0.03967	0.05646	0.05055	0.04551	0.04370	0.01448
Correlations							
Manuf	1.00000	0.57355	0.87834	0.82756	0.88197	0.98477	0.16079
Utils		1.00000	0.47703	0.70224	0.57722	0.62514	0.13743
Shops			1.00000	0.78144	0.81622	0.90795	0.16097
Money				1.00000	0.80349	0.87786	0.16394
Other					1.00000	0.92649	0.16489
Index						1.00000	0.17260
CAY							1.00000

Table 3: Summary Statistics

This table shows the summary statistics (sample mean, standard deviation, and correlations) for monthly gross returns on the five Fama-French industry portfolios, the S&P 500 index (which is used as factor in the estimation of the conditional heteroscedasticity model), as well as the consumption-wealth ratio CAY (which is used as conditioning variable), for the sample period from January 1959 to July 2001. Note that CAY was normalized to have zero mean.

Figure 1: Efficient Frontier

This figure shows the estimated efficient frontier for the fixed-weight case (dashed line), and for conditional (bold-faced line) and generalized (light-weight line) returns. The base assets are monthly returns on the five industry portfolios of Fama and French (shown as circles), and the conditioning variable is CAY.

Figure 2: Efficient Portfolio Weights

This graph shows the weights of two efficient returns as functions of conditioning information. The left-hand panels show the weights for conditional returns while the right-hand panels show the weights of generalized returns. The top row of panels shows the weights of the GMV for the two sets of returns, while the bottom row are the weights of r_t^* , the minimum second moment return. The base assets are the five Fama-French industry portfolios and the conditioning variable is CAY.

		UE			GHT		
Panel (1)		$E(r^*)$	$E((r^*)^2)$	$E(z^*)$	$E(r^*)$	$E((r^*)^2)$	$E(z^*)$
		α_2	α_1	α_3	α_2	α_1	α_3
	Mean	0.9766	0.9861	0.0315	0.9775	0.9853	0.0322
		(09397)	(09487)	(00680)	(09337)	(09426)	(00741)
	Standard Deviation	0.0129	0.0134	0.0127	0.0135	0.0139	0.0133
		(00181)	(00184)	(00178)	(00211)	(00214)	(00209)
	Lower 95% Conf.	0.9532	0.9620	0.0131	0.9512	0.9601	0.0133
		(09084)	(09167)	(00403)	(08966)	(09050)	(00423)
	Upper 95% Conf.	0.9953	1.0055	0.0546	0.9952	1.0053	0.0566
		(09679)	(09776)	(00987)	(09658)	(09753)	(01107)
Panel (2)		$E(r^0)$	$\sigma^2(r^0)$	$E(z^0)$	$E(r^0)$	$\sigma^2(r^0)$	$E(z^0)$
		γ_1	γ_2	γ_3	γ_1	γ_2	γ_3
	Mean	1.0084	0.0013	0.0327	1.0084	0.0012	0.0355
		(10082)	(00013)	(00733)	(10084)	(00011)	(00806)
	Standard Deviation	0.0016	0.0001	0.0137	0.0016	0.0001	0.0144
		(00015)	(00001)	(00206)	(00015)	(00001)	(00246)
	Lower 95% Conf.	1.0057	0.0011	0.0133	1.0058	0.0011	0.0134
		(10058)	(00011)	(00420)	(10060)	(00010)	(00442)
	Upper 95% Conf.	1.0110	0.0014	0.0577	1.0110	0.0014	0.0600
		(10107)	(00014)	(01095)	(10108)	(00012)	(01245)

Table 4: Sampling Properties of Unconditional Moments

This table compares the sampling properties of the simulated estimates of the unconditional moments of the efficient returns that span the frontier and are used to compute the bounds. Panel (1) shows the sampling properties of the unconditional moments of r_t^* and z_t^* , while Panel (2) does the same thing for the moments of r_t^0 and z_t^0 . The figures in brackets are obtained using a hypothetical predictor variable with higher variance.

Figure 3: Comparison with Fixed-Weight Bounds

This graph plots the UE (bold-faced line) and GHT (light-weight line) bounds in relation to the fixed-weight (dashed line) bounds, together with the 95% percent confidence intervals (error bars) of the *fixed-weight* bounds.

Figure 4: Comparison of UE and GHT Bounds

This graph compares the sampling variability of the GHT and UE bounds, estimated from simulated time series. The left-hand panels plot the UE bounds (solid line) and the GHT bounds (dashed line), together with the 95% confidence intervals (error bars) around the UE bounds. The right-hand panel does the same thing with the roles of UE and GHT bounds reversed. The top row of panels use CAY as the predictor variable while the bottom panels use the hypothetical conditioning variable with higher variance.

Figure 5: Empirical Distribution

This graph plots the empirical distribution of the simulated estimates of the discount factor bound at $E(m_t) = .998$. The bold-faced lines represent the UE bounds and the lighter ones the GHT bounds.

The top panel uses CAY as the predictive variable while the bottom panel uses the hypothetical conditioning variable with higher variance.

Figure 6: Small Sample Bias Correction

This graph shows the effect of the small-sample bias correction due to Ferson and Siegel (2003) on the bounds. The top row of panels shows the bounds without bias correction (solid lines), estimated on simulated samples of 60 observations only, while the bottom row shows the bias-adjusted bounds estimated from the same samples. The left-hand panel plots the UE bounds, while the right hand panel shows the GHT bounds. The dashed lines indicate the location of the ‘true’ bounds, obtained from a simulated sample of 1 million observations.

Figure 7: Out-of-Sample Analysis

This figure plots the in-sample and out-of-sample estimated of the UE and GHT bounds. The left-hand panels plot the simulated in-sample estimates of the UE bounds (solid lines), together with their 95% confidence intervals (error bars), while the GHT bounds are shown in the right-hand panels. The dashed lines indicate the location of the corresponding out-of-sample estimates. The top row of panels use CAY as the conditioning variable while the bottom panels use the hypothetical variable with higher variance.

	UE			GHT		
	(a)	(b)	(c)	(a)	(b)	(c)
True Value	0.0477 (00890)	0.0486	0.0474	0.0489 (01010)	0.0499	0.0486
Simulation Mean	0.0554 (00946)	0.0457	0.0481	0.0574 (01088)	0.0475	0.0495
Standard Deviation	0.0207 (00246)	0.0092	0.0187	0.0216 (00299)	0.0102	0.0193
Lower 95% Conf.	0.0248 (00568)	0.0333	0.0209	0.0255 (00634)	0.0342	0.0215
Upper 95% Conf.	0.0924 (01374)	0.0629	0.0820	0.0961 (01613)	0.0666	0.0844

Table 5: Effect of Measurement Error

This table reports the simulation results for the (bias adjusted) UE and GHT bounds at $E(m_t) = 0.998$. Column (a) shows the benchmark results from the simulated estimation. Column (b) reports the results in the case where the unconditional mean μ_0 is assumed to be estimated without error, and column (c) reports the corresponding values assuming that β is correctly estimated. The ‘true values’ are obtained from an estimation along a simulated time series of 1 million observations. The figures in brackets indicate the corresponding values for the hypothetical conditioning variable with higher variance.